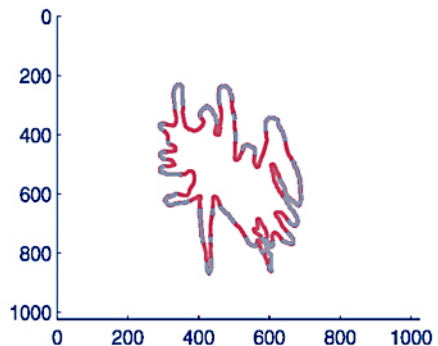
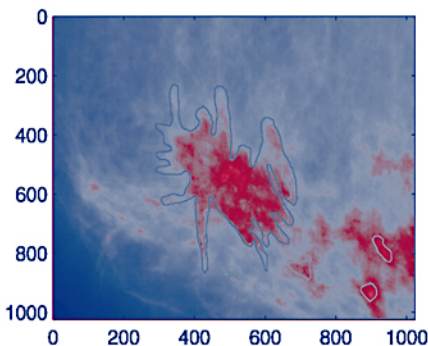
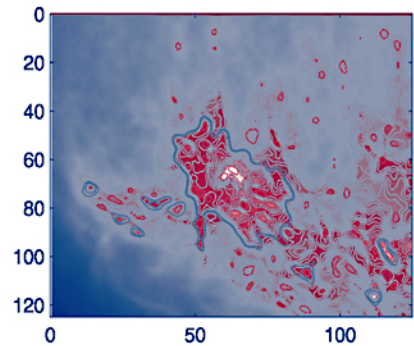
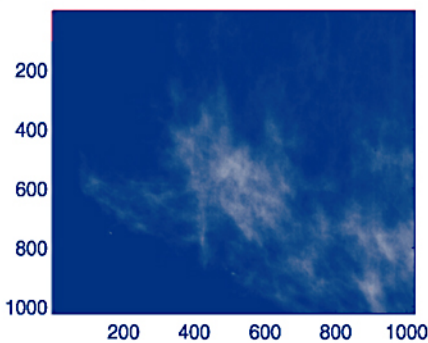


Biomedical Image Analysis

Rangaraj M. Rangayyan



The **BIOMEDICAL ENGINEERING** Series
Series Editor Michael R. Neuman

Biomedical Image Analysis

Biomedical Engineering Series

Edited by Michael R. Neuman

Published Titles

Electromagnetic Analysis and Design in Magnetic Resonance Imaging, Jianming Jin

Endogenous and Exogenous Regulation and Control of Physiological Systems, Robert B. Northrop

Artificial Neural Networks in Cancer Diagnosis, Prognosis, and Treatment, Raouf N.G. Naguib and Gajanan V. Sherbet

Medical Image Registration, Joseph V. Hajnal, Derek Hill, and David J. Hawkes

Introduction to Dynamic Modeling of Neuro-Sensory Systems, Robert B. Northrop

Noninvasive Instrumentation and Measurement in Medical Diagnosis, Robert B. Northrop

Handbook of Neuroprosthetic Methods, Warren E. Finn and Peter G. LoPresti

Signals and Systems Analysis in Biomedical Engineering, Robert B. Northrop

Angiography and Plaque Imaging: Advanced Segmentation Techniques, Jasjit S. Suri and Swamy Laxminarayan

Analysis and Application of Analog Electronic Circuits to Biomedical Instrumentation, Robert B. Northrop

Biomedical Image Analysis, Rangaraj M. Rangayyan

The **BIOMEDICAL ENGINEERING** Series
Series Editor Michael R. Neuman

Biomedical Image Analysis

Rangaraj M. Rangayyan

*University of Calgary
Calgary, Alberta, Canada*



CRC PRESS

Boca Raton London New York Washington, D.C.

Library of Congress Cataloging-in-Publication Data

Catalog record is available from the Library of Congress

This book contains information obtained from authentic and highly regarded sources. Reprinted material is quoted with permission, and sources are indicated. A wide variety of references are listed. Reasonable efforts have been made to publish reliable data and information, but the author and the publisher cannot assume responsibility for the validity of all materials or for the consequences of their use.

Neither this book nor any part may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, microfilming, and recording, or by any information storage or retrieval system, without prior permission in writing from the publisher.

The consent of CRC Press LLC does not extend to copying for general distribution, for promotion, for creating new works, or for resale. Specific permission must be obtained in writing from CRC Press LLC for such copying.

Direct all inquiries to CRC Press LLC, 2000 N.W. Corporate Blvd., Boca Raton, Florida 33431.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation, without intent to infringe.

Visit the CRC Press Web site at www.crcpress.com

© 2005 by CRC Press LLC

No claim to original U.S. Government works
International Standard Book Number 0-8493-9695-6
Printed in the United States of America 1 2 3 4 5 6 7 8 9 0
Printed on acid-free paper

*To
my wife Mayura,
my daughter Vidya,
and my son Adarsh...
for etching in my mind
the most beautiful images
that I will treasure forever!*

Preface

Background and Motivation

The science of medical imaging owes much of its existence to the discovery of X rays by W.C. Röntgen over 100 years ago, in 1895. However, it was the development of practical computed tomography scanners in the early 1970s by G. Hounsfield and others that brought computers into medical imaging and clinical practice. Since then, computers have become integral components of modern medical imaging systems and hospitals, performing a variety of tasks from data acquisition and image generation to image display and analysis.

With the widespread acceptance of computed tomography came an implicit invitation to apply computers and computing to a host of other medical imaging situations. As new imaging modalities were developed, the need for computing in image generation, manipulation, display, and analysis grew by many folds. Computers are now found in virtually every medical imaging system, including radiography, ultrasound, nuclear medicine, and magnetic resonance imaging systems. The strengths of computer applications in medical imaging have been recognized to such an extent that radiology departments in many hospitals are changing over to “totally digital” departments, using computers for image archival and communication as well. The humble X-ray film that launched the field of radiology may soon vanish, thereby contributing to better management of the environment.

The increase in the number of modalities of medical imaging and in their practical use has been accompanied by an almost natural increase in the scope and complexity of the associated problems, requiring further advanced techniques for their solution. For example, physiological imaging with radioisotopes in nuclear medicine imaging comes with a host of problems such as noise due to scatter, effects of attenuation along the path of propagation of the gamma rays through the body, and severe blurring due to the collimators used. Radiation dose concerns limit the strength and amount of the isotopes that may be used, contributing to further reduction in image quality. Along with the increase in the acceptance of mammography as a screening tool has come the need to efficiently process such images using computer vision techniques. The use of high-resolution imaging devices for digital mammography and digital radiography, and the widespread adoption of picture archival and

communication systems, have created the need for higher levels of lossless data compression. The use of multiple modalities of medical imaging for improved diagnosis of a particular type of disease or disorder has raised the need to combine diverse images of the same organ, or the results thereof, into a readily comprehensible visual display.

The major strength in the application of computers to medical imaging lies in the potential use of image processing and computer vision techniques for quantitative or objective analysis. (See the July 1972 and May 1979 issues of the *Proceedings of the IEEE* for historical reviews and articles on digital image processing.) Medical images are primarily visual in nature; however, visual analysis of images by human observers is usually accompanied by limitations associated with interpersonal variations, errors due to fatigue, errors due to the low rate of incidence of a certain sign of abnormality in a screening application, environmental distractions, etc. The interpretation of an image by an expert bears the weight of the experience and expertise of the analyst; however, such analysis is almost always subjective. Computer analysis of image features, if performed with the appropriate logic, has the potential to add objective strength to the interpretation of the expert. It thus becomes possible to improve the diagnostic confidence and accuracy of even an expert with many years of experience.

Developing an algorithm for medical image analysis, however, is not an easy task; quite often, it might not even be a straightforward process. The engineer or computer analyst is often bewildered by the variability of features in biomedical signals, images, and systems that is far higher than that encountered in physical systems or observations. Benign diseases often mimic the features of malignant diseases; malignancies may exhibit characteristic patterns, which, however, are not always guaranteed to appear. Handling all of the possibilities and the degrees of freedom in a biomedical system is a major challenge in most applications. Techniques proven to work well with a certain system or set of images may not work in another seemingly similar situation.

The Problem-solving Approach

The approach I have taken in presenting the material in this book is primarily that of problem solving. Engineers are often said to be (with admiration, I believe) problem solvers. However, the development of a problem statement and gaining of a good understanding of the problem could require a significant amount of preparatory work. I have selected a logical series of problems, from the many I have encountered in my research work, for presentation in this book. Each chapter deals with a certain type of problem with biomedical images. Each chapter begins with a statement of the problem, and includes

illustrations of the problem with real-life images. Image processing or analysis techniques are presented, starting with relatively simple “textbook methods”, followed by more sophisticated methods directed at specific problems. Each chapter concludes with applications to significant and practical problems. The book is illustrated copiously, in due consideration of the visual nature of the subject matter.

The methods presented in the book are at a fairly high level of technical and mathematical sophistication. A good background in one-dimensional signal and system analysis [1, 2, 3] is very much required in order to follow the procedures and analyses. Familiarity with the theory of linear systems, signals, and transforms such as the Laplace and Fourier, in both continuous and discrete versions, will be assumed. We shall only briefly study a few representative medical imaging techniques. We will study in more detail the problems present with medical images after they have been acquired, and concentrate on how to solve the problems. Some preparatory reading on medical imaging equipment and techniques [3, 4, 5, 6] may be useful, but not always essential.

The Intended Audience

The book is primarily directed at engineering students in their final year of undergraduate studies or in their (post-)graduate studies. Electrical and Computer Engineering students with a rich background in signals and systems [1, 2, 3] will be well prepared for the material in the book. Students in other engineering disciplines or in computer science, physics, mathematics, or geophysics should also be able to appreciate the material in this book. A course on digital signal processing or digital filters [7] would form a useful link, but a capable student without this topic may not face much difficulty. Additional study of a book on digital image processing [8, 9, 10, 11, 12, 13] could assist in developing a good understanding of general image processing methods, but is not required.

Practicing engineers, computer scientists, information technologists, medical physicists, and data-processing specialists working in diverse areas such as telecommunications, seismic and geophysical applications, biomedical applications, hospital information systems, remote sensing, mapping, and geomatics may find this book useful in their quest to learn advanced techniques for image analysis. They could draw inspiration from other applications of data processing or analysis, and satisfy their curiosity regarding computer applications in medicine and computer-aided medical diagnosis.

Teaching and Learning Plans

An introduction to the nature of biomedical images is provided in Chapter 1. The easy-to-read material in this chapter gives a general overview of the imaging techniques that are commonly used to acquire biomedical images; for detailed treatment of medical imaging, refer to Macovski [5], Robb [14], Barrett and Swindell [3], Huda and Slone [6], and Cho et al. [4]. A good understanding of the basics of image data acquisition procedures is essential in order to develop appropriate methods for further treatment of the images.

Several concepts related to image quality and information content are described in Chapter 2, along with the related basics of image processing such as the Fourier transform and the modulation transfer function. The notions, techniques, and measures introduced in this chapter are extensively used in the book and in the field of biomedical image analysis; a clear understanding of this material is an important prerequisite to further study of the subject.

Most of the images acquired in practice suffer loss of quality due to artifacts and practical limitations. Several methods for the characterization and removal of artifacts and noise are presented in Chapter 3. Preprocessing of images to remove artifacts without causing distortion or loss of the desired information is an important step in the analysis of biomedical images.

Imaging and image processing techniques aimed toward the improvement of the general quality or the desired features in images are described in Chapter 4. Methods for contrast enhancement and improvement of the visibility of the details of interest are presented with illustrative examples.

The important task of detecting regions of interest is the subject of Chapter 5, the largest chapter in the book. Several approaches for the segmentation and extraction of parts of images are described, along with methods to improve initial approximations or results.

Objective analysis of biomedical images requires the extraction of numerical features that characterize the most significant properties of the regions of interest. Methods to characterize shape, texture, and oriented patterns are described in Chapters 6, 7, and 8, respectively. Specific features are required for each application, and the features that have been found to be useful in one application may not suit a new application under investigation. Regardless, a broad understanding of this subject area is essential in order to possess the arsenal of feature extraction techniques that is required when attacking a new problem.

The material in the book through Chapter 8 provides resources that are more than adequate for a one-semester course with 40 to 50 hours of lectures. Some of the advanced and specialized topics in these chapters may be omitted, depending upon the methods and pace of presentation, as well as the level of comprehension of the students.

The specialized topic of image reconstruction from projections is dealt with in Chapter 9. The mathematical details related to the derivation of tomographic images are presented, along with examples of application. This chapter may be skipped in an introductory course, but included in an advanced course.

Chapter 10 contains descriptions of methods for the restoration of images with known models of image degradation. The advanced material in this chapter may be omitted in an introductory course, but forms an important subject area for those who wish to explore the subject to its full depth.

The subject of image data compression and coding is treated in detail in Chapter 11. With due regard to the importance of quality and fidelity in the treatment of health-related information, the focus of the chapter is on lossless compression. This subject may also be considered to be an advanced topic of specialized interest, and limited to an advanced course.

Finally, the most important and significant tasks in biomedical image analysis — pattern analysis, pattern classification, and diagnostic decision — are described in Chapter 12. The mathematical details of pattern classification techniques are presented, along with procedures for their incorporation in medical diagnosis and clinical assessment. Since this subject forms the culmination of biomedical image analysis, it is recommended that parts of this chapter be included even in an introductory course.

The book includes adequate material for two one-semester courses or a full-year course on biomedical image analysis. The subject area is still a matter of research and development: instructors should endeavor to augment their courses with material selected from the latest developments published in advanced journals such as the *IEEE Transactions on Medical Imaging* as well as the proceedings of the SPIE series of conferences on medical imaging. The topics of biometrics, multimodal imaging, multisensor fusion, image-guided therapy and surgery, and advanced visualization, which are not dealt with in this book, may also be added if desired.

Each chapter includes a number of study questions and problems to facilitate preparation for tests and examinations. Several laboratory exercises are also provided at the end of each chapter, which could be used to formulate hands-on exercises with real-life and/or synthetic images. Selected data files related to some of the problems and exercises at the end of each chapter are available at the site

www.enel.ucalgary.ca/People/Ranga/enel697

It is strongly recommended that the first one or two laboratory sessions in the course be visits to a local hospital, health sciences center, or clinical laboratory to view biomedical image acquisition and analysis in a practical (clinical) setting. Images acquired from local sources (with the permissions and approvals required) could form interesting and motivating material for laboratory exercises, and should be used to supplement the data files provided. A few invited lectures and workshops by physiologists, radiologists,

pathologists, and other medical professionals should also be included in the course so as to provide the students with a nonengineering perspective on the subject.

Practical experience with real-life images is a key element in understanding and appreciating biomedical image analysis. This aspect could be difficult and frustrating at times, but provides professional satisfaction and educational fun!

It is my humble hope that this book will assist students and researchers who seek to enrich their lives and those of others with the wonderful powers of biomedical image analysis. Electrical and Computer Engineering is indeed a great field in the service of humanity.

Rangaraj Mandayam Rangayyan
Calgary, Alberta, Canada
November, 2004

About the Author

Rangaraj (Raj) Mandayam Rangayyan was born in Mysore, Karnataka, India, on July 21, 1955. He received the Bachelor of Engineering degree in Electronics and Communication in 1976 from the University of Mysore at the People's Education Society College of Engineering, Mandya, Karnataka, India, and the Ph.D. degree in Electrical Engineering from the Indian Institute of Science, Bangalore, Karnataka, India, in 1980. He was with the University of Manitoba, Winnipeg, Manitoba, Canada, from 1981 to 1984. He joined the University of Calgary, Calgary, Alberta, Canada, in 1984.

He is, at present, a Professor with the Department of Electrical and Computer Engineering (and an Adjunct Professor of Surgery and Radiology) at the University of Calgary. His research interests are in the areas of digital signal and image processing, biomedical signal analysis, medical imaging and image analysis, and computer vision. His research projects have addressed topics such as mammographic image enhancement and analysis for computer-aided diagnosis of breast cancer; region-based image processing; knee-joint vibration signal analysis for noninvasive diagnosis of articular cartilage pathology; directional analysis of collagen fibers and blood vessels in ligaments; restoration of nuclear medicine images; analysis of textured images by cepstral filtering and sonification; and several other applications of biomedical signal and image analysis.

He has lectured extensively in many countries, including India, Canada, United States, Brazil, Argentina, Uruguay, Chile, United Kingdom, The Netherlands, France, Spain, Italy, Finland, Russia, Romania, Egypt, Malaysia, Singapore, Thailand, Hong Kong, China, and Japan. He has collaborated with many research groups in Brazil, Spain, France, and Romania.

He was an Associate Editor of the *IEEE Transactions on Biomedical Engineering* from 1989 to 1996; the Program Chair and Editor of the Proceedings of the IEEE Western Canada Exhibition and Conference on "Telecommunication for Health Care: Telemetry, Teleradiology, and Telemedicine", July 1990, Calgary, Alberta, Canada; the Canadian Regional Representative to the Administrative Committee of the IEEE Engineering in Medicine and Biology Society (EMBS), 1990 to 1993; a Member of the Scientific Program Committee and Editorial Board, International Symposium on Computerized Tomography, Novosibirsk, Siberia, Russia, August 1993; the Program Chair and Co-editor of the *Proceedings of the 15th Annual International Conference of the IEEE EMBS*, October 1993, San Diego, CA; and Program Co-chair,

20th Annual International Conference of the IEEE EMBS, Hong Kong, October 1998.

His research work was recognized with the 1997 and 2001 Research Excellence Awards of the Department of Electrical and Computer Engineering, the 1997 Research Award of the Faculty of Engineering, and by appointment as a “University Professor” in 2003, at the University of Calgary. He was awarded the Killam Resident Fellowship in 2002 by the University of Calgary in support of writing this book. He was recognized by the IEEE with the award of the Third Millennium Medal in 2000, and was elected as a Fellow of the IEEE in 2001, Fellow of the Engineering Institute of Canada in 2002, Fellow of the American Institute for Medical and Biological Engineering in 2003, and Fellow of SPIE: the International Society for Optical Engineering in 2003.




Photo by Trudie Lee.

Acknowledgments

Writing this book on the multifaceted subject of biomedical image analysis has been challenging, yet yielding more knowledge; tiring, yet stimulating the thirst to understand and appreciate more of the subject matter; and difficult, yet satisfying when a part was brought to a certain stage of completion.

A number of very important people have shaped me and my educational background. My mother, Srimati Padma Srinivasan Rangayyan, and my father, Sri Srinivasan Mandayam Rangayyan, encouraged me to keep striving to gain higher levels of education and to set and achieve higher goals all the time. I have been very fortunate to have been taught and guided by a number of dedicated teachers, the most important of them being Professor Ivaturi Surya Narayana Murthy, my Ph.D. supervisor, who introduced me to the topic of biomedical signal analysis at the Indian Institute of Science, Bangalore, Karnataka, India. I offer my humble prayers, respect, and admiration to their spirits.

My basic education was imparted by many influential teachers at Saint Joseph's Convent, Saint Joseph's Indian High School, and Saint Joseph's College in Mandya and Bangalore, Karnataka, India. My engineering education was provided by the People's Education Society College of Engineering, Mandya, affiliated with the University of Mysore. I was initiated into research in biomedical engineering at the Indian Institute of Science — India's premier research institute and one of the very highly acclaimed research institutions in the world. I express my gratitude to all of my teachers.

My postdoctoral training with Richard Gordon at the University of Manitoba, Winnipeg, Manitoba, Canada, made a major contribution to my comprehension of the field of biomedical imaging and image analysis; I express my sincere gratitude to him. My association with clinical researchers and practitioners at the University of Calgary and the University of Manitoba has been invaluable in furthering my understanding of the subject matter of this book. I express my deep gratitude to Cyril Basil Frank, Gordon Douglas Bell, Joseph Edward Leo Desautels, Leszek Hahn, and Reinhard Kloiber of the University of Calgary.

My understanding and appreciation of the subject of biomedical signal and image analysis has been boosted by the collaborative research and studies performed with my many graduate students, postdoctoral fellows, research associates, and colleagues. I place on record my gratitude to Fábio José Ayres,

Sridhar Krishnan, Naga Ravindra Mudigonda, Margaret Hilary Alto, Hanford John Deglint, Thanh Minh Nguyen, Ricardo José Ferrari, Liang Shen, Roseli de Deus Lopes, Antonio César Germano Martins, Marcelo Knörich Zuffo, Begoña Acha Piñero, Carmen Serrano Gotarredona, Laura Roa, Annie France Frère, Graham Stewart Boag, Vicente Odone Filho, Marcelo Valente, Sílvia Delgado Olabarriaga, Christian Roux, Basel Solaiman, Olivier Menut, Denise Guliato, Fabrício Adorno, Mário Ribeiro, Mihai Ciuc, Vasile Buzuloiu, Titus Zaharia, Constantin Vertan, Margaret Sarah Rose, Salahuddin Elkadiki, Kevin Eng, Nema Mohamed El-Faramawy, Arup Das, Farshad Faghieh, William Alexander Rolston, Yiping Shen, Zahra Marjan Kazem Moussavi, Joseph Provine, Hieu Ngoc Nguyen, Djamel Boulfelfel, Tamer Farouk Rabie, Katherine Olivia Lady, Yuanting Zhang, Zhi-Qiang Liu, Raman Bhalachandra Paranjape, Joseph André Rodrigue Blais, Robert Charles Bray, Gopinath Ramaswamaiah Kuduvalli, Sanjeev Tavathia, William Mark Morrow, Timothy Chi Hung Hon, Subhasis Chaudhuri, Paul Soble, Kirby Jaman, Atam Prakash Dhawan, and Richard Joseph Lehner. In particular, I thank Liang, Naga, Ricardo, Gopi, Djamel, Hilary, Tamer, Antonio, Bill Rolston, Bill Morrow, and Joseph for permitting me to use significant portions of their theses; Naga for producing the cover illustration; and Fábio, Hilary, Liang, Mihai, Gopi, Joseph, Ricardo, and Hanford for careful proofreading of drafts of the book. Sections of the book were reviewed by Cyril Basil Frank, Joseph Edward Leo Desautels, Leszek Hahn, Richard Frayne, Norm Bartley, Randy Hoang Vu, Ilya Kamenetsky, Vijay Devabhaktuni, and Sanjay Srinivasan. I express my gratitude to them for their comments and advice. I thank Leonard Bruton and Abu Sesay for discussions on some of the topics described in the book. I also thank the students of my course ENEL 697 Digital Image Processing over the past several years for their comments and feedback.

The book has benefited significantly from illustrations and text provided by a number of researchers worldwide, as identified in the references and permissions cited. I thank them all for enriching the book with their gifts of knowledge and kindness. Some of the test images used in the book were obtained from the Center for Image Processing Research, Rensselaer Polytechnic Institute, Troy, NY, www.ipl.rpi.edu; I thank them for the resource.

The research projects that have provided me with the background and experience essential in order to write the material in this book have been supported by many agencies. I thank the Natural Sciences and Engineering Research Council of Canada, the Alberta Heritage Foundation for Medical Research, the Alberta Breast Cancer Foundation, Control Data Corporation, Kids Cancer Care Foundation of Alberta, the University of Calgary, the University of Manitoba, and the Indian Institute of Science for supporting my research projects.

I thank the Killam Trusts and the University of Calgary for awarding me a Killam Resident Fellowship to facilitate work on this book. I gratefully acknowledge support from the Alberta Provincial Biomedical Engineering Graduate Programme, funded by a grant from the Whitaker Foundation, toward

student assistantship for preparation of some of the exercises and illustrations for this book and the related courses ENEL 563 Biomedical Signal Analysis and ENEL 697 Digital Image Processing at the University of Calgary. I am pleased to place on record my gratitude for the generous support from the Department of Electrical and Computer Engineering and the Faculty of Engineering at the University of Calgary in terms of supplies, services, and relief from other duties.

I thank Steven Leikeim for help with computer-related issues and problems.

My association with the IEEE Engineering in Medicine and Biology Society (EMBS) in many positions has benefited me considerably in numerous ways. In particular, the period as an Associate Editor of the *IEEE Transactions on Biomedical Engineering* was rewarding, as it provided me with a wonderful opportunity to work with many leading researchers and authors of scientific articles. I thank IEEE EMBS and SPIE for lending professional support to my career on many fronts.

Writing this book has been a monumental task, often draining me of all of my energy. The infinite source of inspiration and recharging of my energy has been my family — my wife Mayura, my daughter Vidya, and my son Adarsh. While supporting me with their love and affection, they have had to bear the loss of my time and effort at home. I express my sincere gratitude to my family for their love and support, and place on record their contribution toward the preparation of this book.

I thank CRC Press and its associates for inviting me to write this book and for completing the publication process in a friendly and efficient manner.

*Rangaraj Mandayam Rangayyan
Calgary, Alberta, Canada
November, 2004*

Contents

Preface	vii
About the Author	xiii
Acknowledgments	xv
Symbols and Abbreviations	xxix
1 The Nature of Biomedical Images	1
1.1 Body Temperature as an Image	2
1.2 Transillumination	6
1.3 Light Microscopy	6
1.4 Electron Microscopy	10
1.5 X-ray Imaging	15
1.5.1 Breast cancer and mammography	22
1.6 Tomography	27
1.7 Nuclear Medicine Imaging	36
1.8 Ultrasonography	43
1.9 Magnetic Resonance Imaging	47
1.10 Objectives of Biomedical Image Analysis	53
1.11 Computer-aided Diagnosis	55
1.12 Remarks	57
1.13 Study Questions and Problems	57
1.14 Laboratory Exercises and Projects	58
2 Image Quality and Information Content	61
2.1 Difficulties in Image Acquisition and Analysis	61
2.2 Characterization of Image Quality	64
2.3 Digitization of Images	65
2.3.1 Sampling	65
2.3.2 Quantization	66
2.3.3 Array and matrix representation of images	69
2.4 Optical Density	72
2.5 Dynamic Range	73
2.6 Contrast	75
2.7 Histogram	78
2.8 Entropy	84
2.9 Blur and Spread Functions	90
2.10 Resolution	99

2.11	The Fourier Transform and Spectral Content	99
2.11.1	Important properties of the Fourier transform	110
2.12	Modulation Transfer Function	122
2.13	Signal-to-Noise Ratio	131
2.14	Error-based Measures	138
2.15	Application: Image Sharpness and Acutance	139
2.16	Remarks	145
2.17	Study Questions and Problems	145
2.18	Laboratory Exercises and Projects	149
3	Removal of Artifacts	151
3.1	Characterization of Artifacts	151
3.1.1	Random noise	151
3.1.2	Examples of noise PDFs	159
3.1.3	Structured noise	164
3.1.4	Physiological interference	165
3.1.5	Other types of noise and artifact	166
3.1.6	Stationary versus nonstationary processes	166
3.1.7	Covariance and cross-correlation	168
3.1.8	Signal-dependent noise	169
3.2	Synchronized or Multiframe Averaging	171
3.3	Space-domain Local-statistics-based Filters	174
3.3.1	The mean filter	176
3.3.2	The median filter	177
3.3.3	Order-statistic filters	181
3.4	Frequency-domain Filters	193
3.4.1	Removal of high-frequency noise	194
3.4.2	Removal of periodic artifacts	199
3.5	Matrix Representation of Image Processing	202
3.5.1	Matrix representation of images	203
3.5.2	Matrix representation of transforms	206
3.5.3	Matrix representation of convolution	212
3.5.4	Illustrations of convolution	215
3.5.5	Diagonalization of a circulant matrix	218
3.5.6	Block-circulant matrix representation of a 2D filter	221
3.6	Optimal Filtering	224
3.6.1	The Wiener filter	225
3.7	Adaptive Filters	228
3.7.1	The local LMMSE filter	228
3.7.2	The noise-updating repeated Wiener filter	234
3.7.3	The adaptive 2D LMS filter	235
3.7.4	The adaptive rectangular window LMS filter	237
3.7.5	The adaptive-neighborhood filter	241
3.8	Comparative Analysis of Filters for Noise Removal	251
3.9	Application: Multiframe Averaging in Confocal Microscopy	270

3.10	Application: Noise Reduction in Nuclear Medicine Imaging	271
3.11	Remarks	277
3.12	Study Questions and Problems	281
3.13	Laboratory Exercises and Projects	283
4	Image Enhancement	285
4.1	Digital Subtraction Angiography	286
4.2	Dual-energy and Energy-subtraction X-ray Imaging	287
4.3	Temporal Subtraction	291
4.4	Gray-scale Transforms	291
4.4.1	Gray-scale thresholding	291
4.4.2	Gray-scale windowing	292
4.4.3	Gamma correction	294
4.5	Histogram Transformation	301
4.5.1	Histogram equalization	301
4.5.2	Histogram specification	305
4.5.3	Limitations of global operations	310
4.5.4	Local-area histogram equalization	310
4.5.5	Adaptive-neighborhood histogram equalization	311
4.6	Convolution Mask Operators	314
4.6.1	Unsharp masking	314
4.6.2	Subtracting Laplacian	316
4.6.3	Limitations of fixed operators	323
4.7	High-frequency Emphasis	325
4.8	Homomorphic Filtering for Enhancement	328
4.8.1	Generalized linear filtering	328
4.9	Adaptive Contrast Enhancement	338
4.9.1	Adaptive-neighborhood contrast enhancement	338
4.10	Objective Assessment of Contrast Enhancement	346
4.11	Application: Contrast Enhancement of Mammograms	350
4.11.1	Clinical evaluation of contrast enhancement	354
4.12	Remarks	357
4.13	Study Questions and Problems	358
4.14	Laboratory Exercises and Projects	361
5	Detection of Regions of Interest	363
5.1	Thresholding and Binarization	364
5.2	Detection of Isolated Points and Lines	365
5.3	Edge Detection	367
5.3.1	Convolution mask operators for edge detection	367
5.3.2	The Laplacian of Gaussian	370
5.3.3	Scale-space methods for multiscale edge detection	380
5.3.4	Canny's method for edge detection	390
5.3.5	Fourier-domain methods for edge detection	390
5.3.6	Edge linking	392

5.4	Segmentation and Region Growing	393
5.4.1	Optimal thresholding	395
5.4.2	Region-oriented segmentation of images	396
5.4.3	Splitting and merging of regions	397
5.4.4	Region growing using an additive tolerance	397
5.4.5	Region growing using a multiplicative tolerance	400
5.4.6	Analysis of region growing in the presence of noise	401
5.4.7	Iterative region growing with multiplicative tolerance	402
5.4.8	Region growing based upon the human visual system	405
5.4.9	Application: Detection of calcifications by multitol- erance region growing	410
5.4.10	Application: Detection of calcifications by linear pre- diction error	414
5.5	Fuzzy-set-based Region Growing to Detect Breast Tumors	417
5.5.1	Preprocessing based upon fuzzy sets	419
5.5.2	Fuzzy segmentation based upon region growing	421
5.5.3	Fuzzy region growing	429
5.6	Detection of Objects of Known Geometry	434
5.6.1	The Hough transform	435
5.6.2	Detection of straight lines	437
5.6.3	Detection of circles	440
5.7	Methods for the Improvement of Contour or Region Estimates	444
5.8	Application: Detection of the Spinal Canal	449
5.9	Application: Detection of the Breast Boundary in Mammo- grams	451
5.9.1	Detection using the traditional active deformable con- tour model	456
5.9.2	Adaptive active deformable contour model	464
5.9.3	Results of application to mammograms	476
5.10	Application: Detection of the Pectoral Muscle in Mammo- grams	481
5.10.1	Detection using the Hough transform	481
5.10.2	Detection using Gabor wavelets	487
5.10.3	Results of application to mammograms	495
5.11	Application: Improved Segmentation of Breast Masses by Fuzzy-set-based Fusion of Contours and Regions	500
5.12	Remarks	524
5.13	Study Questions and Problems	527
5.14	Laboratory Exercises and Projects	527
6	Analysis of Shape	529
6.1	Representation of Shapes and Contours	529
6.1.1	Signatures of contours	530
6.1.2	Chain coding	530
6.1.3	Segmentation of contours	534

6.1.4	Polygonal modeling of contours	537
6.1.5	Parabolic modeling of contours	543
6.1.6	Thinning and skeletonization	548
6.2	Shape Factors	549
6.2.1	Compactness	551
6.2.2	Moments	555
6.2.3	Chord-length statistics	560
6.3	Fourier Descriptors	562
6.4	Fractional Concavity	569
6.5	Analysis of Spicularity	570
6.6	Application: Shape Analysis of Calcifications	575
6.7	Application: Shape Analysis of Breast Masses and Tumors	578
6.8	Remarks	581
6.9	Study Questions and Problems	581
6.10	Laboratory Exercises and Projects	582
7	Analysis of Texture	583
7.1	Texture in Biomedical Images	584
7.2	Models for the Generation of Texture	584
7.2.1	Random texture	589
7.2.2	Ordered texture	589
7.2.3	Oriented texture	590
7.3	Statistical Analysis of Texture	596
7.3.1	The gray-level co-occurrence matrix	597
7.3.2	Haralick's measures of texture	600
7.4	Laws' Measures of Texture Energy	603
7.5	Fractal Analysis	605
7.5.1	Fractal dimension	608
7.5.2	Fractional Brownian motion model	609
7.5.3	Fractal analysis of texture	609
7.5.4	Applications of fractal analysis	611
7.6	Fourier-domain Analysis of Texture	612
7.7	Segmentation and Structural Analysis of Texture	621
7.7.1	Homomorphic deconvolution of periodic patterns	623
7.8	Audification and Sonification of Texture in Images	625
7.9	Application: Analysis of Breast Masses Using Texture and Gradient Measures	627
7.9.1	Adaptive normals and ribbons around mass margins	629
7.9.2	Gradient and contrast measures	632
7.9.3	Results of pattern classification	635
7.10	Remarks	637
7.11	Study Questions and Problems	637
7.12	Laboratory Exercises and Projects	638

8	Analysis of Oriented Patterns	639
8.1	Oriented Patterns in Images	639
8.2	Measures of Directional Distribution	641
8.2.1	The rose diagram	641
8.2.2	The principal axis	641
8.2.3	Angular moments	642
8.2.4	Distance measures	643
8.2.5	Entropy	643
8.3	Directional Filtering	644
8.3.1	Sector filtering in the Fourier domain	646
8.3.2	Thresholding of the component images	649
8.3.3	Design of fan filters	651
8.4	Gabor Filters	657
8.4.1	Multiresolution signal decomposition	660
8.4.2	Formation of the Gabor filter bank	664
8.4.3	Reconstruction of the Gabor filter bank output	665
8.5	Directional Analysis via Multiscale Edge Detection	666
8.6	Hough-Radon Transform Analysis	671
8.6.1	Limitations of the Hough transform	671
8.6.2	The Hough and Radon transforms combined	673
8.6.3	Filtering and integrating the Hough-Radon space	676
8.7	Application: Analysis of Ligament Healing	679
8.7.1	Analysis of collagen remodeling	680
8.7.2	Analysis of the microvascular structure	684
8.8	Application: Detection of Breast Tumors	699
8.8.1	Framework for pyramidal decomposition	707
8.8.2	Segmentation based upon density slicing	710
8.8.3	Hierarchical grouping of iso-intensity contours	712
8.8.4	Results of segmentation of masses	712
8.8.5	Detection of masses in full mammograms	719
8.8.6	Analysis of mammograms using texture flow-field	726
8.8.7	Adaptive computation of features in ribbons	732
8.8.8	Results of mass detection in full mammograms	735
8.9	Application: Bilateral Asymmetry in Mammograms	742
8.9.1	The fibroglandular disc	743
8.9.2	Gaussian mixture model of breast density	744
8.9.3	Delimitation of the fibroglandular disc	747
8.9.4	Motivation for directional analysis of mammograms	755
8.9.5	Directional analysis of fibroglandular tissue	757
8.9.6	Characterization of bilateral asymmetry	766
8.10	Application: Architectural Distortion in Mammograms	775
8.10.1	Detection of spiculated lesions and distortion	775
8.10.2	Phase portraits	779
8.10.3	Estimating the orientation field	780
8.10.4	Characterizing orientation fields with phase portraits	782

8.10.5	Feature extraction for pattern classification	785
8.10.6	Application to segments of mammograms	785
8.10.7	Detection of sites of architectural distortion	786
8.11	Remarks	791
8.12	Study Questions and Problems	796
8.13	Laboratory Exercises and Projects	796
9	Image Reconstruction from Projections	797
9.1	Projection Geometry	797
9.2	The Fourier Slice Theorem	798
9.3	Backprojection	801
9.3.1	Filtered backprojection	804
9.3.2	Discrete filtered backprojection	806
9.4	Algebraic Reconstruction Techniques	813
9.4.1	Approximations to the Kaczmarz method	820
9.5	Imaging with Diffracting Sources	825
9.6	Display of CT Images	825
9.7	Agricultural and Forestry Applications	829
9.8	Microtomography	831
9.9	Application: Analysis of the Tumor in Neuroblastoma	834
9.9.1	Neuroblastoma	834
9.9.2	Tissue characterization using CT	838
9.9.3	Estimation of tissue composition from CT images	839
9.9.4	Results of application to clinical cases	844
9.9.5	Discussion	845
9.10	Remarks	854
9.11	Study Questions and Problems	854
9.12	Laboratory Exercises and Projects	855
10	Deconvolution, Deblurring, and Restoration	857
10.1	Linear Space-invariant Restoration Filters	857
10.1.1	Inverse filtering	858
10.1.2	Power spectrum equalization	860
10.1.3	The Wiener filter	863
10.1.4	Constrained least-squares restoration	872
10.1.5	The Metz filter	874
10.1.6	Information required for image restoration	875
10.1.7	Motion deblurring	875
10.2	Blind Deblurring	877
10.2.1	Iterative blind deblurring	878
10.3	Homomorphic Deconvolution	885
10.3.1	The complex cepstrum	885
10.3.2	Echo removal by Radon-domain cepstral filtering	886
10.4	Space-variant Restoration	891
10.4.1	Sectioned image restoration	893

10.4.2	Adaptive-neighborhood deblurring	894
10.4.3	The Kalman filter	898
10.5	Application: Restoration of Nuclear Medicine Images	919
10.5.1	Quality control	922
10.5.2	Scatter compensation	922
10.5.3	Attenuation correction	923
10.5.4	Resolution recovery	924
10.5.5	Geometric averaging of conjugate projections	926
10.5.6	Examples of restoration of SPECT images	934
10.6	Remarks	949
10.7	Study Questions and Problems	953
10.8	Laboratory Exercises and Projects	954
11	Image Coding and Data Compression	955
11.1	Considerations Based on Information Theory	956
11.1.1	Noiseless coding theorem for binary transmission	957
11.1.2	Lossy versus lossless compression	957
11.1.3	Distortion measures and fidelity criteria	959
11.2	Fundamental Concepts of Coding	960
11.3	Direct Source Coding	961
11.3.1	Huffman coding	961
11.3.2	Run-length coding	969
11.3.3	Arithmetic coding	969
11.3.4	Lempel–Ziv coding	974
11.3.5	Contour coding	977
11.4	Application: Source Coding of Digitized Mammograms	978
11.5	The Need for Decorrelation	980
11.6	Transform Coding	984
11.6.1	The discrete cosine transform	987
11.6.2	The Karhunen–Loève transform	989
11.6.3	Encoding of transform coefficients	992
11.7	Interpolative Coding	1001
11.8	Predictive Coding	1004
11.8.1	Two-dimensional linear prediction	1005
11.8.2	Multichannel linear prediction	1009
11.8.3	Adaptive 2D recursive least-squares prediction	1026
11.9	Image Scanning Using the Peano-Hilbert Curve	1033
11.9.1	Definition of the Peano-scan path	1035
11.9.2	Properties of the Peano-Hilbert curve	1040
11.9.3	Implementation of Peano scanning	1040
11.9.4	Decorrelation of Peano-scanned data	1041
11.10	Image Coding and Compression Standards	1043
11.10.1	The JBIG Standard	1046
11.10.2	The JPEG Standard	1049
11.10.3	The MPEG Standard	1050

11.10.4	The ACR/ NEMA and DICOM Standards	1050
11.11	Segmentation-based Adaptive Scanning	1051
11.11.1	Segmentation-based coding	1051
11.11.2	Region-growing criteria	1052
11.11.3	The SLIC procedure	1055
11.11.4	Results of image data compression with SLIC	1055
11.12	Enhanced JBIG Coding	1062
11.13	Lower-limit Analysis of Lossless Data Compression	1066
11.13.1	Memoryless entropy	1070
11.13.2	Markov entropy	1071
11.13.3	Estimation of the true source entropy	1071
11.14	Application: Teleradiology	1079
11.14.1	Analog teleradiology	1080
11.14.2	Digital teleradiology	1082
11.14.3	High-resolution digital teleradiology	1084
11.15	Remarks	1086
11.16	Study Questions and Problems	1086
11.17	Laboratory Exercises and Projects	1087
12	Pattern Classification and Diagnostic Decision	1089
12.1	Pattern Classification	1091
12.2	Supervised Pattern Classification	1095
12.2.1	Discriminant and decision functions	1095
12.2.2	Distance functions	1097
12.2.3	The nearest-neighbor rule	1104
12.3	Unsupervised Pattern Classification	1104
12.3.1	Cluster-seeking methods	1105
12.4	Probabilistic Models and Statistical Decision	1110
12.4.1	Likelihood functions and statistical decision	1110
12.4.2	Bayes classifier for normal patterns	1118
12.5	Logistic Regression	1120
12.6	The Training and Test Steps	1125
12.6.1	The leave-one-out method	1125
12.7	Neural Networks	1126
12.8	Measures of Diagnostic Accuracy	1132
12.8.1	Receiver operating characteristics	1135
12.8.2	McNemar's test of symmetry	1138
12.9	Reliability of Features, Classifiers, and Decisions	1140
12.9.1	Statistical separability and feature selection	1141
12.10	Application: Image Enhancement for Breast Cancer Screen- ing	1143
12.10.1	Case selection, digitization, and presentation	1145
12.10.2	ROC and statistical analysis	1147
12.10.3	Discussion	1159

12.11	Application: Classification of Breast Masses and Tumors via Shape Analysis	1160
12.12	Application: Content-based Retrieval and Analysis of Breast Masses	1166
12.12.1	Pattern classification of masses	1167
12.12.2	Content-based retrieval	1169
12.12.3	Extension to telemedicine	1177
12.13	Remarks	1182
12.14	Study Questions and Problems	1184
12.15	Laboratory Exercises and Projects	1185
	References	1187
	Index	1262

Symbols and Abbreviations

Note: Bold-faced letters represent the vector or matrix form of the variable in the corresponding plain letters.

Variables or symbols used within limited contexts are not listed here; they are described within their context.

The mathematical symbols listed may stand for other entities or variables in different applications; only the common associations used in this book are listed for ready reference.

$a(p, q)$, a	autoregressive model or filter coefficients
arctan	inverse tangent, \tan^{-1}
arg	argument of
atan	inverse tangent, \tan^{-1}
au	arbitrary units
AADCM	adaptive active deformable contour model
ACF	autocorrelation function
ACR	American College of Radiology
ADC	analog-to-digital converter
ALZ	adaptive Lempel-Ziv coding
AMTA	AMT acutance
ANCE	adaptive-neighborhood contrast enhancement
AND	adaptive-neighborhood deblurring
ANN	artificial neural network
ANNS	adaptive-neighborhood noise subtraction
AR	autoregressive model or filter
ARMA	autoregressive, moving-average model or filter
ARW	adaptive rectangular window
A_z	area under the ROC curve
b	background intensity
b	bit
$b(m, n)$	moving-average model or filter coefficients
bps	bits per second
B	byte
BIBO	bounded-input – bounded-output stability
BI-RADS™	Breast Imaging Reporting and Data System
BP	backprojection
cd	candela

<i>cm</i>	centimeter
<i>C</i>	contrast
C	covariance matrix
<i>C_i</i>	Curie
<i>cf, C_o</i>	compactness
<i>C_i</i>	the i^{th} class in a pattern classification problem
<i>C_{xy}</i>	covariance between x and y
CAD	computer-aided diagnosis
CBP	convolution backprojection
CBIR	content-based image retrieval
CC	cranio-caudal
CCD	charge-coupled device
CCF	cross-correlation function
CCITT	Comité Consultatif International Téléphonique et Télégraphique
CD	compact disk
CLS	constrained least squares
CMTA	cascaded modulation transfer acutance
CMYK	[cyan, magenta, yellow, black] representation of color
CNR	contrast-to-noise ratio
CNS	central nervous system
CR	computed radiography
CREW	compression with reversible embedded wavelets
CRT	cathode ray tube
CSD	cross-spectral density, cross-spectrum
CT	computed tomography
CV	coefficient of variation
<i>dB</i>	decibel
<i>d_E</i>	Euclidean dimension
<i>d_f</i>	fractal dimension
<i>dpi</i>	dots per inch
DAC	digital-to-analog converter
DC	direct current; zero frequency
DCT	discrete cosine transform
DFT	discrete Fourier transform
DICOM	Digital Imaging and Communications in Medicine
DoG	difference of Gaussians
DPCM	differential pulse code modulation
DR	digital radiography
DSA	digital subtraction angiography
DWT	directional wavelet transform
$e(n), E(\omega)$	model or estimation error
<i>eV</i>	electron volt
$\exp(x)$	exponential function, e^x
ECG	electrocardiogram, electrocardiography

EEG	electroencephalogram
EM	electromagnetic
EM	expectation-maximization
E_x	total energy of the signal x
$E[\]$	statistical expectation operator
f	foreground intensity
f_c	cutoff frequency (usually at -3 dB) of a filter
f_{cc}	fractional concavity
ff	shape factor obtained using Fourier descriptors
fps	frames per second
f_s	sampling frequency
$f(m, n)$	a digital image, typically original or undistorted
$f(x, y)$	an image, typically original or undistorted
FBP	filtered backprojection
FFT	fast Fourier transform
FID	free-induction decay
FIR	finite impulse response (filter)
FM	frequency modulation
FN	false negative
FNF	false-negative fraction
FOM	figure of merit
FP	false positive
FPF	false-positive fraction
FT	Fourier transform
FWHM	full width at half the maximum
$g(m, n)$	a digital image, typically processed or distorted
$g(x, y)$	an image, typically processed or distorted
$h(m, n)$	impulse response of a system
$h(p)$	measure of information
$h(x, y)$	impulse response of a system
H	hydrogen
H	Hurst coefficient
H	magnetic field strength
H	entropy
$H_{f,g}$	joint entropy of f and g
$H_{f g}$	conditional entropy of f given g
H	Hermitian (complex-conjugate) transposition of a matrix
$H(k, l)$	discrete Fourier transform of $h(m, n)$
$H(u, v)$	frequency response of a filter, Fourier transform of $h(x, y)$
HINT	hierarchical interpolation
HU	Hounsfield unit
HVS	human visual system
Hz	Hertz
i	index of a series

I	the identity matrix
$I_{f g}$	mutual information
IEPA	image edge-profile acutance
IFT	inverse Fourier transform
IIR	infinite impulse response (filter)
ISO	International Organization for Standardization
j	$\sqrt{-1}$
JBIG	Joint Bi-level Image (experts) Group
JM	Jeffries-Matusita distance
JND	just-noticeable difference
JPEG	Joint Photographic Experts Group
k	kilo (1,000)
(k, l)	indices in the discrete Fourier (frequency) domain
kVp	kilo-volt peak
K	kilo (1,024)
KESF	knife-edge spread function
KLT	Karhunen-Loève transform
ln	natural logarithm (base e)
lp/mm	line pairs per millimeter
L	an image processing operator or transform in matrix form
L_{ij}	loss function in pattern classification
LEAP	low-energy all-purpose collimator
LEGP	low-energy general-purpose collimator
LLMMSE	local linear minimum mean-squared error
LMMSE	linear minimum mean-squared error
LMS	least mean squares
LMSE	Laplacian mean-squared error
LoG	Laplacian of Gaussian
LP	linear prediction (model)
LSF	line spread function
LSI	linear shift-invariant
LUT	look-up table
LZW	Lempel-Ziv-Welch code
m	meter
m	mean
m	mean vector of a pattern class
max	maximum
mA	milliamperere
mf	shape factor using moments
min	minimum
mm	millimeter
(m, n)	indices in the discrete space (image) domain
mod	modulus or modulo
modem	modulator – demodulator

M	number of samples or pixels
MA	moving average (filter)
MAP	maximum-a-posteriori probability
MCL	medial collateral ligament
MIAS	Mammographic Image Analysis Society, London, England
MDL	minimum description length
ME	maximum entropy
MLO	medio-lateral oblique
MMSE	minimum mean-squared error
MPEG	Moving Picture Experts Group
MR	magnetic resonance
MRI	magnetic resonance imaging
MS	mean-squared
MSE	mean-squared error
MTF	modulation (magnitude) transfer function
n	an index
nm	nanometer
N	number of samples or pixels
NE	normalized error
NEMA	National Electrical Manufacturers Association
NMR	nuclear magnetic resonance
NMSE	normalized mean-squared error
NPV	negative predictive value
NSHP	nonsymmetric half plane
OD	optical density
OTF	optical transfer function
$p_f(l)$	normalized histogram or PDF of image f
$p_{f,g}(l_1, l_2)$	joint PDF of images f and g
$p_{f g}(l_1, l_2)$	conditional PDF of f given g
pixel	picture cell or element
p_m	m^{th} ray sum in ART
pps	pulses per second
(p, q)	indices of a 2D array
$p_\theta(t)$	projection (Radon transform) of an image at angle θ
$p(x)$	probability density function of the random variable x
$p(x C_i)$	likelihood function of class C_i or state-conditional PDF of x
P	model order
$P(x)$	probability of the event x
$P_f(l)$	histogram of image f
$P(C_i x)$	posterior probability that x belongs to the class C_i
$F_\theta(w)$	Fourier transform of the projection $p_\theta(t)$
\mathcal{P}	a predicate
PA	posterior–anterior
PACS	picture archival and communication system

PCA	principal-component analysis
PCG	phonocardiogram (heart sound signal)
PDF	probability density function
PET	positron emission tomography
PMSE	perceptual mean-squared error
PMT	photomultiplier tube
<i>PPV</i>	positive predictive value
PSD	power spectral density, power spectrum
PSE	power spectrum equalization
PSF	point spread function
PSV	prediction selection values in JPEG
$q_\theta(t)$	filtered projection of an image at angle θ
Q	model order
QP	quarter plane
$r_j(\mathbf{x})$	average risk or loss in pattern classification
(r, s)	temporary indices of a 2D array
\mathcal{R}^+	the set of nonnegative real numbers
RBST	rubber-band straightening transform
RD	relative dispersion
RDM	radial distance measures
RF	radio-frequency
RGB	[red, green, blue] color representation
RLS	recursive least-squares
RMS	root mean-squared
ROC	receiver operating characteristics
ROI	region of interest
ROS	region of support
RUKF	reduced-update Kalman filter
s	second
s	space variable in the projection (Radon) space
$S_f(u, v)$	power spectral density of the image f
SAR	synthetic-aperture radar
SD	standard deviation
SEM	scanning electron microscope
SI	spiculation index
SLIC	segmentation-based lossless image coding
SMTA	system modulation transfer acutance
SNR	signal-to-noise ratio
SPECT	single-photon emission computed tomography
SQF	subjective quality factor
SQRI	square-root integral
STFT	short-time Fourier transform
SVD	singular value decomposition
S^+	sensitivity of a test
S^-	specificity of a test

t	time variable
t	space variable in the projection (Radon) space
T	Tesla (strength of a magnetic field)
T	a threshold
T	as a superscript, vector or matrix transposition
Tc	technetium
Tl	thallium
T_1	longitudinal relaxation time constant in MRI
T_2	transverse magnetization time constant in MRI
T^+	positive test result
T^-	negative test result
TEM	transmission electron microscope
Th	threshold
TN	true negative
TNF	true-negative fraction
TP	true positive
TPF	true-positive fraction
Tr	trace of a matrix
TSE	total squared error
TV	television
$u(x, y)$	unit step function
(u, v)	frequency coordinates in the continuous Fourier domain
UHF	ultra high frequency
voxel	volume cell or element
V	volt
VLSI	very-large-scale integrated circuit
w	filter tap weight; weighting function
w	frequency variable related to projections
\mathbf{w}	filter or weight vector
WHT	Walsh-Hadamard transform
W_N	Fourier transform kernel function $W_N = \exp(-j\frac{2\pi}{N})$
\mathbf{W}	Fourier transform operator in matrix form
(x, y)	image coordinates in the space domain
\mathbf{x}	a feature vector in pattern classification
YIQ	[luminance, in-phase, quadrature] color representation
\mathbf{z}	a prototype feature vector in pattern classification
\mathcal{Z}	the set of all integers
\emptyset	null set
1D	one-dimensional
2D	two-dimensional
3D	three-dimensional
4D	four-dimensional
γ_{xy}	correlation coefficient between x and y
Γ_{xy}	coherence between x and y : Fourier transform of γ_{xy}

$\Gamma(p)$	a fuzzy membership function
δ	Dirac delta (impulse) function
$\Delta x, \Delta y$	sampling intervals along the x and y axes
ε, ϵ	model error, total squared error
η	a random variable or noise process
η	scale factor in fractal analysis
θ	an angle
θ	a threshold
θ, Θ	cross-correlation function
(θ, t)	the Radon (projection) space
λ	forgetting factor in the RLS filter
μ	the mean (average) of a random variable
μ	X-ray attenuation coefficient
μm	micrometer
μCT	micro-computed tomography
ρ	correlation coefficient
σ	the standard deviation of a random variable
σ^2	the variance of a random variable
σ_{fg}	covariance between images f and g
σ_{fg}	covariance between images f and g in matrix form
φ	basis function of a transform
ϕ_f	autocorrelation of image f in array form
ϕ_f	autocorrelation of image f in matrix form
ϕ_{fg}	cross-correlation between images f and g in array form
ϕ_{fg}	cross-correlation between images f and g in matrix form
Φ_f	Fourier transform of ϕ_f ; power spectral density of f
∇	gradient operator
$\cdot, \bullet, \langle, \rangle$	dot product
$!$	factorial
$*$	when in-line, convolution
$*$	as a superscript, complex conjugation
$\#$	number of
$\bar{\quad}$	average or normalized version of the variable under the bar
$\bar{\quad}$	complement of the variable under the bar
$\hat{\quad}$	complex cepstrum of the signal (function of space)
$\hat{\quad}$	complex logarithm of the signal (function of frequency)
$\tilde{\quad}$	estimate of the variable under the symbol
\sim	a variant of a function
$' , '' , '''$	first, second, and third derivatives of the preceding function
$'$	a variant of a function
\times	cross product when the related entities are vectors
\forall	for all
\in	belongs to or is in (the set)
$\{ \}$	a set
\subset	subset

\supset	superset
\cap	intersection
\cup	union
\equiv	equivalent to
$ $	given, conditional upon
\rightarrow	maps to
\leftarrow	gets (updated as)
\Rightarrow	leads to
\Leftrightarrow	transform pair
$[\]$	closed interval, including the limits
$(\)$	open interval, not including the limits
$ $	absolute value or magnitude
$ $	determinant of a matrix
$\ \ $	norm of a vector or matrix
$\lceil x \rceil$	ceiling operator; the smallest integer $\geq x$
$\lfloor x \rfloor$	floor operator; the largest integer $\leq x$

The Nature of Biomedical Images

The human body is composed of many *systems*, such as the cardiovascular system, the musculo-skeletal system, and the central nervous system. Each system is made up of several subsystems that carry on many *physiological processes*. For example, the visual system performs the task of focusing visual or pictorial information on to the retina, transduction of the image information into neural signals, and encoding and transmission of the neural signals to the visual cortex. The visual cortex is responsible for interpretation of the image information. The cardiac system performs the important task of rhythmic pumping of blood through the arterial network of the body to facilitate the delivery of nutrients, as well as pumping of blood through the pulmonary system for oxygenation of the blood itself. The anatomical features of the organs related to a physiological system often demonstrate characteristics that reflect the functional aspects of its processes as well as the well-being or integrity of the system itself.

Physiological processes are complex phenomena, including neural or hormonal stimulation and control; inputs and outputs that could be in the form of physical material or information; and action that could be mechanical, electrical, or biochemical. Most physiological processes are accompanied by or manifest themselves as *signals* that reflect their nature and activities. Such signals could be of many types, including biochemical in the form of hormones or neurotransmitters, electrical in the form of potential or current, and physical in the form of pressure or temperature.

Diseases or defects in a physiological system cause alterations in its normal processes, leading to *pathological processes* that affect the performance, health, and general well-being of the system. A pathological process is typically associated with signals and anatomical features that are different in some respects from the corresponding normal patterns. If we possess a good understanding of a system of interest, it becomes possible to observe the corresponding signals and features and assess the state of the system. The task is not difficult when the signal is simple and appears at the outer surface of the body. However, most systems and organs are placed well within the body and enclosed in protective layers (for good reason!). Investigating or probing such systems typically requires the use of some form of penetrating radiation or invasive procedure.

1.1 Body Temperature as an Image

Most infections cause a rise in the temperature of the body, which may be sensed easily, albeit in a relative and *qualitative* manner, via the palm of one's hand. Objective or *quantitative* measurement of temperature requires an instrument, such as a thermometer.

A single measurement f of temperature is a *scalar*, and represents the thermal state of the body at a particular physical location in or on the body denoted by its spatial coordinates (x, y, z) and at a particular or single instant of time t . If we record the temperature continuously in some form, such as a strip-chart record, we obtain a signal as a one-dimensional (1D) function of time, which may be expressed in the continuous-time or analog form as $f(t)$. The units applicable here are $^{\circ}C$ (degrees Celsius) for the temperature variable, and s (seconds) for the temporal variable t . If some means were available to measure the temperature of the body at *every spatial position*, we could obtain a three-dimensional (3D) distribution of temperature as $f(x, y, z)$. Furthermore, if we were to perform the 3D measurement at every instant of time, we would obtain a 3D function of time as $f(x, y, z, t)$; this entity may also be referred to as a four-dimensional (4D) function.

When oral temperature, for example, is measured at discrete instants of time, it may be expressed in *discrete-time* form as $f(nT)$ or $f(n)$, where n is the index or measurement sample number of the array of values, and T represents the uniform interval between the time instants of measurement. A discrete-time signal that can take amplitude values only from a limited list of *quantized* levels is called a *digital* signal; this distinction between discrete-time and digital signals is often ignored.

If one were to use a thermal camera and take a picture of a body, a two-dimensional (2D) representation of the heat radiated from the body would be obtained. Although the temperature distribution within the body (and even on the surface of the body) is a 3D entity, the picture produced by the camera is a 2D snapshot of the heat radiation field. We then have a 2D spatial function of temperature — an *image* — which could be represented as $f(x, y)$. The units applicable here are $^{\circ}C$ for the temperature variable itself, and mm (millimeters) for the spatial variables x and y . If the image were to be sampled in space and represented on a discrete spatial grid, the corresponding data could be expressed as $f(m\Delta x, n\Delta y)$, where Δx and Δy are the sampling intervals along the horizontal and vertical axes, respectively (in spatial units such as mm). It is common practice to represent a digital image simply as $f(m, n)$, which could be interpreted as a 2D array or a matrix of values. It should be noted at the outset that, while images are routinely treated as arrays, matrices, and related mathematical entities, they are almost always representative of physical or other measures of organs or of physiological pro-

cesses that impose practical limitations on the range, degrees of freedom, and other properties of the image data.

Examples: In intensive-care monitoring, the tympanic (ear drum) temperature is often measured using an infrared sensor. Occasionally, when catheters are being used for other purposes, a temperature sensor may also be introduced into an artery or the heart to measure the *core* temperature of the body. It then becomes possible to obtain a continuous measurement of temperature, although only a few samples taken at intervals of a few minutes may be stored for subsequent analysis. Figure 1.1 illustrates representations of temperature measurements as a scalar, an array, and a signal that is a function of time. It is obvious that the graphical representation facilitates easier and faster comprehension of trends in the temperature than the numerical format. Long-term recordings of temperature can facilitate the analysis of temperature-regulation mechanisms [15, 16].


Infrared (with wavelength in the range 3,000–5,000 *nm*) or thermal sensors may also be used to capture the heat radiated or emitted from a body or a part of a body as an image. Thermal imaging has been investigated as a potential tool for the detection of breast cancer. A tumor is expected to be more vascularized than its neighboring tissues, and hence could be at a slightly higher temperature. The skin surface near the tumor may also demonstrate a relatively high temperature. Temperature differences of the order of 2°C have been measured between surface regions near breast tumors and neighboring tissues. Figure 1.2 shows thermal images of a patient with benign fibrocysts and a patient with breast cancer; the local increase in temperature due to a tumor is evident in the latter case. Thermography can help in the diagnosis of advanced cancer, but has limited success in the detection of early breast cancer [17, 18]. Recent improvements in detectors and imaging techniques have created a renewed interest in the application of thermography for the detection of breast cancer [19, 20, 21, 22, 23].

Infrared imaging via a telethermographic camera has been applied to the detection of varicocele, which is the most common cause of infertility in men [24, 25, 26]. In normal men, the testicular temperature is about 3–4 °C below the core body temperature. In the case of varicocele, dilation of the testicular veins reduces the venous return from the scrotum, causes stagnation of blood and edema, and leads to increased testicular temperature. In the experiments conducted by Merla et al. [25], a cold patch was applied to the subject's scrotum, and the thermal recovery curves were analyzed. The results obtained showed that the technique was successful in detecting subclinical varicocele. Vlaisavljević [26] showed that telethermography can provide better diagnostic accuracy in the detection of varicocele than contact thermography.

33.5 °C
(a)

Time (hours)	08	10	12	14	16	18	20	22	24
Temperature (°C)	33.5	33.3	34.5	36.2	37.3	37.5	38.0	37.8	38.0

(b)



(c)

FIGURE 1.1

Measurements of the temperature of a patient presented as (a) a scalar with one temperature measurement f at a time instant t ; (b) an array $f(n)$ made up of several measurements at different instants of time; and (c) a signal $f(t)$ or $f(n)$. The horizontal axis of the plot represents time in hours; the vertical axis gives temperature in degrees Celsius. Data courtesy of Foothills Hospital, Calgary.




FIGURE 1.2

Body temperature as a 2D image $f(x, y)$ or $f(m, n)$. The images illustrate the distribution of surface temperature measured using an infrared camera operating in the $3,000 - 5,000 \text{ nm}$ wavelength range. (a) Image of a patient with pronounced vascular features and benign fibrocysts in the breasts. (b) Image of a patient with a malignant mass in the upper-outer quadrant of the left breast. Images courtesy of P. Hoekstra, III, Therma-Scan, Inc., Huntington Woods, MI.

The thermal images shown in Figure 1.2 serve to illustrate an important distinction between two major categories of medical images:

- anatomical or physical images, and
- functional or physiological images.

The images illustrate the notion of body temperature as a signal or image. Each point in the images in Figure 1.2 represents body temperature, which is related to the ongoing physiological or pathological processes at the corresponding location in the body. A thermal image is, therefore, a functional image. An ordinary photograph obtained with reflected light, on the other hand, would be a purely anatomical or physical image. More sophisticated techniques that provide functional images related to circulation and various physiological processes are described in the following sections.

1.2 Transillumination

Transillumination, diaphanography, and diaphanoscopy involve the shining of visible light or near-infrared radiation through a part of the body, and viewing or imaging the transmitted radiation. The technique has been investigated for the detection of breast cancer, the attractive feature being the use of nonionizing radiation [27]. The use of near-infrared radiation appears to have more potential than visible light, due to the observation that nitrogen-rich compounds preferentially absorb (or attenuate) infrared radiation. The fat and fibroglandular tissue in the mature breast contain much less nitrogen than malignant tissues. Furthermore, the hemoglobin in blood has a high nitrogen content, and tumors are more vascularized than normal tissues. For these reasons, breast cancer appears as a relatively dark region in a transilluminated image.

The effectiveness of transillumination is limited by scatter and ineffective penetration of light through a large organ such as the breast. Transillumination has been found to be useful in differentiating between cystic (fluid-filled) and solid lesions; however, the technique has had limited success in distinguishing malignant tumors from benign masses [18, 28, 29].

1.3 Light Microscopy

Studies of the fine structure of biological cells and tissues require significant magnification for visualization of the details of interest. Useful magnification

of up to $\times 1,000$ may be obtained via light microscopy by the use of combinations of lenses. However, the resolution of light microscopy is reduced by the following factors [30]:

- **Diffraction:** The bending of light at edges causes blurring; the image of a pinhole appears as a blurred disc known as the Airy disc.
- **Astigmatism:** Due to nonuniformities in lenses, a point may appear as an ellipse.
- **Chromatic aberration:** Electromagnetic (EM) waves of different wavelength or energy that compose the ordinarily used white light converge at different focal planes, thereby causing enlargement of the focal point. This effect may be corrected for by using monochromatic light. See Section 3.9 for a description of confocal microscopy.
- **Spherical aberration:** The rays of light arriving at the periphery of a lens are refracted more than the rays along the axis of the lens. This causes the rays from the periphery and the axis not to arrive at a common focal point, thereby resulting in blurring. The effect may be reduced by using a small aperture.
- **Geometric distortion:** Poorly crafted lenses may cause geometric distortion such as the pin-cushion effect and barrel distortion.

Whereas the best resolution achievable by the human eye is of the order of $0.1 - 0.2 \text{ mm}$, light microscopes can provide resolving power up to about $0.2 \mu\text{m}$.


Example: Figure 1.3 shows a rabbit ventricular myocyte in its relaxed state as seen through a light microscope at a magnification of about $\times 600$. The experimental setup was used to study the contractility of the myocyte with the application of electrical stimuli [31].

Example: Figure 1.4 shows images of three-week-old scar tissue and forty-week-old healed tissue samples from rabbit ligaments at a magnification of about $\times 300$. The images demonstrate the alignment patterns of the nuclei of fibroblasts (stained to appear as the dark objects in the images): the three-week-old scar tissue has many fibroblasts that are scattered in different directions, whereas the forty-week-old healed sample has fewer fibroblasts that are well-aligned along the length of the ligament (the horizontal edge of the image). The appearance of the forty-week-old sample is closer to that of normal samples than that of the three-week-old sample. Images of this nature have been found to be useful in studying the healing and remodeling processes in ligaments [32].



FIGURE 1.3

A single ventricular myocyte (of a rabbit) in its relaxed state. The width (thickness) of the myocyte is approximately $15 \mu\text{m}$. Image courtesy of R. Clark, Department of Physiology and Biophysics, University of Calgary.



(a)



(b)

FIGURE 1.4

(a) Three-week-old scar tissue sample, and (b) forty-week-old healed tissue sample from rabbit medial collateral ligaments. Images courtesy of C.B. Frank, Department of Surgery, University of Calgary.


1.4 Electron Microscopy

Accelerated electrons possess EM wave properties, with the wavelength λ given by $\lambda = \frac{h}{mv}$, where h is Planck's constant, m is the mass of the electron, and v is the electron's velocity; this relationship reduces to $\lambda = \frac{1.23}{\sqrt{V}}$, where V is the accelerating voltage [30]. At a voltage of 60 kV , an electron beam has an effective wavelength of about 0.005 nm , and a resolving power limit of about 0.003 nm . Imaging at a low kV provides high contrast but low resolution, whereas imaging at a high kV provides high resolution due to smaller wavelength but low contrast due to higher penetrating power. In addition, a high- kV beam causes less damage to the specimen as the faster electrons pass through the specimen in less time than with a low- kV beam. Electron microscopes can provide useful magnification of the order of $\times 10^6$, and may be used to reveal the ultrastructure of biological tissues. Electron microscopy typically requires the specimen to be fixed, dehydrated, dried, mounted, and coated with a metal.


Transmission electron microscopy: A transmission electron microscope (TEM) consists of a high-voltage electron beam generator, a series of EM lenses, a specimen holding and changing system, and a screen-film holder, all enclosed in vacuum. In TEM, the electron beam passes through the specimen, is affected in a manner similar to light, and the resulting image is captured through a screen-film combination or viewed via a phosphorescent viewing screen.

Example: Figure 1.5 shows TEM images of collagen fibers (in cross-section) in rabbit ligament samples. The images facilitate analysis of the diameter distribution of the fibers [33]. Scar samples have been observed to have an almost uniform distribution of fiber diameter in the range 60–70 nm , whereas normal samples have an average diameter of about 150 nm over a broader distribution. Methods for the detection and analysis of circular objects are described in Sections 5.6.1, 5.6.3, and 5.8.

Example: In patients with hematuria, the glomerular basement membrane of capillaries in the kidney is thinner (< 200 nm) than the normal thickness of the order of 300 nm [34]. Investigation of this feature requires needle-core biopsy of the kidney and TEM imaging. Figure 1.6 shows a TEM image of a capillary of a normal kidney in cross-section. Figure 1.7 (a) shows an image of a sample with normal membrane thickness; Figure 1.7 (b) shows an image of a sample with reduced and variable thickness. Although the ranges of normal and abnormal membrane thickness have been established by several studies [34], the diagnostic decision process is subjective; methods for objective and quantitative analysis are desired in this application.



(a)



(b)

FIGURE 1.5

TEM images of collagen fibers in rabbit ligament samples at a magnification of approximately $\times 30,000$. (a) Normal and (b) scar tissue. Images courtesy of C.B. Frank, Department of Surgery, University of Calgary.




FIGURE 1.6

TEM image of a kidney biopsy sample at a magnification of approximately $\times 3,500$. The image shows the complete cross-section of a capillary with normal membrane thickness. Image courtesy of H. Benediktsson, Department of Pathology and Laboratory Medicine, University of Calgary.




FIGURE 1.7

TEM images of kidney biopsy samples at a magnification of approximately $\times 8,000$. (a) The sample shows normal capillary membrane thickness. (b) The sample shows reduced and varying membrane thickness. Images courtesy of H. Benediktsson, Department of Pathology and Laboratory Medicine, University of Calgary.

Scanning electron microscopy: A scanning electron microscope (SEM) is similar to a TEM in many ways, but uses a finely focused electron beam with a diameter of the order of 2 nm to scan the surface of the specimen. The electron beam is not transmitted through the specimen, which could be fairly thick in SEM. Instead, the beam is used to scan the surface of the specimen in a raster pattern, and the secondary electrons that are emitted from the surface of the sample are detected and amplified through a photomultiplier tube (PMT), and used to form an image on a cathode-ray tube (CRT). An SEM may be operated in different modes to detect a variety of signals emitted from the sample, and may be used to obtain images with a depth of field of several mm .

Example: Figure 1.8 illustrates SEM images of collagen fibers in rabbit ligament samples (freeze-fractured surfaces) [35]. The images are useful in analyzing the angular distribution of fibers and the realignment process during healing after injury. It has been observed that collagen fibers in a normal ligament are well aligned, that fibers in scar tissue lack a preferred orientation, and that organization and alignment return toward their normal patterns during the course of healing [36, 37, 35]. Image processing methods for directional analysis are described in Chapter 8.




FIGURE 1.8

SEM images of collagen fibers in rabbit ligament samples at a magnification of approximately $\times 4,000$. (a) Normal and (b) scar tissue. Reproduced with permission from C.B. Frank, B. MacFarlane, P. Edwards, R. Rangayyan, Z.Q. Liu, S. Walsh, and R. Bray, “A quantitative analysis of matrix alignment in ligament scars: A comparison of movement versus immobilization in an immature rabbit model”, *Journal of Orthopaedic Research*, 9(2): 219 – 227, 1991. © Orthopaedic Research Society.

1.5 X-ray Imaging

The medical diagnostic potential of X rays was realized soon after their discovery by Röntgen in 1895. (See Robb [38] for a review of the history of X-ray imaging.) In the simplest form of X-ray imaging or radiography, a 2D projection (shadow or silhouette) of a 3D body is produced on film by irradiating the body with X-ray photons [4, 3, 5, 6]. This mode of imaging is referred to as projection or planar imaging. Each ray of X-ray photons is attenuated by a factor depending upon the integral of the linear attenuation coefficient along the path of the ray, and produces a corresponding gray level (or signal) at the point hit on the film or the detecting device used.

Considering the ray path marked as AB in Figure 1.9, let N_i denote the number of X-ray photons incident upon the body being imaged, within a specified time interval. Let us assume that the X rays are mutually parallel, with the X-ray source at a large distance from the subject or object being imaged. Let N_o be the corresponding number of photons exiting the body. Then, we have


$$N_o = N_i \exp \left[- \int_{\text{ray AB}} \mu(x, y) ds \right], \quad (1.1)$$

or

$$\int_{\text{ray AB}} \mu(x, y) ds = \ln \left(\frac{N_i}{N_o} \right). \quad (1.2)$$

The equations above are modified versions of Beer's law (also known as the Beer-Lambert law) on the attenuation of X rays due to passage through a medium. The ray AB lies in the sectional plane PQRS; the mutually parallel rays within the plane PQRS are represented by the coordinates (t, s) that are at an angle θ with respect to the (x, y) coordinates indicated in Figure 1.9, with the s axis being parallel to the rays. Then, $s = -x \sin \theta + y \cos \theta$. The variable of integration ds represents the elemental distance along the ray, and the integral is along the ray path AB from the X-ray source to the detector. (See Section 9.1 for further details on this notation.) The quantities N_i and N_o are Poisson variables; it is assumed that their values are large for the equations above to be applicable. The function $\mu(x, y)$ represents the linear attenuation coefficient at (x, y) in the sectional plane PQRS. The value of $\mu(x, y)$ depends upon the density of the object or its constituents along the ray path, as well as the frequency (or wavelength or energy) of the radiation used. Equation 1.2 assumes the use of monochromatic or monoenergetic X rays.

A measurement of the exiting X rays (that is, N_o , and N_i for reference) thus gives us only an integral of $\mu(x, y)$ over the ray path. The internal details of the body along the ray path are compressed onto a single point on the film or a single measurement. Extending the same argument to all ray paths,

**FIGURE 1.9**

An X-ray image or a typical radiograph is a 2D projection or planar image of a 3D object. The entire object is irradiated with X rays. The projection of a 2D cross-sectional plane PQRS of the object is a 1D profile P'Q' of the 2D planar image. See also Figures 1.19 and 9.1. Reproduced, with permission, from R.M. Rangayyan and A. Kantzas, "Image reconstruction", *Wiley Encyclopedia of Electrical and Electronics Engineering*, Supplement 1, Editor: John G. Webster, Wiley, New York, NY, pp 249–268, 2000. © This material is used by permission of John Wiley & Sons, Inc.

we see that the radiographic image so produced is a 2D planar image of the 3D object, where the internal details are superimposed. In the case that the rays are parallel to the x axis (as in Figure 1.9), we have $\theta = 90^\circ$, $s = -x$, $ds = -dx$, and the planar image

$$g(y, z) = \int -\mu(x, y, z) dx. \quad (1.3)$$

Ignoring the negative sign, we see that the 3D object is reduced to (or integrated into) a 2D planar image by the process of radiographic imaging.

The most commonly used detector in X-ray imaging is the screen-film combination [5, 6]. The X rays exiting from the body being imaged strike a fluorescent (phosphor) screen made of compounds of rare-earth elements such as lanthanum oxybromide or gadolinium oxysulfide, where the X-ray photons are converted into visible-light photons. A light-sensitive film that is placed in contact with the screen (in a light-tight cassette) records the result. The film contains a layer of silver-halide emulsion with a thickness of about $10 \mu\text{m}$. The exposure or blackening of the film depends upon the number of light photons that reach the film.

A thick screen provides a high efficiency of conversion of X rays to light, but causes loss of resolution due to blurring (see Figure 1.10). The typical thickness of the phosphor layer in screens is in the range $40 - 100 \mu\text{m}$. Some

receiving units make use of a film with emulsion on both sides that is sandwiched between two screens: the second screen (located after the film along the path of propagation of the X rays) converts the X-ray photons not affected by the first screen into light, and thereby increases the efficiency of the receiver. Thin screens may be used in such dual-screen systems to achieve higher conversion efficiency (and lower dose to the patient) without sacrificing resolution.




FIGURE 1.10

Blur caused by a thick screen. Light emanating from point A in the screen is spread over a larger area on the film than that from point B.

A fluoroscopy system uses an image intensifier and a video camera in place of the film to capture the image and display it on a monitor as a movie or video [5, 6]. Images are acquired at a rate of 2 – 8 *frames/s* (*fps*), with the X-ray beam pulsed at 30 – 100 *ms* per frame. In computed radiography (CR), a photo-stimulable phosphor plate (made of europium-activated barium fluorohalide) is used instead of film to capture and temporarily hold the image pattern. The latent image pattern is then scanned using a laser and digitized. In digital radiography (DR), the film or the entire screen-film combination is replaced with solid-state electronic detectors [39, 40, 41, 42].

Examples: Figures 1.11 (a) and (b) show the posterior-anterior (PA, that is, back-to-front) and lateral (side-to-side) X-ray images of the chest of a patient. Details of the ribs and lungs, as well as the outline of the heart, are visible in the images. Images of this type are useful in visualizing and discriminating between the air-filled lungs, the fluid-filled heart, the ribs, and vessels. The size of the heart may be assessed in order to detect enlargement of the heart. The images may be used to detect lesions in the lungs and fracture of the ribs or the spinal column, and to exclude the presence of fluid in the thoracic cage. The use of two views assists in localizing lesions: use of the PA view only, for example, will not provide information to decide if a tumor is located toward the posterior or anterior of the patient.



(a)



(b)

FIGURE 1.11

(a) Posterior-anterior and (b) lateral chest X-ray images of a patient. Images courtesy of Foothills Hospital, Calgary.

The following paragraphs describe some of the physical and technical considerations in X-ray imaging [4, 5, 6, 43, 44].

- **Target and focal spot:** An electron beam with energy in the range of $20 - 140 \text{ keV}$ is used to produce X rays for diagnostic imaging. The typical target materials used are tungsten and molybdenum. The term “focal spot” refers to the area of the target struck by the electron beam to generate X rays; however, the nominal focal spot is typically expressed in terms of its diameter in mm as observed in the imaging plane (on the film). A small focal spot is desired in order to obtain a sharp image, especially in magnification imaging. (See also Section 2.9 and Figure 2.18.) Typical focal spot sizes in radiography lie in the range of $0.1 - 2 \text{ mm}$. A focal spot size of $0.1 - 0.3 \text{ mm}$ is desired in mammography.
- **Energy:** The penetrating capability of an X-ray beam is mainly determined by the accelerating voltage applied to the electron beam that impinges the target in the X-ray generator. The commonly used indicator of penetrating capability (often referred to as the “energy” of the X-ray beam) is kVp , standing for kilo-volt-peak. The higher the kVp , the more penetrating the X-ray beam will be. The actual unit of energy of an X-ray photon is the electron volt or eV , which is the energy gained by an electron when a potential of 1 V is applied to it. The kVp measure relates to the highest possible X-ray photon energy that may be achieved at the voltage used.

Low-energy X-ray photons are absorbed at or near the skin surface, and do not contribute to the image. In order to prevent such unwanted radiation, a filter is used at the X-ray source to absorb low-energy X rays. Typical filter materials are aluminum and molybdenum.

Imaging of soft-tissue organs such as the breast is performed with low-energy X rays in the range of $25 - 32 \text{ kVp}$ [45]. The use of a higher kVp would result in low differential attenuation and poor tissue-detail visibility or contrast. A few other energy levels used in projection radiography are, for imaging the abdomen: $60 - 100 \text{ kVp}$; chest: $80 - 120 \text{ kVp}$; and skull: $70 - 90 \text{ kVp}$. The kVp to be used depends upon the distance between the X-ray source and the patient, the size (thickness) of the patient, the type of grid used, and several other factors.

- **Exposure:** For a given tube voltage (kVp), the total number of X-ray photons released at the source is related to the product of the tube current (mA) and the exposure time (s), together expressed as the product mAs . As a result, for a given body being imaged, the number of photons that arrive at the film is also related to the mAs quantity. A low mAs results in an under-exposed film (faint or light image), whereas a high mAs results in an over-exposed or dark image (as well as increased

X-ray dose to the patient). Typical exposure values lie in the range of 2 – 120 *mAs*. Most imaging systems determine automatically the required exposure for a given mode of imaging, patient size, and *kVp* setting. Some systems use an initial exposure of the order of 5 *ms* to estimate the penetration of the X rays through the body being imaged, and then determine the required exposure.

- **Beam hardening:** The X rays used in radiographic imaging are typically not monoenergetic; that is, they possess X-ray photons over a certain band of frequencies or EM energy levels. As the X rays propagate through a body, the lower-energy photons get absorbed preferentially, depending upon the length of the ray path through the body and the attenuation characteristics of the tissues along the path. Thus, the X rays that pass through the object at longer distances from the source will possess relatively fewer photons at lower-energy levels than at the point of entry into the object (and hence a relatively higher concentration of higher-energy photons). This effect is known as beam hardening, and leads to incorrect estimation of the attenuation coefficient in computed tomography (CT) imaging. The effect of beam hardening may be reduced by prefiltering or prehardening the X-ray beam and narrowing its spectrum. The use of monoenergetic X rays from a synchrotron or a laser obviates this problem.
- **Scatter and the use of grids:** As an X-ray beam propagates through a body, photons are lost due to absorption and scattering at each point in the body. The angle of the scattered photon at a given point along the incoming beam is a random variable, and hence the scattered photon contributes to noise at the point where it strikes the detector. Furthermore, scattering results in the loss of contrast of the part of the object where X-ray photons were scattered from the main beam. The noise effect of the scattered radiation is significant in gamma-ray emission imaging, and requires specific methods to improve the quality of the image [4, 46]. The effect of scatter may be reduced by the use of grids, collimation, or energy discrimination due to the fact that the scattered (or secondary) photons usually have lower energy levels than the primary photons.

A grid consists of an array of X-ray absorbing strips that are mutually parallel if the X rays are in a parallel beam, as in chest imaging (see Figures 1.12 and 1.13), or are converging toward the X-ray source in the case of a diverging beam (as in breast imaging, see Figure 1.15). Lattice or honeycomb grids with parallel strips in criss-cross patterns are also used in mammography. X-ray photons that arrive via a path that is not aligned with the grids will be stopped from reaching the detector.

A typical grid contains thin strips of lead or aluminum with a strip density of 25 – 80 *lines/cm* and a grid height:strip width ratio in the range

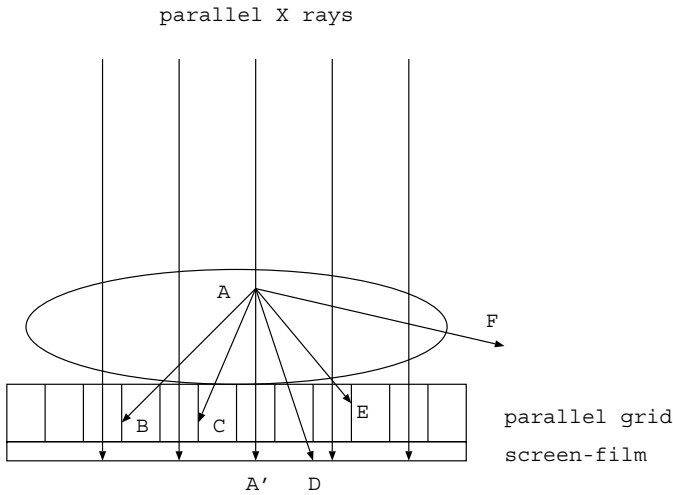


FIGURE 1.12

Use of parallel grids to reduce scatter. X rays that are parallel to the grids reach the film; for example, line AA'. Scattered rays AB, AC, and AE have been blocked by the grids; however, the scattered ray AD has reached the film in the illustration.

of 5:1 to 12:1. The space between the grids is filled with low-attenuation material such as wood. A stationary grid produces a line pattern that is superimposed upon the image, which would be distracting. Figure 1.13 (a) shows a part of an image of a phantom with the grid artifact clearly visible. (An image of the complete phantom is shown in Figure 1.14.) Grid artifact is prevented in a reciprocating grid, where the grid is moved about 20 grid spacings during exposure: the movement smears the grid shadow and renders it invisible on the image. Figure 1.13 (b) shows an image of the same object as in part (a), but with no grid artifact. Low levels of grid artifact may appear in images if the bucky that holds the grid does not move at a uniform pace or starts moving late or ends movement early with respect to the X-ray exposure interval. A major disadvantage of using grids is that it requires approximately two times the radiation dose required for imaging techniques without grids. Furthermore, the contrast of fine details is reduced due to the smeared shadow of the grid.

- **Photon detection noise:** The interaction between an X-ray beam and a detector is governed by the same rules as for interaction with any other matter: photons are lost due to scatter and absorption, and some photons may pass through unaffected (or undetected). The small size of the detectors in DR and CT imaging reduces their detection

efficiency. Scattered and undetected photons cause noise in the measurement; for detailed analysis of noise in X-ray detection, refer to Barrett and Swindell [3], Macovski [5], and Cho et al. [4]. More details on noise in medical images and techniques to remove noise are presented in Chapter 3.

- **Ray stopping by heavy implants:** If the body being imaged contains extremely heavy parts or components, such as metal screws or pins in bones and surgical clips that are nearly X-ray-opaque and entirely stop the incoming X-ray photons, no photons would be detected at the corresponding point of exit from the body. The attenuation coefficient for the corresponding path would be indefinite, or within the computational context, infinity. Then, a reconstruction algorithm would not be able to redistribute the attenuation values over the points along the corresponding ray path in the reconstructed image. This leads to streaking artifacts in CT images.


Two special techniques for enhanced X-ray imaging — digital subtraction angiography (DSA) and dual-energy imaging — are described in Sections 4.1 and 4.2, respectively.

1.5.1 Breast cancer and mammography

Breast cancer: Cancer is caused when a single cell or a group of cells escapes from the usual controls that regulate cellular growth, and begins to multiply and spread. This activity results in a mass, tumor, or neoplasm. Many masses are benign; that is, the abnormal growth is restricted to a single, circumscribed, expanding mass of cells. Some tumors are malignant; that is, the abnormal growth invades the surrounding tissues and may spread, or metastasize, to distant areas of the body. Although benign masses may lead to complications, malignant tumors are usually more serious, and it is for these tumors that the term “cancer” is used. The majority of breast tumors will have metastasized before reaching a palpable size.

Although curable, especially when detected at early stages, breast cancer is a major cause of death in women. An important factor in breast cancer is that it tends to occur earlier in life than other types of cancer and other major diseases [47, 48]. Although the cause of breast cancer has not yet been fully understood, early detection and removal of the primary tumor are essential and effective methods to reduce mortality, because, at such a point in time, only a few of the cells that departed from the primary tumor would have succeeded in forming secondary tumors [49]. When breast tumors are detected by the affected women themselves (via self-examination), most of the tumors would have metastasized [50].

If breast cancer can be detected by some means at an early stage, while it is clinically localized, the survival rate can be dramatically increased. However,



(a)



(b)

FIGURE 1.13

X-ray images of a part of a phantom: (a) with, and (b) without grid artifact. Image courtesy of L.J. Hahn, Foothills Hospital, Calgary. See also Figure 1.14.




FIGURE 1.14

X-ray image of the American College of Radiology (ACR) phantom for mammography. The pixel-value range [117, 210] has been linearly stretched to the display range [0, 255] to show the details. Image courtesy of S. Bright, Sunnybrook & Women's College Health Sciences Centre, Toronto, ON, Canada. See also Figure 1.13.

such early breast cancer is generally not amenable to detection by physical examination and breast self-examination. The primary role of an imaging technique is thus the detection of lesions in the breast [29]. Currently, the most effective method for the detection of early breast cancer is X-ray mammography. Other modalities, such as ultrasonography, transillumination, thermography, CT, and magnetic resonance imaging (MRI) have been investigated for breast cancer diagnosis, but mammography is the only reliable procedure for detecting nonpalpable cancers and for detecting many minimal breast cancers when they appear to be curable [18, 28, 29, 51]. Therefore, mammography has been recommended for periodic screening of asymptomatic women. Mammography has gained recognition as the single most successful technique for the detection of early, clinically occult breast cancer [52, 53, 54, 55, 56].

X-ray imaging of the breast: The technique of using X rays to obtain images of the breast was first reported by Warren in 1930, after he had examined 100 women using sagittal views [57]. Because of the lack of a reproducible method for obtaining satisfactory images, this technique did not make much progress until 1960, when Egan [58] reported on high- mA and low- kVp X-ray sources that yielded reproducible images on industrial film. It was in the mid-1960s that the first modern X-ray unit dedicated to mammography was developed. Since then, remarkable advances have led to a striking improvement in image quality and a dramatic reduction in radiation dose.

A major characteristic of mammograms is low contrast, which is due to the relatively homogeneous soft-tissue composition of the breast. Many efforts have been focused on developing methods to enhance contrast. In an alternative imaging method known as xeromammography, a selenium-coated aluminum plate is used as the detector [6]. The plate is initially charged to about 1,000 V. Exposure to the X rays exiting the patient creates a charge pattern on the plate due to the liberation of electrons and ions. The plate is then sprayed with an ionized toner, the pattern of which is transferred to plastic-coated paper. Xeromammograms provide wide latitude and edge enhancement, which lead to improved images as compared to screen-film mammography. However, xeromammography results in a higher dose to the subject, and has not been in much use since the 1980s.

A typical mammographic imaging system is shown schematically in Figure 1.15. Mammography requires high X-ray beam quality (a narrow-band or nearly monochromatic beam), which is controlled by the tube target material (molybdenum) and beam filtration with molybdenum. Effective breast compression is an important factor in reducing scattered radiation, creating as uniform a density distribution as possible, eliminating motion, and separating mammary structures, thereby increasing the visibility of details in the image. The use of grids specifically designed for mammography can further reduce scattered radiation and improve subject contrast, which is especially significant when imaging thick, dense breasts [59].

Generally, conventional screen-film mammography is performed with the breast directly in contact with the screen-film cassette, producing essentially




FIGURE 1.15
A typical mammography setup.

life-size images. The magnification technique, on the other hand, interposes an air gap between the breast and the film, so that the projected radiographic image is enlarged. Magnification produces fine-detail images containing additional anatomical information that may be useful in refining mammographic diagnosis, especially in cases where conventional imaging demonstrates uncertain or equivocal findings [60]. As in the grid method, the advantages of magnification imaging are achieved at the expense of increased radiation exposure. Therefore, the magnification technique is not used routinely.

Screen-film mammography is now the main tool for the detection of early breast cancer. The risk of radiation is still a matter of concern, although there is no direct evidence of breast cancer risk from the low-dose radiation exposure of mammography. Regardless, technological advances in mammography continue to be directed toward minimizing radiation exposure while maintaining the high quality of the images.

Examples: Figures 1.16 (a) and (b) show the cranio-caudal (CC) and medio-lateral-oblique (MLO) views of the same breast of a subject. The MLO view demonstrates architectural distortion due to a spiculated tumor near the upper right-hand corner edge.

Mammograms are analyzed by radiologists specialized in mammography. A normal mammogram typically depicts converging patterns of fibroglandular tissues and vessels. Any feature that causes a departure from or distortion with reference to the normal pattern is viewed with suspicion and analyzed with extra attention. Features such as calcifications, masses, localized increase in density, architectural distortion, and asymmetry between the left and right breast images are carefully analyzed.

Several countries and states have instituted breast cancer screening programs where asymptomatic women within a certain age group are invited to participate in regular mammographic examinations. Screen Test — Alberta Program for the Early Detection of Breast Cancer [61] is an example of such a program. Several applications of image processing and pattern analysis techniques for mammographic image analysis and breast cancer detection are described in the chapters to follow.

1.6 Tomography

The problem of visualizing the details of the interior of the human body noninvasively has always been of interest, and within a few years after the discovery of X rays by Röntgen in 1895, techniques were developed to image sectional planes of the body. The techniques of laminagraphy, planigraphy, or “classical” tomography [38, 62] used synchronous movement of the X-ray source and film in such a way as to produce a relatively sharp image of a




FIGURE 1.16

(a) Cranio-caudal (CC) and (b) medio-lateral oblique (MLO) mammograms of the same breast of a subject. Images courtesy of Screen Test — Alberta Program for the Early Detection of Breast Cancer [61].

single focal plane of the object, with the images of all other planes being blurred. Figure 1.17 illustrates a simple linear-motion system, where the X-ray source and film cassette move along straight-line paths so as to maintain the longitudinal (coronal) plane, indicated by the straight line AB, in focus. It is seen that the X rays along the paths X1-A and X2-A strike the same physical spot A1 = A2 on the film, and that the rays along the paths X1-B and X2-B strike the same spot B1 = B2. On the other hand, for the point C in a different plane, the rays along the paths X1-C and X2-C strike different points C1 ≠ C2 on the film. Therefore, the details in the plane AB remain in focus and cause a strong image, whereas the details in the other planes are smeared all over the film. The smearing of information from the other planes of the object causes loss of contrast in the plane of interest. The development of CT imaging rendered film-based tomography obsolete.

Example: Figure 1.18 shows a tomographic image of a patient in a longitudinal (coronal) plane through the chest. Images of this nature provided better visualization and localization of lesions than regular X-ray projection images, and permitted the detection of masses in bronchial tubes and air ducts.

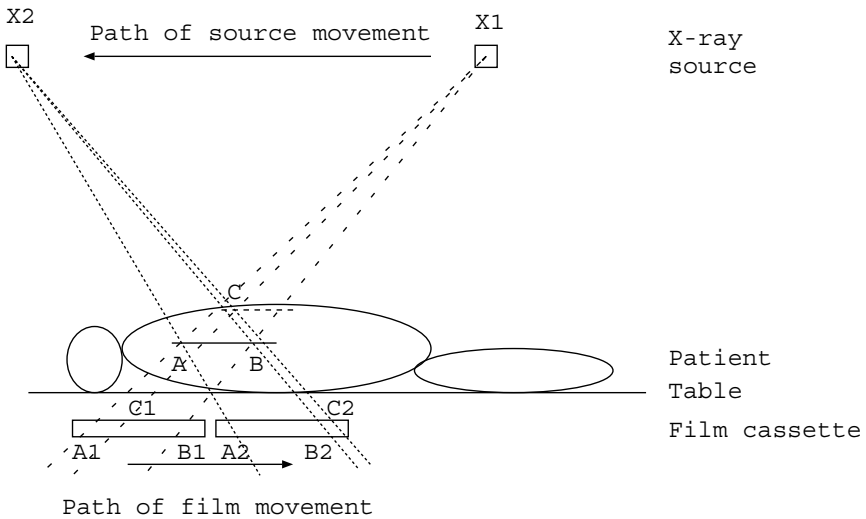


FIGURE 1.17

Synchronized movement of the X-ray source and film to obtain a tomographic image of the focal plane indicated as AB. Adapted from Robb [38].

Computed tomography: The technique of CT imaging was developed during the late 1960s and the early 1970s, producing images of cross-sections of the human head and body as never seen before (noninvasively and nondestructively).



FIGURE 1.18

Tomographic image of a patient in a longitudinal sectional plane through the chest. Reproduced with permission from R.A. Robb, "X-ray computed tomography: An engineering synthesis of multiscientific principles", *CRC Critical Reviews in Biomedical Engineering*, 7:264-333, March 1982. © CRC Press.

tively!). In the simplest form of CT imaging, only the desired cross-sectional plane of the body is irradiated using a finely collimated ray of X-ray photons (see Figure 1.19), instead of irradiating the entire body with a 3D beam of X rays as in ordinary radiography (Figure 1.9). The fundamental radiographic equation for CT is the same as Equation 1.2. Ray integrals are measured at many positions and angles around the body, scanning the body in the process. The principle of image reconstruction from projections, described in detail in Chapter 9, is then used to compute an image of a section of the body: hence the name *computed* tomography. (See Robb [38] for an excellent review of the history of CT imaging; see also Rangayyan and Kantzas [63] and Rangayyan [64].)

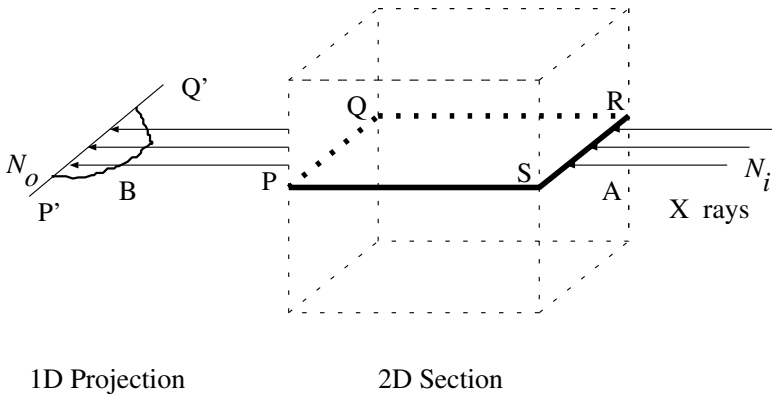


FIGURE 1.19

In the basic form of CT imaging, only the cross-sectional plane of interest is irradiated with X rays. The projection of the 2D cross-sectional plane PQRS of the object is the 1D profile $P'Q'$ shown. Compare this case with the planar imaging case illustrated in Figure 1.9. See also Figure 9.1. Reproduced, with permission, from R.M. Rangayyan and A. Kantzas, “Image reconstruction”, *Wiley Encyclopedia of Electrical and Electronics Engineering*, Supplement 1, Editor: John G. Webster, Wiley, New York, NY, pp 249–268, 2000. © This material is used by permission of John Wiley & Sons, Inc.

Figure 1.20 depicts some of the scanning procedures employed: Figure 1.20 (a) shows the translate-rotate scanning geometry for parallel-ray projections; Figure 1.20 (b) shows the translate-rotate scanning geometry with a small fan-beam detector array; Figure 1.20 (c) shows the rotate-only scanning geometry for fan-beam projections; and Figure 1.20 (d) shows the rotate-only scanning geometry for fan-beam projections using a ring of detectors. A more recently developed scanner specialized for cardiovascular imaging [65, 66] completely eliminates mechanical scanning movement to reduce the scanning time by

employing electronically steered X-ray microbeams and rings of detectors, as illustrated schematically in Figure 1.21.

The fundamental principle behind CT, namely, image reconstruction from projections, has been known for close to 100 years, since the exposition of the topic by Radon [67, 68] in 1917. More recent developments in the subject arose in the 1950s and 1960s from the works of a number of researchers in diverse applications. Some of the important publications in this area are the works of Cormack on the representation of a function by its line integrals [69, 70]; Bracewell and Riddle on the reconstruction of brightness distributions of astronomical bodies from fan-beam scans at various angles [71]; Crowther et al. [72] and De Rosier and Klug [73] on the reconstruction of 3D images of viruses from electron micrographs; Ramachandran and Lakshminarayanan [74] on the convolution backprojection technique; and Gordon et al. [75] on algebraic reconstruction techniques. Pioneering works on the development of practical scanners for medical applications were performed by Oldendorf [76], Hounsfield [77], and Ambrose [78]. X-ray CT was well established as a clinical diagnostic tool by the early 1970s.

Once a sectional image is obtained, the process may be repeated to obtain a series of sectional images of the 3D body or object being investigated. CT imaging a 3D body may be accomplished by reconstructing one 2D section at a time through the use of 1D projections. In the Dynamic Spatial Reconstructor [38, 79], 2D projection images are obtained on a fluorescent screen via irradiation of the entire portion of interest of the body. In single-photon emission computed tomography (SPECT), several 2D projection images are obtained using a gamma camera [4, 5, 46, 80, 81]. In these cases, the projection data of several sectional planes are acquired simultaneously: each row of a given 2D projection or planar image provides the 1D projection data of a sectional plane of the body imaged (see Figure 1.9). Many sectional images may then be reconstructed from the set of 2D planar images acquired.

Other imaging modalities used for projection data collection are ultrasound (time of flight or attenuation), magnetic resonance (MR), nuclear emission (gamma rays or positrons), and light [4, 5, 80, 82, 83, 84, 85]. Techniques using nonionizing radiation are of importance in imaging pregnant women. Whereas the physical parameter imaged may differ between these modalities, once the projection data are acquired, the mathematical image reconstruction procedure could be almost the same. A few special considerations in imaging with diffracting sources are described in Section 9.5. The characteristics of the data acquired in nuclear medicine, ultrasound, and MR imaging are described in Sections 1.7, 1.8, and 1.9, respectively.

Examples: Figure 1.22 shows a CT scan of the head of a patient. The image displays, among other features, the ventricles of the brain. CT images of the head are useful in the detection of abnormalities in the brain and skull, including bleeding in brain masses, calcifications, and fractures in the cranial vault.

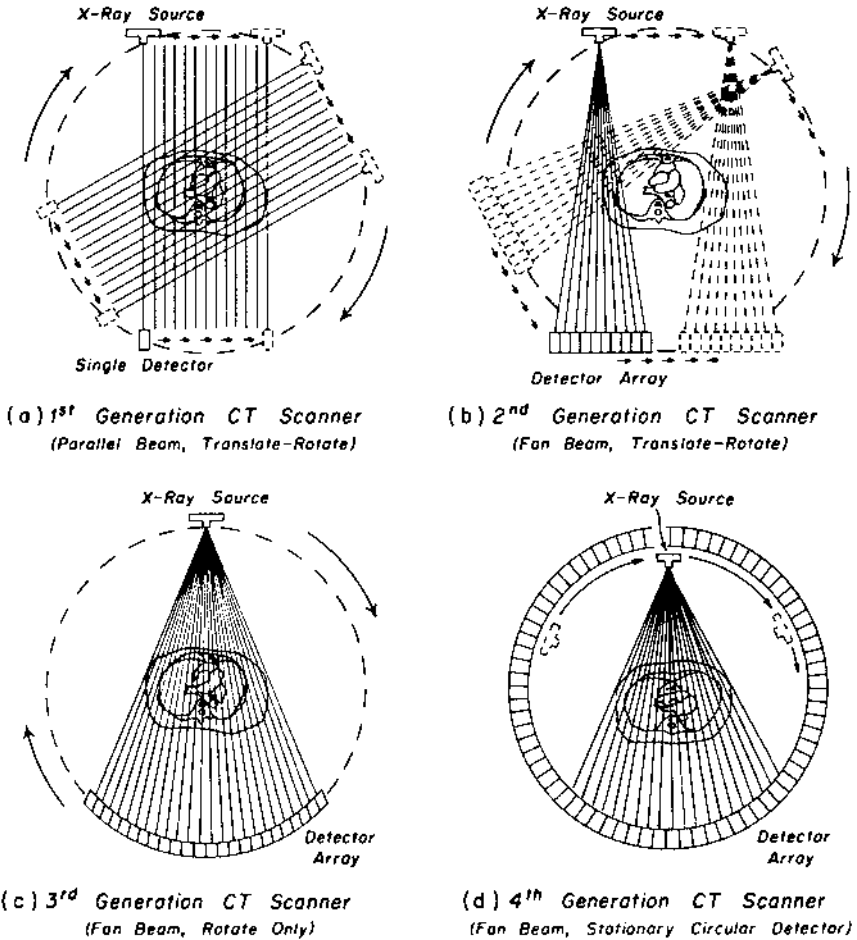


FIGURE 1.20

(a) Translate-rotate scanning geometry for parallel-ray projections; (b) translate-rotate scanning geometry with a small fan-beam detector array; (c) rotate-only scanning geometry for fan-beam projections; (d) rotate-only scanning geometry for fan-beam projections using a ring of detectors. Reproduced with permission from R.A. Robb, "X-ray computed tomography: An engineering synthesis of multiscientific principles", *CRC Critical Reviews in Biomedical Engineering*, 7:264-333, March 1982. © CRC Press.





FIGURE 1.21

Electronic steering of an X-ray beam for motion-free scanning and CT imaging. Reproduced with permission from D.P. Boyd, R.G. Gould, J.R. Quinn, and R. Sparks, "Proposed dynamic cardiac 3-D densitometer for early detection and evaluation of heart disease", *IEEE Transactions on Nuclear Science*, 26(2):2724-2727, 1979. © IEEE. See also Robb [38].

**FIGURE 1.22**

CT image of a patient showing the details in a cross-section through the head (brain). Image courtesy of Foothills Hospital, Calgary.

Figure 1.23 illustrates the CT image of the abdomen of a patient. The image shows the stomach, gall bladder, liver, spleen, kidneys, intestines, and the spinal column. The air-fluid interface in the stomach is clearly visible. Images of this type are useful in detecting several abnormalities in the abdomen, including gallstones, kidney stones, and tumors in the liver.

Figure 1.24 shows two renditions of the same CT image of the chest of a patient: image (a) has been scaled (or windowed; details provided in Sections 4.4.2 and 9.6) to illustrate the details of the lungs, whereas image (b) has been scaled to display the mediastinum in relatively increased detail. Image (a) shows the details of the distal branches of the pulmonary arteries. CT images of the chest facilitate the detection of the distortion of anatomy due to intrathoracic or extrapulmonary fluid collection, or due to ruptured lungs. They also aid in the detection of lung tumors and blockage of the pulmonary arteries due to thrombi.

CT is an imaging technique that has revolutionized the field of medical diagnostics. CT has also found applications in many other areas such as non-destructive evaluation of industrial and biological specimens, radioastronomy, light and electron microscopy, optical interferometry, X-ray crystallography, petroleum engineering, and geophysical exploration. Indirectly, it has also led to new developments in its predecessor techniques in radiographic imag-



FIGURE 1.23

CT image of a patient showing the details in a cross-section through the abdomen. Image courtesy of Foothills Hospital, Calgary.

ing. Details of the mathematical principles related to reconstruction from projections and more illustrations of CT imaging are provided in Chapter 9.

1.7 Nuclear Medicine Imaging

The use of radioactivity in medical imaging began in the 1950s; nuclear medicine has now become an integral part of most medical imaging centers [3, 4, 5, 6, 86]. In nuclear medicine imaging, a small quantity of a radiopharmaceutical is administered into the body orally, by intravenous injection, or by inhalation. The radiopharmaceutical is designed so as to be absorbed by and localized in a specific organ of interest. The gamma-ray photons emitted from the body as a result of radioactive decay of the radiopharmaceutical are used to form an image that represents the distribution of radioactivity in the organ.

Nuclear medicine imaging is used to map physiological function such as perfusion and ventilation of the lungs, and blood supply to the musculature of the heart, liver, spleen, and thyroid gland. Nuclear medicine has also proven to be useful in detecting brain and bone tumors. Whereas X-ray images



(a)



(b)

FIGURE 1.24

CT image of a patient scaled to (a) show the details of the lungs; and (b) display the mediastinum in detail — the details of the lungs are not visible in this rendition. Images courtesy of Alberta Children’s Hospital, Calgary.

provide information related to density and may be used to detect altered anatomy, nuclear medicine imaging helps in examining altered physiological (or pathological) functioning of specific organs in a body.

The most commonly used isotopes in nuclear medicine imaging are technetium as ^{99m}Tc which emits gamma-ray photons at 140 keV , and thallium as ^{201}Tl at 70 keV or 167 keV . Iodine as ^{131}I is also used for thyroid imaging.

The first imaging device used in nuclear medicine was the rectilinear scanner, which consisted of a single-bore collimator connected to a gamma-ray counter or detector. The scanner was coupled to a mechanical system that performed a raster scan over the area of interest, making a map of the radiation distribution in the area. The amount of radioactivity detected at each position was either recorded on film or on a storage oscilloscope. A major difficulty with this approach is that scanning is time consuming.

The scintillation gamma camera or the Anger camera uses a large thallium-activated sodium iodide [$\text{NaI}(\text{Tl})$] detector, typically 40 cm in diameter and 10 mm in thickness. The gamma camera consists of three major parts: a collimator, a detector, and a set of PMTs. Figure 1.25 illustrates the Anger camera in a schematic sectional view. The following paragraphs describe some of the important components of a nuclear medicine imaging system.

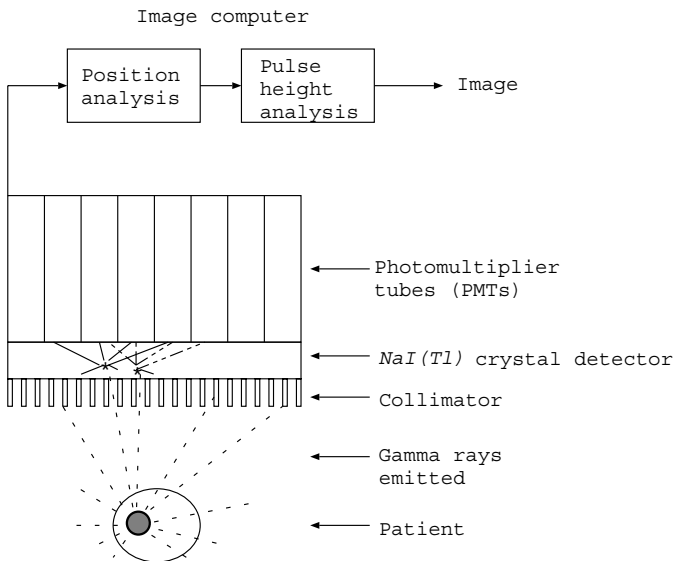


FIGURE 1.25

Schematic (vertical sectional) representation of a nuclear medicine imaging system with an Anger camera.

- **Collimator:** A collimator consists of an array of holes separated by lead septa. The function of the collimator is to allow passage of the gamma rays that arrive along a certain path of propagation, and to block (absorb) all gamma rays outside a narrow solid angle of acceptance. Collimators are usually made of lead alloys, but other materials such as tantalum, tungsten, and gold have also been used. Different geometries have been used in the design of collimator holes, including triangular, square, hexagonal, and round patterns. Hexagonal holes have been observed to be the most efficient, and are commonly used.

Two key factors in collimator design are geometric efficiency, which is the fraction of the gamma-ray photons from the source that are transmitted through the collimator to the detector, and geometric (spatial) resolution. In general, for a given type of collimator, the higher the efficiency, the poorer is the resolution. The resolution of a collimator is increased if the size of the holes is reduced or if the collimator thickness is increased; however, these measures decrease the number of photons that will reach the crystal, and hence reduce the sensitivity and efficiency of the system. The efficiency of a typical collimator is about 0.01%; that is, only 1 in every 10,000 photons emitted is passed by the collimator to the crystal.

The most commonly used type of collimator is the parallel-hole collimator, which usually serves for general purpose imaging, particularly for large organs. Other designs include diverging, converging, fan-beam, and pin-hole collimators.

- **Detector:** At the back of the collimator is attached a detector, which is usually a $NaI(Tl)$ crystal of 6 – 13 mm thickness. The crystal absorbs the gamma-ray photons that pass through the collimator holes, and reemits their energy as visible light (scintillation). The thickness of the crystal determines the absorbed fraction of the gamma-ray photons by the photoelectric effect. A thick crystal has better absorption than a thin crystal; however, a thick crystal scatters and absorbs the light before it reaches the back surface of the crystal. A crystal of thickness 10 mm absorbs about 92% of the photons received at 140 keV.
- **Photomultiplier tubes:** The crystal is optically coupled at its back surface to an array of PMTs. The PMTs are usually hexagonal in section, and are arranged so as to cover the imaging area. Scintillations within the crystal are converted by the photocathodes at the front of the PMTs to photoelectrons, which are accelerated toward each of a series of dynodes held at successively higher potentials until they reach the anode at the back of the tube. During this process, the photoelectrons produce a number of secondary electrons at each dynode, leading to a current gain of the order of 10^6 .

- **Image computer:** The current pulses produced by the PMTs in response to scintillations in the crystal are applied to a resistor matrix that computes the points of arrival of the corresponding gamma-ray photons. The amplitudes of the pulses represent the energy deposited by the gamma rays. A pulse-height analyzer is used to select pulses that are within a preset energy window corresponding to the peak energy of the gamma rays. The pulse-selection step reduces the effect of scattered rays at energy levels outside the energy window used.

Whereas the major advantage of nuclear medicine imaging lies in its capability of imaging the functional aspects of the human body rather than the anatomy, its major disadvantages are poor spatial resolution and high noise content. The intrinsic resolution of a typical gamma camera (crystal) is 3–5 mm; the net resolution including the effect of the collimator, expressed as the full width at half the maximum (FWHM) of the image of a thin line source (the line spread function or LSF) is 7.5–10 mm (see Figure 2.21). The main causes of noise are quantum mottle due to the low number of photons used to create images, and the random nature of gamma ray emission. Structured noise may also be caused by nonuniformities in the gamma camera.

In general, the contrast of nuclear medicine images is high as the radiopharmaceutical is designed so as to be selectively absorbed by and localized in the organ of interest. Regardless, other organs and tissues in the body will also absorb some amounts of the radiopharmaceutical, and emit gamma rays that will appear as background counts and degrade contrast. Such effects may be labeled as physiological or anatomical artifacts. The contrast of an image is also diminished by septal penetration in the collimator and scatter.

Multi-camera imaging and scanning systems: The data acquired by a gamma camera represent a projection or planar image, which corresponds to an integration of the 3D body being imaged along the paths of the gamma rays. It should be noted that gamma rays are emitted by the body in all directions during the imaging procedure. Modern imaging systems use two or three gamma cameras (“heads”) to acquire simultaneously multiple projection images. Projection images acquired at diametrically opposite positions may be averaged to reduce artifacts [87]; see Section 10.5.5.

SPECT imaging: SPECT scanners usually gather 64 or 128 projections spanning 180° or 360° around the patient, depending upon the application. Individual scan lines from the projection images may then be processed through a reconstruction algorithm to obtain 2D sectional images, which may further be combined to create a 3D image of the patient. Coronal, sagittal, and oblique sections may then be created from the 3D dataset. Circular scanning is commonly used to acquire projection images of the body at different angles. Circular scanning provides projections at equal angular intervals, as required by certain reconstruction algorithms. However, some clinical studies use elliptical scanning so as to keep the camera close to the body, in consideration of the fact that the spatial resolution deteriorates rapidly with distance. In

such situations, the CT reconstruction algorithm should be adapted to the nonuniformly spaced data.

Examples: Figure 1.26 illustrates a nuclear medicine projection (planar) image of the chest of a patient. The region of high intensity (activity) on the right-hand side of the image represents the heart. Observe the high level of noise in the image. Figure 1.27 illustrates several series of SPECT images of the left ventricle of the patient before and after stress (exercise on a treadmill). The SPECT images display oblique sectional views of the myocardium in three orientations: the short axis, the horizontal long axis, and the vertical long axis of the left ventricle. Ischemic regions demonstrate lower intensity than normal regions due to reduced blood supply. The generally noisy appearance of SPECT images as well as the nature of the artifacts in nuclear medicine images are illustrated by the images.

See Section 3.10 for details regarding gated blood-pool imaging; see Section 10.5 for several examples of SPECT images, and for discussions on the restoration of SPECT images.






FIGURE 1.26

A planar or projection image of a patient used for myocardial SPECT imaging. The two horizontal lines indicate the limits of the data used to reconstruct the SPECT images shown in Figure 1.27. Image courtesy of Foothills Hospital, Calgary.


Positron emission tomography (PET): Certain isotopes of carbon (^{11}C), nitrogen (^{13}N), oxygen (^{15}O), and fluorine (^{18}F) emit positrons and are suitable for nuclear medicine imaging. PET is based upon the simultaneous detection of the two annihilation photons produced at 511 keV and emitted in opposite directions when a positron loses its kinetic energy and combines with an electron [4, 5, 6, 88]. The process is also known as coincidence detection. Collimation in PET imaging is electronic, which substantially increases



(a)



(b)



(c)

FIGURE 1.27

SPECT imaging of the left ventricle. (a) Short-axis images. (b) Horizontal long axis images. (c) Vertical long axis images. In each case, the upper panel shows four SPECT images after exercise (stress), and the lower panel shows the corresponding views before exercise (rest). Images courtesy of Foothills Hospital, Calgary.

the efficiency of the detection process as compared to SPECT imaging. A diaphragm is used only to define the plane of sectional (CT) imaging.

Several modes of data collection are possible for PET imaging, including stationary, rotate-only, and rotate-translate gantries [88]. In one mode of data collection, a ring of bismuth-germanate detectors is used to gather emission statistics that correspond to a projection of a transversal section. A reconstruction algorithm is then used to create 2D and 3D images.

The spatial resolution of PET imaging is typically 5 *mm*, which is almost two times better than that of SPECT imaging. PET is useful in functional imaging and physiological analysis of organs [89, 90].

1.8 Ultrasonography

Ultrasound in the frequency range of 1 – 20 *MHz* is used in diagnostic ultrasonography [4, 5, 6]. The velocity of propagation of ultrasound through a medium depends upon its compressibility: lower compressibility results in higher velocity. Typical velocities in human tissues are 330 *m/s* in air (the lungs); 1,540 *m/s* in soft tissue; and 3,300 *m/s* in bone. A wave of ultrasound may get reflected, refracted, scattered, or absorbed as it propagates through a body. Most modes of diagnostic ultrasonography are based upon the reflection of ultrasound at tissue interfaces. A gel is used to minimize the presence of air between the transducer and the skin in order to avoid reflection at the skin surface. Typically, pulses of ultrasound of about 1 μ s duration at a repetition rate of about 1,000 *pps* (pulses per second) are applied, and the resulting echoes are used for locating tissue interfaces and imaging.

Large, smooth surfaces in a body cause specular reflection, whereas rough surfaces and regions cause nonspecular reflection or diffuse scatter. The normal liver, for example, is made up of clusters of parenchyma that are of the order of 2 *mm* in size. Considering an ultrasound signal at 1 *MHz* and assuming a propagation velocity of 1,540 *m/s*, we get the wavelength to be 1.54 *mm*, which is of the order of the size of the parenchymal clusters. For this reason, ultrasound is scattered in all directions by the liver, which appears with a speckled texture in ultrasound scans [3, 6]. Fluid-filled regions such as cysts have no internal structure, generate no echoes except at their boundaries, and appear as black regions on ultrasound images. The almost-complete absorption of ultrasound by bone causes shadowing in images: tissues and organs past bones and dense objects along the path of propagation of the beam are not imaged in full and accurate detail. The quality of ultrasonographic images is also affected by multiple reflections, speckle noise due to scattering, and spatial distortion due to refraction. The spatial resolution in ultrasound images is limited to the order of 0.5 – 3 *mm*.

Some of the commonly used modes of ultrasonography are briefly described below.

- **A mode:** A single transducer is used in this mode. The amplitude (hence the name “A”) of the echoes received is displayed on the vertical axis, with the corresponding depth (related to the time of arrival of the echo) being displayed on the horizontal axis. The A mode is useful in distance measurement (ranging), with applications in the detection of retinal detachment and the detection of shift of the midline of the brain.
- **M mode:** This mode produces a display with time on the horizontal axis and echo depth on the vertical axis. The M mode is useful in the study of movement or motion (hence the name), with applications in cardiac valve motion analysis.
- **B mode:** An image of a 2D section or slice of the body is produced in this mode by using a single transducer to scan the region of interest or by using an array of sequentially activated transducers. Real-time imaging is possible in the B mode with 15 – 40 *fps*. The B mode is useful in studying large organs, such as the liver, and in fetal imaging.
- **Doppler ultrasound:** This mode is based upon the change in frequency of the investigating beam caused by a moving target (the Doppler effect), and is useful in imaging blood flow. A particular application lies in the detection of turbulence and retrograde flow, which is useful in the diagnosis of stenosis or insufficiency of cardiac valves and plaques in blood vessels [91]. Doppler imaging may be used to obtain a combination of anatomic information with B-mode imaging and flow information obtained using pulsed Doppler.
- **Special probes:** A variety of probes have been developed for ultrasonography of specific organs and for special applications, some of which are endovaginal probes for fetal imaging, transrectal probes for imaging the prostate [92], transesophageal probes for imaging the heart via the esophagus, and intravascular probes for the study of blood vessels.

Examples: Echocardiography is a major application of ultrasonography for the assessment of the functional integrity of heart valves. An array of ultrasound transducers is used in the B mode in echocardiography, so as to obtain a video illustrating the opening and closing activities of the valve leaflets. Figure 1.28 illustrates two frames of an echocardiogram of a subject with a normal mitral valve, which is captured in the open and closed positions in the two frames. Figure 1.29 shows the M-mode ultrasound image of the same subject, illustrating the valve leaflet movement against time. Echocardiography is useful in the detection of stenosis and loss of flexibility of the cardiac valves due to calcification.




FIGURE 1.28

Two frames of the echocardiogram of a subject with normal function of the mitral valve. (a) Mitral valve in the fully open position. (b) Mitral valve in the closed position. Images courtesy of Foothills Hospital, Calgary.




FIGURE 1.29

M-mode ultrasound image of a subject with normal function of the mitral valve. The horizontal axis represents time. The echo signature of the mitral valve leaflets as they open and close is illustrated. Image courtesy of Foothills Hospital, Calgary.

In spite of limitations in image quality and resolution, ultrasonography is an important medical imaging modality due to the nonionizing nature of the medium. For this reason, ultrasonography is particularly useful in fetal imaging. Figure 1.30 shows a B-mode ultrasound image of a fetus. The outline of the head and face as well as the spinal column are clearly visible in the image. Images of this nature may be used to measure the size of the head and head-to-sacrum length of the fetus. Ultrasonography is useful in the detection of abnormalities in fetal development, especially distension of the stomach, hydrocephalus, and complications due to maternal (or gestational) diabetes such as the lack of development of the sacrum.




FIGURE 1.30

B-mode ultrasound (3.5 *MHz*) image of a fetus (sagittal view). Image courtesy of Foothills Hospital, Calgary.

Ultrasonography is also useful in tomographic imaging [83, 93], discriminating between solid masses and fluid-filled cysts in the breast [53], and tissue characterization [5].

1.9 Magnetic Resonance Imaging

MRI is based on the principle of nuclear magnetic resonance (NMR): the behavior of nuclei under the influence of externally applied magnetic and EM (radio-frequency or RF) fields [4, 6, 84, 94, 95, 96]. A nucleus with an odd number of protons or an odd number of neutrons has an inherent nuclear spin and exhibits a magnetic moment; such a nucleus is said to be NMR-active. The commonly used modes of MRI rely on hydrogen (1H or proton), carbon (^{13}C), fluorine (^{19}F), and phosphorus (^{31}P) nuclei.

In the absence of an external magnetic field, the vectors of magnetic moments of active nuclei have random orientations, resulting in no net magnetism. When a strong external magnetic field H_o is applied, some of the nuclear spins of active nuclei align with the field (either parallel or antiparallel to the field). The number of spins that align with the field is a function of H_o and inversely related to the absolute temperature. At the commonly used field strength of 1.5 T (Tesla), only a relatively small fraction of spins align with the field. The axis of the magnetic field is referred to as the z axis. Parallel alignment corresponds to a lower energy state than antiparallel alignment, and hence there will be more nuclei in the former state. This state of forced alignment results in a net magnetization vector. (At equilibrium and 1.5 T, only about seven more spins in a million are aligned in the parallel state; hence, MRI is a low-sensitivity imaging technique.)

The magnetic spin vector of each active nucleus precesses about the z axis at a frequency known as the Larmor frequency, given by

$$\omega_o = \gamma H_o, \quad (1.4)$$

where γ is the gyromagnetic ratio of the nucleus considered (for protons, $\gamma = 42.57 \text{ MHz } T^{-1}$). From the view-point of classical mechanics and for 1H , this form of precession is comparable to the rotation of a spinning top's axis around the vertical.

MRI involves controlled perturbation of the precession of nuclear spins, and measurement of the RF signals emitted when the perturbation is stopped and the nuclei return to their previous state of equilibrium. MRI is an intrinsically 3D imaging procedure. The traditional CT scanner requires mechanical scanning and provides 2D cross-sectional images in a slice-by-slice manner: slices at other orientations, if required, have to be computed from a set of 2D slices covering the required volume. In MRI, however, images may be obtained directly in any transversal, coronal, sagittal, or oblique section by using appropriate gradients and pulse sequences. Furthermore, no mechanical scanning is involved: slice selection and scanning are performed electronically by the use of magnetic field gradients and RF pulses.

The main components and principles of MRI are as follows [84]:

- **A magnet** that provides a strong, uniform field of the order of $0.5 - 4 T$. This causes some active nuclei to align in the direction of the field (parallel or antiparallel) and precess about the axis of the field. The rate of precession is proportional to the strength of the magnetic field H_o . The stronger the magnetic field, the more spins are aligned in the parallel state versus the antiparallel state, and the higher will be the signal-to-noise ratio (SNR) of the data acquired.
- **An RF transmitter** to deliver an RF electromagnetic pulse \mathbf{H}_1 to the body being imaged. The RF pulse provides the perturbation mentioned earlier: it causes the axis of precession of the net spin vector to deviate or “flip” from the z axis. In order for this to happen, the frequency of the RF field must be the same as that of precession of the active nuclei, such that the nuclei can absorb energy from the RF field (hence the term “resonance” in MRI). The frequency of RF-induced rotation is given by

$$\omega_1 = \gamma H_1. \quad (1.5)$$

When the RF perturbation is removed, the active nuclei return to their unperturbed states (alignment with H_o) through various relaxation processes, emitting energy in the form of RF signals.

- **A gradient system** to apply to the body a controlled space-variant and time-variant magnetic field

$$h(t, \mathbf{x}) = \mathbf{G}(t) \cdot \mathbf{x}, \quad (1.6)$$

where \mathbf{x} is a vector representing the spatial coordinates, \mathbf{G} is the gradient applied, and t is time. The components of \mathbf{G} along the z direction as well as in the x and y directions (the plane $x - y$ is orthogonal to the z axis) are controlled individually; however, the component of magnetic field change is only in the z direction. The gradient causes nuclei at different positions to precess at different frequencies, and provides for spatial coding of the signal emitted from the body. The Larmor frequency at \mathbf{x} is given by

$$\omega(\mathbf{x}) = \gamma(H_o + \mathbf{G} \cdot \mathbf{x}). \quad (1.7)$$

Nuclei at specific positions or planes in the body may be excited selectively by applying RF pulses of specific frequencies. The combination of the gradient fields and the RF pulses applied is called the pulse sequence.

- **An RF detector system** to detect the RF signals emitted from the body. The RF signal measured outside the body represents the sum of the RF signals emitted by active nuclei from a certain part or slice of the body, as determined by the pulse sequence. The spectral spread of the RF signal due to the application of gradients provides information on the location of the corresponding source nuclei.

- **A computing and imaging system** to reconstruct images from the measured data, as well as process and display the images. Depending upon the pulse sequence applied, the RF signal sensed may be formulated as the 2D or 3D Fourier transform of the image to be reconstructed [4, 84, 94, 95]. The data measured correspond to samples of the 2D Fourier transform of a sectional image at points on concentric squares or circles [4]. The Fourier or backprojection methods of image reconstruction from projections (described in Chapter 9) may then be used to obtain the image. (The Fourier method is the most commonly used method for reconstruction of MR images.)


Various pulse sequences may be used to measure different parameters of the tissues in the body being imaged. The image obtained is a function of the nuclear spin density in space and the corresponding parameters of the relaxation processes involved. Longitudinal magnetization refers to the component of the magnetization vector along the direction of the external magnetic field. The process by which longitudinal magnetization returns to its state of equilibrium (that is, realignment with the external magnetic field) after an excitation pulse is referred to as longitudinal relaxation. The time constant of longitudinal relaxation is labeled as T_1 .

A 90° RF pulse causes the net magnetization vector to be oriented in the plane perpendicular to the external magnetic field: this is known as transverse magnetization. When the excitation is removed, the affected nuclei return to their states of equilibrium, emitting a signal, known as the free-induction decay (FID) signal, at the Larmor frequency. The decay time constant of transverse magnetization is labeled as T_2 . Values of T_1 for various types of tissues range from 200 *ms* to 2,000 *ms*; T_2 values range from 80 *ms* to 180 *ms*. Several other parameters may be measured by using different MRI pulse sequences: the resulting images may have different appearances and clinical applications.

MRI is suitable for functional imaging. The increased supply of oxygen (or oxygenated blood) to certain regions of the brain due to related stimuli may be recorded on MR images. The difference between the prestimulus and post-stimulus images may then be used to analyze the functional aspects of specific regions of the brain.

Examples: Figure 1.31 shows a sagittal MR image of a patient's knee, illustrating the bones and cartilages that form the knee joint. Images of this type assist in the detection of bruised bones, bleeding inside the distal end of the femur, torn cartilages, and ruptured ligaments.

Figures 1.32 (a)–(c) illustrate the sagittal, coronal, and transversal (cross-sectional) views of the MR image of a patient's head. The images show the details of the structure of the brain. MRI is useful in functional imaging of the brain.

**FIGURE 1.31**

Sagittal section of the MR image of a patient's knee. Image courtesy of Foothills Hospital, Calgary.




FIGURE 1.32

(a) Sagittal, (b) coronal, and (c) transversal (cross-sectional) MR images of a patient's head. Images courtesy of Foothills Hospital, Calgary.

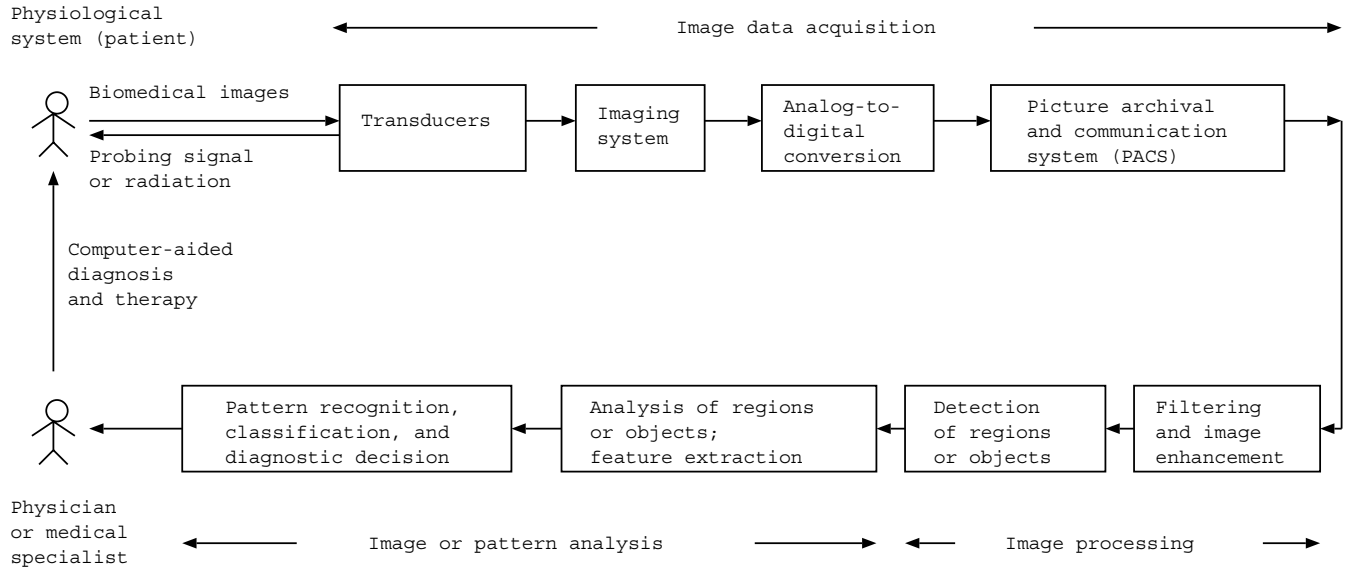


FIGURE 1.33
Computer-aided diagnosis and therapy based upon biomedical image analysis.

1.10 Objectives of Biomedical Image Analysis

The representation of biomedical images in electronic form facilitates computer processing and analysis of the data. Figure 1.33 illustrates the typical steps and processes involved in computer-aided diagnosis (CAD) and therapy based upon biomedical images analysis.

The human-instrument system: The components of a *human-instrument system* [97, 98, 99, 100, 101, 102] and some related notions are described in the following paragraphs.

- The subject (or patient): It is important always to bear in mind that the main purpose of biomedical imaging and image analysis is to provide a certain benefit to the subject or patient. All systems and procedures should be designed so as not to cause undue inconvenience to the subject, and not to cause any harm or danger. In applying invasive or risky procedures, it is extremely important to perform a risk-benefit analysis and determine if the anticipated benefits of the procedure are worth placing the subject at the risks involved.
- Transducers: films, scintillation detectors, fluorescent screens, solid-state detectors, piezoelectric crystals, X-ray generators, ultrasound generators, EM coils, electrodes, sensors.
- Signal-conditioning equipment: PMTs, amplifiers, filters.
- Display equipment: oscilloscopes, strip-chart or paper recorders, computer monitors, printers.
- Recording, data processing, and transmission equipment: films, analog-to-digital converters (ADCs), digital-to-analog converters (DACs), digital tapes, compact disks (CDs), diskettes, computers, telemetry systems, picture archival and communication systems (PACS).
- Control devices: power supply stabilizers and isolation equipment, patient intervention systems.

The major objectives of biomedical instrumentation [97, 98, 99, 100, 101, 102] in the context of imaging and image analysis are:

- Information gathering — measurement of phenomena to interpret an organ, a process, or a system.
- Screening — investigating a large asymptomatic population for the incidence of a certain disease (with the aim of early detection).
- Diagnosis — detection or confirmation of malfunction, pathology, or abnormality.

- Monitoring — obtaining periodic information about a system.
- Therapy and control — modification of the behavior of a system based upon the outcome of the activities listed above to ensure a specific result.
- Evaluation — objective analysis to determine the ability to meet functional requirements, obtain proof of performance, perform quality control, or quantify the effect of treatment.

Invasive versus noninvasive procedures: Image acquisition procedures may be categorized as invasive or noninvasive procedures. Invasive procedures involve the placement of devices or materials inside the body, such as the insertion of endoscopes, catheter-tip sensors, and X-ray contrast media. Noninvasive procedures are desirable in order to minimize risk to the subject.

Note that making measurements or imaging with X rays, ultrasound, etc. could strictly be classified as invasive procedures because they involve penetration of the body with externally administered radiation, even though the radiation is invisible and there is no visible puncturing or invasion of the body.

Active versus passive procedures: Image acquisition procedures may be categorized as active or passive procedures. Active data acquisition procedures require external stimuli to be applied to the subject, or require the subject to perform a certain activity to stimulate the system of interest in order to elicit the desired response. For example, in SPECT investigations of myocardial ischemia, the patient performs vigorous exercise on a treadmill. An ischemic zone is better delineated in SPECT images taken when the cardiac system is under stress than when at rest. While such a procedure may appear to be innocuous, it does carry risks in certain situations for some subjects: stressing an unwell system beyond a certain limit may cause pain; in the extreme situation, the procedure may cause irreparable damage or death. The investigator should be aware of such risks, factor them in a risk–benefit analysis, and be prepared to manage adverse reactions.

Passive procedures do not require the subject to perform any activity. Acquiring an image of a subject using reflected natural light (with no flash from the camera) or thermal emission could be categorized as a passive and non-contact procedure.

Most organizations require ethical approval by specialized committees for experimental procedures involving human or animal subjects, with the aim of minimizing the risk and discomfort to the subject and maximizing the benefits to both the subject and the investigator.

1.11 Computer-aided Diagnosis

Radiologists, physicians, cardiologists, neuroscientists, pathologists, and other health-care professionals are highly trained and skilled practitioners. Why then would we want to suggest the use of computers for the analysis of biomedical images? The following paragraphs provide some arguments in favor of the application of computers to process and analyze biomedical images.

- Humans are highly skilled and fast in the analysis of visual patterns, but are slow (usually) in arithmetic operations with large numbers of values. A single 64×64 -pixel SPECT image contains a total of 4,096 pixels; a high-resolution mammogram may contain as many as $5,000 \times 4,000 = 20 \times 10^6$ pixels. If images need to be processed to remove noise or extract a parameter, it would not be practical for a person to perform such computation. Computers can perform millions of arithmetic operations per second. It should be noted, however, that the recognition of objects and patterns in images using mathematical procedures typically requires huge numbers of operations that could lead to slow responses in such tasks from low-level computers. A trained human observer, on the other hand, can usually recognize an object or a pattern in an instant.
- Humans could be affected by fatigue, boredom, and environmental factors, and are susceptible to committing errors. Working with large numbers of images in one sitting, such as in breast cancer screening, poses practical difficulties. A human observer could be distracted by other events in the surrounding areas and may miss uncommon signs present in some images. Computers, being inanimate but mathematically accurate and consistent machines, can be designed to perform computationally specific and repetitive tasks.
- Analysis by humans is usually subjective and qualitative. When comparative analysis is required between an image of a subject and another or a reference pattern, a human observer would typically provide a qualitative response. Specific or objective comparison — for example, the comparison of the volume of two regions to the accuracy of the order of even a milliliter — would require the use of a computer. The derivation of quantitative or numerical features from images would certainly demand the use of computers.
- Analysis by humans is subject to interobserver as well as intraobserver variations (with time). Given that most analyses performed by humans are based upon qualitative judgment, they are liable to vary with time for a given observer, or from one observer to another. The former could be due to lack of diligence or due to inconsistent application of

knowledge, and the latter due to variations in training and the level of understanding or competence. Computers can apply a given procedure repeatedly and whenever recalled in a consistent manner. Furthermore, it is possible to encode the knowledge (to be more specific, the logical processes) of many experts into a single computational procedure, and thereby enable a computer with the collective “intelligence” of several human experts in an area of interest.

One of the important points to note in the discussion above is that *quantitative analysis* becomes possible by the application of computers to biomedical images. The logic of medical or clinical diagnosis via image analysis could then be *objectively* encoded and *consistently* applied in routine or repetitive tasks. However, it should be emphasized at this stage that the end-goal of biomedical image analysis should be computer-*aided* diagnosis and not automated diagnosis. A physician or medical specialist typically uses a significant amount of information in addition to images, including the general physical appearance and mental state of the patient, family history, and socio-economic factors affecting the patient, many of which are not amenable to quantification and logical rule-based processes. Biomedical images are, at best, indirect indicators of the state of the patient; many cases may lack a direct or unique image-to-pathology relationship. The results of image analysis need to be integrated with other clinical signs, symptoms, and information by a physician or medical specialist. Above all, the *intuition* of the medical specialist plays an important role in arriving at the final diagnosis. For these reasons, and keeping in mind the realms of practice of various licensed and regulated professions, liability, and legal factors, the final diagnostic decision is best left to the physician or medical specialist. It could be expected that quantitative and objective analysis facilitated by the application of computers to biomedical image analysis will lead to a more accurate diagnostic decision by the physician.

On the importance of quantitative analysis:

“When you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind: it may be the beginning of knowledge, but you have scarcely, in your thoughts, advanced to the stage of *science*.”

– Lord Kelvin (*William Thomson, 1824 – 1907*) [103]

On assumptions made in quantitative analysis:

“Things do not in general run around with their measure stamped on them like the capacity of a freight car; it requires a certain amount of investigation to discover what their measures are ... What most experimenters take for granted before they begin their

experiments is infinitely more interesting than any results to which their experiments lead.”

– *Norbert Wiener (1894 – 1964)*

1.12 Remarks

We have taken a general look at the nature of biomedical images in this chapter, and seen a few images illustrated for the purpose of gaining familiarity with their appearance and typical features. Specific details of the characteristics of several types of biomedical images and their processing or analysis are described in subsequent chapters, along with more examples.

We have also stated the objectives of biomedical imaging and image analysis. Some practical difficulties that arise in biomedical investigations based upon imaging were discussed in order to draw attention to the relevant practical issues. The suitability and desirability of the application of computers for biomedical image analysis were discussed, with emphasis on objective and quantitative analysis toward the end-goal of CAD. The following chapters provide the descriptions of several image processing and analysis techniques for various biomedical applications.

1.13 Study Questions and Problems

1. Give three reasons for the application of computers in medicine for computer-aided diagnosis.
2. List the relative advantages and disadvantages of X-ray, ultrasound, CT, MR, and nuclear medicine imaging for two clinical applications. Indicate where each method is inappropriate or inapplicable.
3. Discuss the factors affecting the choice of X-ray, ultrasound, CT, MR, and nuclear medicine imaging procedures for clinical applications.
4. Describe a few sources of artifact in X-ray, ultrasound, CT, MR, and nuclear medicine imaging.
5. Discuss the sources and the nature of random, periodic, structured, and physiological artifacts in medical images.
6. Describe the difference between anatomical (physical) and functional (physiological) imaging. Give examples.
7. Distinguish between active and passive medical imaging procedures; give examples.

8. Distinguish between invasive and noninvasive medical imaging procedures; give examples.
 9. Discuss factors that affect the resolution in various medical imaging modalities, including X-ray, ultrasound, CT, MR, and nuclear medicine imaging.
-

1.14 Laboratory Exercises and Projects

1. Visit a few medical imaging facilities in your local hospital or health sciences center. View the procedures related to the acquisition of a few medical images, including X-ray, ultrasound, MR, CT, and SPECT images.
Respect the priority, privacy, and confidentiality of patients.
Discuss the imaging protocols and parameters with a medical physicist and the technologists. Develop an understanding of the relationship between the imaging system parameters, image quality, and radiation exposure to the patient.
Request a radiologist to explain how he or she interprets the images. Obtain information on the differences between normal and abnormal (disease) patterns in each mode of imaging.
Collect a few sample images for use in image processing experiments, after obtaining the necessary permissions and ensuring that you carry no patient identification out of the facility.
2. Most medical imaging facilities use phantoms or test objects for quality control of imaging systems. If a phantom is not available, prepare one by attaching strips of different thickness and different metals to a plastic or plexiglass sheet. With the help of a qualified technologist, obtain X-ray images of a phantom at widely different kVp and mAs settings. Study the contrast, noise, and detail visibility in the resulting images.
Digitize the images for use in image processing experiments.
3. Visit a medical (clinical or pathology) laboratory. View several samples and specimens through microscopes.
Respect the priority, privacy, and confidentiality of patients.
Request a technologist or pathologist to explain how he or she interprets the images. Obtain information on the differences between normal and abnormal (disease) patterns in different types of samples and tests.
Collect a few sample images for use in image processing experiments, after obtaining the necessary permissions and ensuring that you carry no patient identification out of the laboratory.
4. Interview physicians, radiologists, pathologists, medical physicists, and medical technologists to find areas where they need and would like to use computing technology, digital image processing, computer vision, pattern recognition, and pattern classification methods to help in their work. Volunteer to assist them in their work! Develop projects for your course of study or research

in biomedical image analysis. Request a specialist in the relevant area to collaborate with you in the project.

5. Prepare a set of test images by collecting at least ten images that contain the following features:
 - a collection of small objects,
 - a collection of large objects,
 - directional (oriented) features,
 - fine texture,
 - coarse texture,
 - geometrical shapes,
 - human faces,
 - smooth features,
 - sharp edges.

Scan a few photos from your family photo album. Limit synthesized images to one or two in the collection. Limit images borrowed from Web sites to two in the collection. Use the images in the exercises provided at the end of each chapter.

For a selection of test images from those that have been used in this book, visit www.enel.ucalgary.ca/People/Ranga/enel697

Image Quality and Information Content

Several factors affect the quality and information content of biomedical images acquired with the modalities described in Chapter 1. A few considerations in biomedical image acquisition and analysis that could have a bearing on image quality are described in Section 2.1. A good understanding of such factors, as well as appropriate characterization of the concomitant loss in image quality, are essential in order to design image processing techniques to remove the degradation and/or improve the quality of biomedical images. The characterization of information content is important for the same purposes as above, as well as in the analysis and design of image transmission and archival systems.

An inherent problem in characterizing quality lies in the fact that image quality is typically judged by human observers in a subjective manner. To quantify the notion of image quality is a difficult proposition. Similarly, the nature of the information conveyed by an image is difficult to quantify due to its multifaceted characteristics in terms of statistical, structural, perceptual, semantic, and diagnostic connotations. However, several measures have been designed to characterize or quantify a few specific attributes of images, which may in turn be associated with various notions of quality as well as information content. The numerical values of such measures of a given image before and after certain processes, or the changes in the attributes due to certain phenomena, could then be used to assess variations in image quality and information content. We shall explore several such measures in this chapter.

2.1 Difficulties in Image Acquisition and Analysis

In Chapter 1, we studied several imaging systems and procedures for the acquisition of many different types of biomedical images. The practical application of these techniques may pose certain difficulties: the investigator often faces conditions that may impose limitations on the quality and information content of the images acquired. The following paragraphs illustrate a few practical difficulties that one might encounter in biomedical image acquisition and analysis.

Accessibility of the organ of interest: Several organs of interest in imaging-based investigation are situated well within the body, encased in protective and difficult-to-access regions, for good reason! For example, the brain is protected by the skull, and the prostate is situated at the base of the bladder near the pelvic outlet. Several limitations are encountered in imaging such organs; special imaging devices and image processing techniques are required to facilitate their visualization. Visualization of the arteries in the brain requires the injection of an X-ray contrast agent and the subtraction of a reference image; see Section 4.1. Special transrectal probes have been designed for 3D ultrasonic imaging of the prostate [92]. Despite the use of such special devices and techniques, images obtained in applications as above tend to be affected by severe artifacts.

Variability of information: Biological systems exhibit great ranges of inherent variability within their different categories. The intrinsic and natural variability presented by biological entities within a given class far exceeds the variability that we may observe in engineering, physical, and manufactured samples. The distinction between a normal pattern and an abnormal pattern is often clouded by significant overlap between the ranges of the features or variables that are used to characterize the two categories; the problem is compounded when multiple abnormalities need to be considered. Imaging conditions and parameters could cause further ambiguities due to the effects of subject positioning and projection. For example, most malignant breast tumors are irregular and spiculated in shape, whereas benign masses are smooth and round or oval. However, some malignant tumors may present smooth shapes, and some benign masses may have rough shapes. A tumor may present a rough appearance in one view or projection, but a smoother profile in another. Furthermore, the notion of shape roughness is nonspecific and open-ended. Overlapping patterns caused by ligaments, ducts, and breast tissue that may lie in other planes, but are integrated on to a single image plane in the process of mammographic imaging, could also affect the appearance of tumors and masses in images. The use of multiple views and spot magnification imaging could help resolve some of these ambiguities, but at the cost of additional radiation dose to the subject.

Physiological artifacts and interference: Physiological systems are dynamic and active. Some activities, such as breathing, may be suspended voluntarily by an adult subject (in a reasonable state of health and well-being) for brief periods of time to permit improved imaging. However, cardiac activity, blood circulation, and peristaltic movement are not under one's volitional control. The rhythmic contractile activity of the heart poses challenges in imaging of the heart. The pulsatile movement of blood through the brain causes slight movements of the brain that could cause artifacts in angiographic imaging; see Section 4.1. Dark shadows may appear in ultrasound images next to bony regions due to significant attenuation of the investigating beam, and hence the lack of echoes from tissues beyond the bony regions along

the path of beam propagation. An analyst should pay attention to potential physiological artifacts when interpreting biomedical images.

Special techniques have been developed to overcome some of the limitations mentioned above in cardiac imaging. Electronic steering of the X-ray beam has been employed to reduce the scanning time required for CT projection data acquisition in order to permit imaging of the heart; see Figure 1.21. State-of-the-art multislice and helical-scan CT scanners acquire the required data in intervals much shorter than the time taken by the initial models of CT scanners. Cardiac nuclear medicine imaging is performed by gating the photon-counting process to a certain specific phase of the cardiac cycle by using the electrocardiogram (ECG) as a reference; see Figure 1.27 and Section 3.10. Although nuclear medicine imaging procedures take several minutes, the almost-periodic activity of the heart permits the cumulative imaging of its musculature or chambers at particular positions repeatedly over several cardiac cycles.

Energy limitations: In X-ray mammography, considering the fact that the organ imaged is mainly composed of soft tissues, a low kVp would be desired in order to maximize image contrast. However, low-energy X-ray photons are absorbed more readily than high-energy photons by the skin and breast tissues, thereby increasing the radiation dose to the patient. A compromise is required between these two considerations. Similarly, in TEM, a high- kV electron beam would be desirable in order to minimize damage to the specimen, but a low- kV beam can provide improved contrast. The practical application of imaging techniques often requires the striking of a trade-off between conflicting considerations as above.

Patient safety: The protection of the subject or patient in a study from electrical shock, radiation hazard, and other potentially dangerous conditions is an unquestionable requirement of paramount importance. Most organizations require ethical approval by specialized committees for experimental procedures involving human or animal subjects, with the aim of minimizing the risk and discomfort to the subject and maximizing the benefits to both the subjects and the investigator. The relative levels of potential risks involved should be assessed when a choice is available between various procedures, and analyzed against their relative benefits. Patient safety concerns may preclude the use of a procedure that may yield better images or results than others, or may require modifications to a procedure that may lead to inferior images. Further image processing steps would then become essential in order to improve image quality or otherwise compensate for the initial compromise.

2.2 Characterization of Image Quality

Biomedical images are typically complex sources of several items of information. Furthermore, the notion of quality cannot be easily characterized with a small number of features or attributes. Because of these reasons, researchers have developed a rather large number of measures to represent quantitatively several attributes of images related to impressions of quality. Changes in measures related to quality may be analyzed for several purposes, such as:

- comparison of images generated by different medical imaging systems;
- comparison of images obtained using different imaging parameter settings of a given system;
- comparison of the results of several image enhancement algorithms;
- assessment of the effect of the passage of an image through a transmission channel or medium; and
- assessment of images compressed by different data compression techniques at different rates of loss of data, information, or quality.

Specially designed phantoms are often used to test medical imaging systems for routine quality control [104, 105, 106, 107, 108]. Bijkerk et al. [109] developed a phantom with gold disks of different diameter and thickness to test mammography systems. Because the signal contrast and location are known from the design of the phantom, the detection performance of trained observers may be used to test and compare imaging systems.

Ideally, it is desirable to use “numerical observers”: automatic tools to measure and express image quality by means of numbers or “figures of merit” (FOMs) that could be objectively compared; see Furuie et al. [110] and Barrett [111] for examples. It is clear that not only are FOMs important, but so is the methodology for their comparison. Kayargadde and Martens [112, 113] discuss the relationships between image quality attributes in a psychometric space and a perceptual space.

Many algorithms have been proposed to explore various attributes of images or imaging systems. The attributes take into consideration either the whole image or a chosen region to calculate FOMs, and are labeled as being “global” or “local”, respectively. Often, the measured attribute is image definition — the clarity with which details are reproduced [114] — which is typically expressed in terms of image sharpness. This notion was first mentioned by Higgins and Jones [115] in the realm of photography, but is valid for image evaluation in a broader context. Rangayyan and Elkadiki [116] present a survey of different methods to measure sharpness in photographic and digital images (see Section 2.15). Because quality is a subjective notion, the results

obtained by algorithms such as those mentioned above need to be validated against the evaluation of test images by human observers. This could be done by submitting the same set of images to human and numerical (computer) evaluation, and then comparing the results [104, 105, 106, 107, 108, 117]. Subjective and objective judgment should agree to some degree under defined conditions in order for the numerical measures to be useful. The following sections describe some of the concepts and measures that are commonly used in biomedical image analysis.

2.3 Digitization of Images

The representation of natural scenes and objects as digital images for processing using computers requires two steps: sampling and quantization. Both of these steps could potentially cause loss of quality and introduce artifacts.

2.3.1 Sampling

Sampling is the process of representing a continuous-time or continuous-space signal on a discrete grid, with samples that are separated by (usually) uniform intervals. The theory and practice of sampling 1D signals have been well established [1, 2, 7]. In essence, a band-limited signal with the frequency of its fastest component being f_m Hz may be represented without loss by its samples obtained at the Nyquist rate of $f_s = 2 f_m$ Hz.

Sampling may be modeled as the multiplication of the given continuous-time or analog signal with a periodic train of impulses. The multiplication of two signals in the time domain corresponds to the convolution of their Fourier spectra. The Fourier transform of a periodic train of impulses is another periodic train of impulses with a period that is equal to the inverse of the period in the time domain (that is, f_s Hz). Therefore, the Fourier spectrum of the sampled signal is periodic, with a period equal to f_s Hz. A sampled signal has infinite bandwidth; however, the sampled signal contains distinct or unique frequency components only up to $f_m = \pm f_s/2$ Hz.

If the signal as above is sampled at a rate lower than f_s Hz, an error known as *aliasing* occurs, where the frequency components above $f_s/2$ Hz appear at lower frequencies. It then becomes impossible to recover the original signal from its sampled version.

If sampled at a rate of at least f_s Hz, the original signal may be recovered from its sampled version by lowpass filtering and extracting the base-band component over the band $\pm f_m$ Hz from the infinite spectrum of the sampled signal. If an ideal (rectangular) lowpass filter were to be used, the equivalent operation in the time domain would be convolution with a sinc function (which

is of infinite duration). This operation is known as *interpolation*. Other interpolating functions of finite duration need to be used in practice, with the equivalent filter extracting the base-band components without significant reduction in gain over the band $\pm f_m$ Hz.

In practice, in order to prevent aliasing errors, it is common to use an *anti-aliasing filter* prior to the sampling of 1D signals, with a pass-band that is close to $f_s/2$ Hz, with the prior knowledge that the signal contains no significant energy or information beyond $f_m \leq f_s/2$ Hz. Analog spectrum analyzers may be used to estimate the bandwidth and spectral content of a given 1D analog signal prior to sampling.

All of the concepts explained above apply to the sampling of 2D signals or images. However, in most real-life applications of imaging and image processing, it is not possible to estimate the frequency content of the images, and also not possible to apply anti-aliasing filters. Adequate sampling frequencies need to be established for each type of image or application based upon prior experience and knowledge. Regardless, even with the same type of images, different sampling frequencies may be suitable or adequate for different applications.

Figure 2.1 illustrates the loss of quality associated with sampling an image at lower and lower numbers of pixels.

Biomedical images originally obtained on film are usually digitized using high-resolution CCD cameras or laser scanners. Several newer biomedical imaging systems include devices for direct digital data acquisition. In digital imaging systems such as CT, sampling is inherent in the measurement process, which is also performed in a domain that is different from the image domain. This adds a further level of complexity to the analysis of sampling. Practical experimentation and experience have helped in the development of guidelines to assist in such applications.

2.3.2 Quantization

Quantization is the process of representing the values of a sampled signal or image using a finite set of allowed values. In a digital representation using n bits per sample and positive integers only, there exist 2^n possible *quantized* levels, spanning the range $[0, 2^n - 1]$. If $n = 8$ bits are used to represent each pixel, there can exist 256 values or *gray levels* to represent the values of the image at each pixel, in the range $[0, 255]$.

It is necessary to map appropriately the range of variation of the given analog signal, such as the output of a charge-coupled device (CCD) detector or a video device, to the input dynamic range of the quantizer. If the lowest level (or lower threshold) of the quantizer is set too high in relation to the range of the original signal, the quantized output will have several samples with the value zero, corresponding to all signal values that are less than the lower threshold. Similarly, if the highest level (or higher threshold) of the quantizer is set too low, the output will have several samples with the highest




FIGURE 2.1

Effect of sampling on the appearance and quality of an image: (a) 225×250 pixels; (b) 112×125 pixels; (c) 56×62 pixels; and (d) 28×31 pixels. All four images have 256 gray levels at 8 bits per pixel.

quantized level, corresponding to all signal values that are greater than the higher threshold. Furthermore, the decision levels of the quantizer should be optimized in accordance with the probability density function (PDF) of the original signal or image.

The Lloyd–Max quantization procedure [8, 9, 118, 119] to optimize a quantizer is derived as follows. Let $p(r)$ represent the PDF of the amplitude or gray levels in the given image, with the values of the continuous or analog variable r varying within the range $[r_{\min}, r_{\max}]$. Let the range $[r_{\min}, r_{\max}]$ be divided into L parts demarcated by the decision levels $R_0, R_1, R_2, \dots, R_L$, with $R_0 = r_{\min}$ and $R_L = r_{\max}$; see Figure 2.2. Let the L output levels of the quantizer represent the values $Q_0, Q_1, Q_2, \dots, Q_{L-1}$, as indicated in Figure 2.2.

The mean-squared error (MSE) in representing the analog signal by its quantized values is given by

$$\overline{\varepsilon^2} = \sum_{l=0}^{L-1} \int_{R_l}^{R_{l+1}} (r - Q_l)^2 p(r) dr. \quad (2.1)$$

Several procedures exist to determine the values of R_l and Q_l that minimize the MSE [8, 9, 118, 119]. A classical result indicates that the output level Q_l should lie at the centroid of the part of the PDF between the decision levels R_l and R_{l+1} , given by

$$Q_l = \frac{\int_{R_l}^{R_{l+1}} r p(r) dr}{\int_{R_l}^{R_{l+1}} p(r) dr}, \quad (2.2)$$

which reduces to

$$Q_l = \frac{R_l + R_{l+1}}{2} \quad (2.3)$$

if the PDF is uniform. It also follows that the decision levels are then given by

$$R_l = \frac{Q_{l-1} + Q_l}{2}. \quad (2.4)$$

It is common to quantize images to 8 *bits/pixel*. However, CT images represent a large dynamic range of X-ray attenuation coefficient, normalized into *HU*, over the range $[-1,000, 1,000]$ for human tissues. Small differences of the order of 10 *HU* could indicate the distinction between normal tissue and diseased tissue. If the range of 2,000 *HU* were to be quantized into 256 levels using an 8-bit quantizer, each quantized level would represent a change of $\frac{2,000}{256} = 7.8125$ *HU*, which could lead to the loss of the distinction as above in noise. For this reason, CT and several other medical images are quantized using 12 – 16 *bits/pixel*.

The use of an inadequate number of quantized gray levels leads to false contours and poor representation of image intensities. Figure 2.3 illustrates the loss of image quality as the number of bits per pixel is reduced from six to one.




FIGURE 2.2

Quantization of an image gray-level signal r with a Gaussian (solid line) or uniform (dashed line) PDF. The quantizer output levels are indicated by Q_l and the decision levels represented by R_l .

The quantized values in a digital image are commonly referred to as *gray levels*, with 0 representing black and 255 standing for white when 8-bit quantization is used. Unfortunately, this goes against the notion of a larger amount of gray being darker than a smaller amount of gray! However, if the quantized values represent optical density (OD), a larger value would represent a darker region than a smaller value. Table 2.1 lists a few variables that bear different relationships with the displayed pixel value.

2.3.3 Array and matrix representation of images

Images are commonly represented as 2D functions of space: $f(x, y)$. A digital image $f(m, n)$ may be interpreted as a discretized version of $f(x, y)$ in a 2D array, or as a matrix; see Section 3.5 for details on matrix representation of images and image processing operations. The notational differences between the representation of an image as a function of space and as a matrix could be a source of confusion.




FIGURE 2.3

Effect of gray-level quantization on the appearance and quality of an image: (a) 64 gray levels (6 bits per pixel); (b) 16 gray levels (4 bits per pixel); (c) four gray levels (2 bits per pixel); and (d) two gray levels (1 bit per pixel). All four images have 225×250 pixels. Compare with the image in Figure 2.1 (a) with 256 gray levels at 8 bits per pixel.

TABLE 2.1

Relationships Between Tissue Type, Tissue Density, X-ray Attenuation Coefficient, Hounsfield Units (HU), Optical Density (OD), and Gray Level [120, 121]. The X-ray Attenuation Coefficient was Measured at a Photon Energy of 103.2 keV [121].

Tissue type	Density gm/cm^3	X-ray attenuation (cm^{-1})	Hounsfield units	Optical density	Gray level (brightness)	Appearance in image
lung	< 0.001	lower	low [-700, -800]	high	low	dark
liver	1.2	0.18	medium [50, 70]	medium	medium	gray
bone	1.9	higher	high [+800, +1, 000]	low	high	white

An $M \times N$ matrix has M rows and N columns; its height is M and width is N ; numbering of the elements starts with $(1, 1)$ at the top-left corner and ends with (M, N) at the lower-right corner of the image. A function of space $f(x, y)$ that has been converted into a digital representation $f(m, n)$ is typically placed in the first quadrant in the Cartesian coordinate system. Then, an $M \times N$ will have a width of M and height of N ; indexing of the elements starts with $(0, 0)$ at the origin at the bottom-left corner and ends with $(M - 1, N - 1)$ at the upper-right corner of the image. Figure 2.4 illustrates the distinction between these two types of representation of an image. Observe that the size of a matrix is expressed as *rows* \times *columns*, whereas the size of an image is usually expressed as *width* \times *height*.

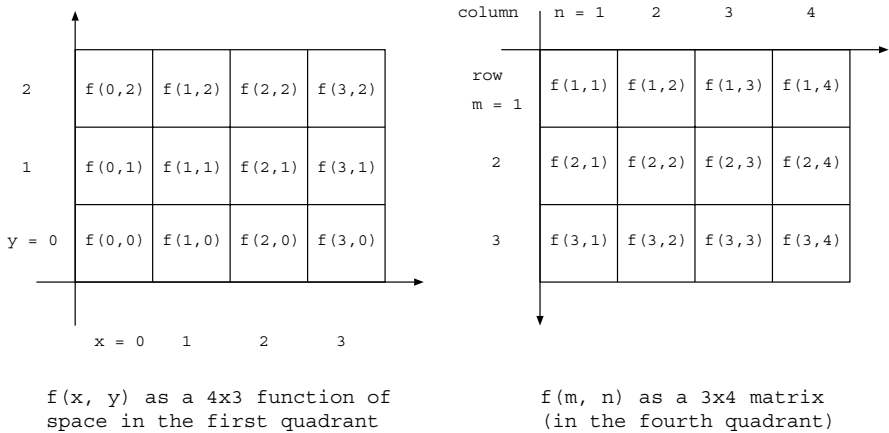


FIGURE 2.4

Array and matrix representation of an image.

2.4 Optical Density

The value of a picture element or cell — commonly known as a *pixel*, or occasionally as a pel — in an image may be expressed in terms of a physical attribute such as temperature, density, or X-ray attenuation coefficient; the intensity of light reflected from the body at the location corresponding to the pixel; or the transmittance at the corresponding location on a film rendition of the image. The last one of the options listed above is popular in medical imaging due to the common use of film as the medium for acquisition and

display of images. The OD at a spot on a film is defined as

$$OD = \log_{10} \left[\frac{I_i}{I_o} \right], \quad (2.5)$$

where I_i is the intensity of the light input and I_o is the intensity of the light transmitted through the film at the spot of interest; see Figure 2.5. A perfectly clear spot will transmit all of the light that is input and will have $OD = 0$; a dark spot that reduces the intensity of the input light by a factor of 1,000 will have $OD = 3$. X-ray films, in particular those used in mammography, are capable of representing gray levels from $OD \approx 0$ to $OD \approx 3.5$.




FIGURE 2.5

Measurement of the optical density at a spot on a film or transparency using a laser microdensitometer.

2.5 Dynamic Range

The dynamic range of an imaging system or a variable is its range or gamut of operation, usually limited to the portion of linear response, and is expressed as the maximum minus the minimum value of the variable or parameter of interest. The dynamic range of an image is usually expressed as the difference between the maximum and minimum values present in the image. X-ray films for mammography typically possess a dynamic range of 0 – 3.5 OD . Modern CRT monitors provide dynamic range of the order of 0 – 600 cd/m^2 in luminance or 1 : 1,000 in sampled gray levels.

Figure 2.6 compares the characteristic curves of two devices. Device A has a larger slope or “gamma” (see Section 4.4.3) than Device B, and hence can provide higher contrast (defined in Section 2.6). Device B has a larger latitude, or breadth of exposure and optical density over which it can operate, than Device A. Plots of film density versus the log of (X-ray) exposure are known as Hurter–Driffield or H-D curves [3].




FIGURE 2.6

Characteristic response curves of two hypothetical imaging devices.

The lower levels of response of a film or electronic display device are affected by a background level that could include the base level of the medium or operation of the device as well as noise. The response of a device typically begins with a nonlinear “toe” region before it reaches its linear range of operation. Another nonlinear region referred to as the “shoulder” region leads to the saturation level of the device. It is desirable to operate within the linear range of a given device.

Air in the lungs and bowels, as well as fat in various organs including the breast, tend to extend the dynamic range of images toward the lower end of the density scale. Bone, calcifications in the breast and in tumors, as well as metallic implants such as screws in bones and surgical clips contribute to high-density areas in images. Mammograms are expected to possess a dynamic range of $0 - 3.5 OD$. CT images may have a dynamic range of about $-1,000$ to $+1,000 HU$. Metallic implants could have HU values beyond the operating range of CT systems, and lead to saturated areas in images: the X-ray beam is effectively stopped by heavy-metal implants.

2.6 Contrast

Contrast is defined in a few different ways [9], but is essentially the difference between the parameter imaged in a region of interest (ROI) and that in a suitably defined background. If the image parameter is expressed in OD , contrast is defined as

$$C_{OD} = f_{OD} - b_{OD}, \quad (2.6)$$

where f_{OD} and b_{OD} represent the foreground ROI and background OD , respectively. Figure 2.7 illustrates the notion of contrast using circular ROIs.




FIGURE 2.7

Illustration of the notion of contrast, comparing a foreground region f with its background b .

When the image parameter has not been normalized, the measure of contrast will require normalization. If, for example, f and b represent the average light intensities emitted or reflected from the foreground ROI and the background, respectively, contrast may be defined as

$$C = \frac{f - b}{f + b}, \quad (2.7)$$

or as

$$C_1 = \frac{f - b}{b}. \quad (2.8)$$

Due to the use of a reference background, the measures defined above are often referred to as “simultaneous contrast”. It should be observed that the contrast of a region or an object depends not only upon its own intensity, but also upon that of its background. Furthermore, the measure is not simply a

difference, but a ratio. The human visual system (HVS) has bandpass filter characteristics, which lead to responses that are proportional to differences between illumination levels rather than to absolute illumination levels [122].

Example: The two squares in Figure 2.8 are of the same value (130 in the scale 0 – 255), but are placed on two different background regions of value 150 on the left and 50 on the right. The lighter background on the left makes the inner square region appear darker than the corresponding inner square on the right. This effect could be explained by the measure of simultaneous contrast: the contrast of the inner square on the left, using the definition in Equation 2.8, is

$$C_l = \frac{130 - 150}{150} = -0.1333, \quad (2.9)$$

whereas that for the inner square on the right is

$$C_r = \frac{130 - 50}{50} = +1.6. \quad (2.10)$$

The values of C_l and C_r using the definition in Equation 2.7 are, respectively, -0.0714 and $+0.444$; the advantage of this formulation is that the values of contrast are limited to the range $[-1, 1]$. The negative contrast value for the inner square on the left indicates that it is darker than the background, whereas it is the opposite for that on the right. (By covering the background regions and viewing only the two inner squares simultaneously, it will be seen that the gray levels of the latter are indeed the same.)

Just-noticeable difference: The concept of *just-noticeable difference* (JND) is important in analyzing contrast, visibility, and the quality of medical images. JND is determined as follows [9, 122]: For a given background level b as in Equation 2.8, the value of an object in the foreground f is increased gradually from the same level as b to a level when the object is just perceived. The value $(f - b)/b$ at the level of minimal perception of the object is the JND for the background level b . The experiment should, ideally, be repeated many times for the same observer, and also repeated for several observers. Experiments have shown that the JND is almost constant, at approximately 0.02 or 2%, over a wide range of background intensity; this is known as Weber's law [122].

Example: The five bars in Figure 2.9 have intensity values of (from left to right) 155, 175, 195, 215, and 235. The bars are placed on a background of 150. The contrast of the first bar (to the left), according to Equation 2.8, is

$$C_l = \frac{155 - 150}{150} = +0.033. \quad (2.11)$$

This contrast value is slightly greater than the nominal JND; the object should be barely perceptible to most observers. The contrast values of the remaining four bars are more than adequate for clear perception.

Example: Calcifications appear as bright spots in mammograms. A calcification that appears against fat and low-density tissue may possess high



FIGURE 2.8

Illustration of the effect of the background on the perception of an object (simultaneous contrast). The two inner squares have the same gray level of 130, but are placed on different background levels of 150 on the left and 50 on the right.



FIGURE 2.9

Illustration of the notion of just-noticeable difference. The five bars have intensity values of (from left to right) 155, 175, 195, 215, and 235, and are placed on a background of 150. The first bar is barely noticeable; the contrast of the bars increases from left to right.

contrast and be easily visible. On the other hand, a similar calcification that appears against a background of high-density breast tissue, or a calcification that is present within a high-density tumor, could possess low contrast, and be difficult to detect. Figure 2.10 shows a part of a mammogram with several calcifications appearing against different background tissue patterns and density. The various calcifications in this image present different levels of contrast and visibility.

Small calcifications and masses situated amidst high-density breast tissue could present low contrast close to the JND in a mammogram. Such features present significant challenges in a breast cancer screening situation. Enhancement of the contrast and visibility of such features could assist in improving the accuracy of detecting early breast cancer [123, 124, 125]; see Sections 4.9.1 and 12.10.

2.7 Histogram

The dynamic range of the gray levels in an image provides global information on the extent or spread of intensity levels across the image. However, the dynamic range does not provide any information on the existence of intermediate gray levels in the image. The histogram of an image provides information on the spread of gray levels over the complete dynamic range of the image across all pixels in the image.

Consider an image $f(m, n)$ of size $M \times N$ pixels, with gray levels $l = 0, 1, 2, \dots, L - 1$. The histogram of the image may be defined as

$$P_f(l) = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} \delta_d[f(m, n) - l], \quad l = 0, 1, 2, \dots, L - 1, \quad (2.12)$$

where the discrete unit impulse function or delta function is defined as [1, 2]

$$\delta_d(k) = \begin{cases} 1 & \text{if } k = 0 \\ 0 & \text{otherwise.} \end{cases} \quad (2.13)$$

The histogram value $P_f(l)$ provides the number of pixels in the image f that possess the gray level l . The sum of all the entries in a histogram equals the total number of pixels in the image:

$$\sum_{l=0}^{L-1} P_f(l) = MN. \quad (2.14)$$

The area under the function $P_f(l)$, when multiplied with an appropriate scaling factor, provides the total intensity, density, or brightness of the image, depending upon the physical parameter represented by the pixel values.




FIGURE 2.10

Part of a mammogram with several calcifications associated with malignant breast disease. The density of the background affects the contrast and visibility of the calcifications. The image has 768×512 pixels at a resolution of $62 \mu\text{m}$; the true width of the image is about 32 mm .

A histogram may be normalized by dividing its entries by the total number of pixels in the image. Then, with the assumption that the total number of pixels is large and that the image is a typical representative of its class or the process that generates images of its kind, the normalized histogram may be taken to represent the PDF $p_f(l)$ of the image-generating process:

$$p_f(l) = \frac{1}{MN} P_f(l). \quad (2.15)$$

It follows that

$$\sum_{l=0}^{L-1} p_f(l) = 1. \quad (2.16)$$


Example: The histogram of the image in Figure 1.3 is shown in Figure 2.11. It is seen that most of the pixels in the image lie in the narrow range of 70–150 out of the available range of 0–255. The effective dynamic range of the image may be taken to be 70–150, rather than 0–255. This agrees with the dull and low-contrast appearance of the image. The full available range of gray levels has not been utilized in the image, which could be due to poor lighting and image acquisition conditions, or due to the nature of the object being imaged.

The gray level of the large, blank background in the image in Figure 1.3 is in the range 80–90: the peak in the histogram corresponds to the general background range. The relatively bright areas of the myocyte itself have gray levels in the range 100–130. The histogram of the myocyte image is almost *unimodal*; that is, it has only one major peak. The peak happens to represent the background in the image rather than the object of interest.


Example: Figure 2.12 (a) shows the histogram of the image in Figure 1.5 (b). The discrete spikes are due to noise in the image. The histogram of the image after smoothing, using the 3×3 mean filter and rounding the results to integers, is shown in part (b) of the figure. The histogram of the filtered image is *bimodal*, with two main peaks spanning the gray level ranges 100–180 and 180–255, representing the collagen fibers and background, respectively. Most of the pixels corresponding to the collagen fibers in cross-section have gray levels below about 170; most of the brighter background pixels have values greater than 200.

Example: Figure 2.13 shows a part of a mammogram with a tumor. The normalized histogram of the image is shown in Figure 2.14. It is seen that the histogram has two large peaks in the range 0–20 representing the background in the image with no breast tissue. Although the image has bright areas, the number of pixels occupying the high gray levels in the range 200–255 is insignificant.


Example: Figure 2.15 shows a CT image of a two-year-old male patient with neuroblastoma (see Section 9.9 for details). The histogram of the image is shown in Figure 2.16 (a). The histogram of the entire CT study of the patient, including 75 sectional images, is shown in Figure 2.16 (b). Observe

**FIGURE 2.11**

Histogram of the image of the ventricular myocyte in Figure 1.3. The size of the image is $480 \times 480 = 230,400$ pixels. Entropy $H = 4.96$ bits.



(a)



(b)

FIGURE 2.12

(a) Histogram of the image of the collagen fibers in Figure 1.5 (b); $H = 7.0$ bits. (b) Histogram of the image after the application of the 3×3 mean filter and rounding the results to integers; $H = 7.1$ bits.




FIGURE 2.13

Part of a mammogram with a malignant tumor (the relatively bright region along the upper-left edge of the image). The size of the image is $700 \times 700 = 490,000$ pixels. The pixel resolution of $62 \mu\text{m}$; the width of the image is about 44 mm . Image courtesy of Foothills Hospital, Calgary.




FIGURE 2.14

Normalized histogram of the mammogram in Figure 2.13. Entropy $H = 6.92$ bits.

that the unit of the pixel variable in the histograms is HU ; however, the gray-level values in the image have been scaled for display in Figure 2.15, and do not directly correspond to the HU values. The histograms are *multimodal*, indicating the presence of several types of tissue in the CT images. The peaks in the histogram in Figure 2.16 (a) in the range $50\text{--}150$ HU correspond to liver and other abdominal organs and tissues. The small peak in the range $200\text{--}300$ HU in the same histogram corresponds to calcified parts of the tumor. The histogram of the full volume includes a small peak in the range $700\text{--}800$ HU corresponding to bone [not shown in Figure 2.16 (b)]. Histograms of this nature provide information useful in diagnosis as well as in the follow up of the effect of therapy. Methods for the analysis of histograms for application in neuroblastoma are described in Section 9.9.

2.8 Entropy

The distribution of gray levels over the full available range is represented by the histogram. The histogram provides quantitative information on the






FIGURE 2.15

CT image of a patient with neuroblastoma. Only one sectional image out of a total of 75 images in the study is shown. The size of the image is $512 \times 512 = 262,144$ pixels. The tumor, which appears as a large circular region on the left-hand side of the image, includes calcified tissues that appear as bright regions. The *HU* range of $[-200, 400]$ has been linearly mapped to the display range of $[0, 255]$; see also Figures 2.16 and 4.4. Image courtesy of Alberta Children's Hospital, Calgary.

**FIGURE 2.16**

(a) Histogram of the CT section image in Figure 2.15. (b) Histogram of the entire CT study of the patient, with 75 sectional images. The histograms are displayed for the range $HU = [-200, 400]$ only.

probability of occurrence of each gray level in the image. However, it is often desirable to express, in a single quantity, the manner in which the values of a histogram or PDF vary over the full available range. Entropy is a statistical measure of information that is commonly used for this purpose [9, 11, 126, 127, 128, 129].

The various pixels in an image may be considered to be symbols produced by a discrete information source with the gray levels as its states. Let us consider the occurrence of L gray levels in an image, with the probability of occurrence of the l^{th} gray level being $p(l)$, $l = 0, 1, 2, \dots, L - 1$. Let us also treat the gray level of a pixel as a random variable. A measure of information conveyed by an event (a pixel or a gray level) may be related to the *statistical uncertainty* of the event giving rise to the information, rather than the semantic or structural content of the signal or image. Given the unlimited scope of applications of imaging and the context-dependent meaning conveyed by images, a statistical approach as above is appropriate to serve the general purpose of analysis of the information content of images.

A measure of information $h(p)$ should be a function of $p(l)$, satisfying the following criteria [9, 11, 126, 127]:

- $h(p)$ should be continuous for $0 < p < 1$.
- $h(p) = \infty$ for $p = 0$: a totally unexpected event conveys maximal information when it does indeed occur.
- $h(p) = 0$ for $p = 1$: an event that is certain to occur does not convey any information.
- $h(p_2) > h(p_1)$ if $p_2 < p_1$: an event with a lower probability of occurrence conveys more information when it does occur than an event with a higher probability of occurrence.
- If two statistically independent image processes (or pixels) f and g are considered, the joint information of the two sources is given by the sum of their individual measures of information: $h_{f,g} = h_f + h_g$.

These requirements are met by $h(p) = -\log(p)$.

When a source generates a number of gray levels with different probabilities, a measure of average information or *entropy* is defined as the expected value of information contained in each possible level:

$$H = \sum_{l=0}^{L-1} p(l) h[p(l)]. \quad (2.17)$$

Using $-\log_2$ in place of h , we obtain the commonly used definition of entropy as

$$H = - \sum_{l=0}^{L-1} p(l) \log_2 [p(l)] \text{ bits}. \quad (2.18)$$

Because the gray levels are considered as individual entities in this definition, that is, no neighboring elements are taken into account, the result is known as the zeroth-order entropy [130].

The entropies of the images in Figures 1.3 and 2.13, with the corresponding histogram or PDF in Figures 2.11 and 2.14, are 4.96 and 6.92 *bits*, respectively. Observe that the histogram in Figure 2.14 has a broader spread than that in Figure 2.11, which accounts for the correspondingly higher entropy.

Differentiating the function in Equation 2.18 with respect to $p(l)$, it can be shown that the maximum possible entropy occurs when all the gray levels occur with the same probability (equal to $\frac{1}{L}$), that is, when the various gray levels are equally likely:

$$H_{\max} = - \sum_{l=0}^{L-1} \frac{1}{L} \log_2 \left[\frac{1}{L} \right] = \log_2 L. \quad (2.19)$$

If the number of gray levels in an image is 2^K , then H_{\max} is K *bits*; the maximum possible entropy of an image with 8-bit pixels is 8 *bits*.

It should be observed that entropy characterizes the *statistical* information content of a source based upon the PDF of the constituent events, which are treated as random variables. When an image is characterized by its entropy, it is important to recognize that the measure is not sensitive to the pictorial, structural, semantic, or application-specific (diagnostic) information in the image. Entropy does not account for the spatial distribution of the gray levels in a given image. Regardless, the entropy of an image is an important measure because it gives a summarized measure of the statistical information content of an image, an image-generating source, or an information source characterized by a PDF, as well as the lower bound on the noise-free transmission rate and storage capacity requirements.

Properties of entropy: A few important properties of entropy [9, 11, 126, 127, 128, 129] are as follows:

- $H_p \geq 0$, with $H_p = 0$ only for $p = 0$ or $p = 1$: no information is conveyed by events that do not occur or occur with certainty.
- The joint information $H_{(p_1, p_2, \dots, p_n)}$ conveyed by n events, with probabilities of occurrence p_1, p_2, \dots, p_n , is governed by $H_{(p_1, p_2, \dots, p_n)} \leq \log(n)$, with equality if and only if $p_i = \frac{1}{n}$ for $i = 1, 2, \dots, n$.
- Considering two images or sources f and g with PDFs $p_f(l_1)$ and $p_g(l_2)$, where l_1 and l_2 represent gray levels in the range $[0, L - 1]$, the average joint information or joint entropy is

$$H_{f,g} = - \sum_{l_1=0}^{L-1} \sum_{l_2=0}^{L-1} p_{f,g}(l_1, l_2) \log_2 [p_{f,g}(l_1, l_2)]. \quad (2.20)$$

If the two sources are statistically independent, the joint PDF $p_{f,g}(l_1, l_2)$ reduces to $p_f(l_1)p_g(l_2)$. Joint entropy is governed by the condition $H_{f,g} \leq H_f + H_g$, with equality if and only if f and g are statistically independent.

- The conditional entropy of an image f given that another image g has been observed is

$$\begin{aligned} H_{f|g} &= - \sum_{l_1=0}^{L-1} \sum_{l_2=0}^{L-1} p_g(l_2) p_{f|g}(l_1, l_2) \log_2 [p_{f|g}(l_1, l_2)] \\ &= - \sum_{l_1=0}^{L-1} \sum_{l_2=0}^{L-1} p_{f,g}(l_1, l_2) \log_2 [p_{f|g}(l_1, l_2)], \end{aligned} \quad (2.21)$$

where $p_{f|g}(l_1, l_2)$ is the conditional PDF of f given g . Then, $H_{f|g} = H_{f,g} - H_g \leq H_f$, with equality if and only if f and g are statistically independent. *Note:* The conditional PDF of f given g is defined as [128, 129]

$$p_{f|g}(l_1, l_2) = \begin{cases} \frac{p_{f,g}(l_1, l_2)}{p_g(l_2)} & \text{if } p_g(l_2) > 0 \\ 1 & \text{otherwise.} \end{cases} \quad (2.22)$$

Higher-order entropy: The formulation of entropy as a measure of information is based upon the premise that the various pixels in an image may be considered to be symbols produced by a discrete information source with the gray levels as its states. From the discussion above, it follows that the definition of entropy in Equation 2.18 assumes that the successive pixels produced by the source are statistically independent. While governed by the limit $H_{\max} = K \text{ bits}$, the entropy of a real-world image (with $K \text{ bits}$ per pixel) encountered in practice could be considerably lower, due to the fact that neighboring pixels in most real images are not independent of one another. Due to this reason, it is desirable to consider sequences of pixels to estimate the true entropy or information content of a given image.

Let $p(\{l_n\})$ represent the probability of occurrence of the sequence $\{l_0, l_1, l_2, \dots, l_n\}$ of gray levels in the image f . The n^{th} -order entropy of f is defined as

$$H_n = - \frac{1}{(n+1)} \sum_{\{l_n\}} p(\{l_n\}) \log_2 [p(\{l_n\})], \quad (2.23)$$

where the summation is over all possible sequences $\{l_n\}$ with $(n+1)$ pixels. (*Note:* Some variations exist in the literature regarding the definition of higher-order entropy. In the definition given above, n refers to the number of neighboring or additional elements considered, not counting the initial or zeroth element; this is consistent with the definition of the zeroth-order entropy in Equation 2.18.) H_n is a monotonically decreasing function of n , and approaches the true entropy of the source as $n \rightarrow \infty$ [9, 126, 127].

Mutual information: A measure that is important in the analysis of transmission of images over a communication system, as well as in the analysis of storage in and retrieval from an archival system, with potential loss of information, is *mutual information*, defined as

$$I_{f|g} = H_f + H_g - H_{f,g} = H_f - H_{f|g} = H_g - H_{g|f}. \quad (2.24)$$

This measure represents the information received or retrieved with the following explanation: H_f is the information input to the transmission or archival system in the form of the image f . $H_{f|g}$ is the information about f given that the received or retrieved image g has been observed. (In this analysis, g is taken to be known, but f is considered to be unknown, although g is expected to be a good representation of f .) Then, if g is completely correlated with f , we have $H_{f|g} = 0$, and $I_{f|g} = H_f$: this represents the case where there is no loss or distortion in image transmission and reception (or in image storage and retrieval). If g is independent of f , $H_{f|g} = H_f$, and $I_{f|g} = 0$: this represents the situation where there is complete loss of information in the transmission or archival process.

Entropy and mutual information are important concepts that are useful in the design and analysis of image archival, coding, and communication systems; this topic is discussed in Chapter 11.

2.9 Blur and Spread Functions

Several components of image acquisition systems cause blurring due to intrinsic and practical limitations. The simplest visualization of blurring is provided by using a single, ideal point to represent the object being imaged; see Figure 2.17 (a). Mathematically, an ideal point is represented by the continuous unit impulse function or the Dirac delta function $\delta(x, y)$, defined as [1, 2, 131]

$$\delta(x, y) = \begin{cases} \text{undefined} & \text{at } x = 0, y = 0 \\ 0 & \text{otherwise,} \end{cases} \quad (2.25)$$


and

$$\int_{x=-\infty}^{\infty} \int_{y=-\infty}^{\infty} \delta(x, y) \, dx \, dy = 1. \quad (2.26)$$

Note: The 1D Dirac delta function $\delta(x)$ is defined in terms of its action within an integral as [3]

$$\int_a^b f(x) \delta(x - x_o) \, dx = \begin{cases} f(x_o) & \text{if } a < x_o < b \\ 0 & \text{otherwise,} \end{cases} \quad (2.27)$$

where $f(x)$ is a function that is continuous at x_o . This is known as the *sifting property* of the delta function, because the value of the function $f(x)$ at the location

**FIGURE 2.17**

(a) An ideal point source. (b) A Gaussian-shaped point spread function.

x_o of the delta function is sifted or selected from all of its values. The expression may be extended to all x as

$$f(x) = \int_{\alpha=-\infty}^{\infty} f(\alpha) \delta(x - \alpha) d\alpha, \quad (2.28)$$

which may also be interpreted as resolving the arbitrary signal $f(x)$ into a weighted combination of mutually orthogonal delta functions. A common definition of the delta function is in terms of its integrated strength as

$$\int_{-\infty}^{\infty} \delta(x) dx = 1, \quad (2.29)$$

with the conditions

$$\delta(x) = \begin{cases} \text{undefined} & \text{at } x = 0 \\ 0 & \text{otherwise.} \end{cases} \quad (2.30)$$

The delta function is also defined as the limiting condition of several ordinary functions, one of which is

$$\delta(x) = \lim_{\epsilon \rightarrow 0} \frac{1}{2\epsilon} \exp\left(-\frac{|x|}{\epsilon}\right). \quad (2.31)$$

The delta function may be visualized as the limit of a function with a sharp peak of undefined value, whose integral over the full extent of the independent variable is maintained as unity while its temporal or spatial extent is compressed toward zero.

The image obtained when the input is a point or impulse function is known as the impulse response or *point spread function* (PSF); see Figure 2.17 (b). Assuming the imaging system to be linear and shift-invariant (or position-invariant or space-invariant, abbreviated as LSI), the image $g(x, y)$ of an object $f(x, y)$ is given by the 2D convolution integral [8, 9, 11, 131]

$$g(x, y) = \int_{\alpha=-\infty}^{\infty} \int_{\beta=-\infty}^{\infty} h(x - \alpha, y - \beta) f(\alpha, \beta) d\alpha d\beta \quad (2.32)$$

$$\begin{aligned}
 &= \int_{\alpha=-\infty}^{\infty} \int_{\beta=-\infty}^{\infty} h(\alpha, \beta) f(x - \alpha, y - \beta) d\alpha d\beta \\
 &= h(x, y) * f(x, y),
 \end{aligned}$$

where $h(x, y)$ is the PSF, α and β are temporary variables of integration, and $*$ represents 2D convolution.

(*Note:* For details on the theory of linear systems and convolution, refer to Lathi [1], Oppenheim et al. [2], Oppenheim and Schaffer [7], and Gonzalez and Woods [8]. In extending the concepts of LSI system theory from time-domain signals to the space domain of images, it should be observed that causality is not a matter of concern in most applications of image processing.)

Some examples of the cause of blurring are:

- **Focal spot:** The physical spot on the anode (target) that generates X rays is not an ideal dimensionless point, but has finite physical dimensions and an area of the order of $1 - 2 \text{ mm}^2$. Several straight-line paths would then be possible from the X-ray source, through a given point in the object being imaged, and on to the film. The image so formed will include not only the main radiographic shadow (the “umbra”), but also an associated blur (the “penumbra”), as illustrated in Figure 2.18. The penumbra causes blurring of the image.
- **Thickness of screen or crystal:** The screen used in screen-film X-ray imaging and the scintillation crystal used in gamma-ray imaging generate visible light when struck by X or gamma rays. Due to the finite thickness of the screen or crystal, a point source of light within the detector will be sensed over a wider region on the film (see Figure 1.10) or by several PMTs (see Figure 1.25): the thicker the crystal or screen, the worse the blurring effect caused as above.
- **Scattering:** Although it is common to assume straight-line propagation of X or gamma rays through the body or object being imaged, this is not always the case in reality. X, gamma, and ultrasound rays do indeed get scattered within the body and within the detector. The effect of rays that are scattered to a direction that is significantly different from the original path will likely be perceived as background noise. However, scattering to a smaller extent may cause unsharp edges and blurring in a manner similar to those described above.

Point, line, and edge spread functions: In practice, it is often not possible or convenient to obtain an image of an ideal point: a microscopic hole in a sheet of metal may not allow adequate X-ray photons to pass through and create a useful image; an infinitesimally small drop of a radiopharmaceutical may not emit sufficient gamma-ray photons to record an appreciable image on a gamma camera. However, it is possible to construct phantoms to represent ideal lines or edges. For use in X-ray imaging, a line phantom may be created





FIGURE 2.18

The effect of a finite focal spot (X-ray-generating portion of the target) on the sharpness of the image of an object.


by cutting a narrow slit in a sheet of metal. In SPECT imaging, it is common to use a thin plastic tube, with diameter of the order of 1 mm and filled with a radiopharmaceutical, to create a line source [86, 132]. Given that the spatial resolution of a typical SPECT system is of the order of several mm , such a phantom may be assumed to represent an ideal straight line with no thickness. An image obtained of such a source is known as the *line spread function* (LSF) of the system. Because any cross-section of an ideal straight line is a point or impulse function, the reconstruction of a cross-section of a line phantom provides the PSF of the system. Observe also that the integration of an ideal point results in a straight line along the path of integration; see Figure 2.19.

In cases where the construction of a line source is not possible or appropriate, one may prepare a phantom representing an ideal edge. Such a phantom is easy to prepare for planar X-ray imaging: one needs to simply image the (ideal and straight) edge of a sheet or slab made of a material with a higher attenuation coefficient than that of the background or table upon which it is placed when imaging. In the case of CT imaging, a 3D cube or parallelepiped with its sides and edges milled to be perfect planes and straight lines, respectively, may be used as the test object. A profile of the image of such a phantom across the ideal edge provides the *edge spread function* (ESF) of the system: see Figure 2.20; see also Section 2.15. The derivative of an edge along the direction perpendicular to the edge is an ideal straight line; see Fig-

**FIGURE 2.19**

The relationship between point (impulse function), line, and edge (step) images. The height of each function represents its strength.

ure 2.19. Therefore, the derivative of the ESF gives the LSF of the system. Then, the PSF may be estimated from the LSF as described above.

**FIGURE 2.20**

Blurring of an ideal sharp edge into an unsharp edge by an imaging system.

In practice, due to the presence of noise and artifacts, it would be desirable to average several measurements of the LSF, which could be performed along the length of the line or edge. If the imaging system is anisotropic, the LSF should be obtained for several orientations of the line source. If the blur of the system varies with the distance between the detector and the source, as is

the case in nuclear medicine imaging with a gamma camera, one should also measure the LSF at several distances.

The mathematical relationships between the PSF, LSF, and ESF may be expressed as follows [3, 9, 11]. Consider integration of the 2D delta function along the x axis as follows:

$$\begin{aligned}
 f_l(x, y) &= \int_{x=-\infty}^{\infty} \delta(x, y) dx \\
 &= \int_{x=-\infty}^{\infty} \delta(x) \delta(y) dx \\
 &= \delta(y) \int_{x=-\infty}^{\infty} \delta(x) dx \\
 &= \delta(y).
 \end{aligned} \tag{2.33}$$

The last integral above is equal to unity; the separability property of the 2D impulse function as $\delta(x, y) = \delta(x) \delta(y)$ has been used above. Observe that although $\delta(y)$ has been expressed as a function of y only, it represents a 2D function of (x, y) that is independent of x in the present case. Considering $\delta(y)$ over the entire 2D (x, y) space, it becomes evident that it is a line function that is placed on the x axis. The line function is thus given by an integral of the impulse function (see Figure 2.19).

The output of an LSI system when the input is the line image $f_l(x, y) = \delta(y)$, that is, the LSF, which we shall denote here as $h_l(x, y)$, is given by

$$\begin{aligned}
 h_l(x, y) &= \int_{\alpha=-\infty}^{\infty} \int_{\beta=-\infty}^{\infty} h(\alpha, \beta) f_l(x - \alpha, y - \beta) d\alpha d\beta \\
 &= \int_{\alpha=-\infty}^{\infty} \int_{\beta=-\infty}^{\infty} h(\alpha, \beta) \delta(y - \beta) d\alpha d\beta \\
 &= \int_{\alpha=-\infty}^{\infty} h(\alpha, y) d\alpha \\
 &= \int_{x=-\infty}^{\infty} h(x, y) dx.
 \end{aligned} \tag{2.34}$$

In the equations above, $h(x, y)$ is the PSF of the system, and the sifting property of the delta function has been used. The final equation above shows that the LSF is the integral (in this case, along the x axis) of the PSF. This result also follows simply from the linearity of the LSI system and that of the operation of integration: given that $h(x, y)$ is the output due to $\delta(x, y)$ as the input, if the input is an integral of the delta function, the output will be the corresponding integral of $h(x, y)$. Observe that, in the present example, $h_l(x, y)$ is independent of x .

Let us now consider the Fourier transform of $h_l(x, y)$. Given that $h_l(x, y)$ is independent of x in the present illustration, we may write it as a 1D function

$h_l(y)$; correspondingly, its Fourier transform will be a 1D function, which we shall express as $H_l(v)$. Then, we have

$$\begin{aligned} H_l(v) &= \int_{y=-\infty}^{\infty} h_l(y) \exp(-j2\pi vy) dy \\ &= \int_{y=-\infty}^{\infty} dy \int_{x=-\infty}^{\infty} dx h(x, y) \exp[-j2\pi(ux + vy)]|_{u=0} \\ &= H(u, v)|_{u=0} \\ &= H(0, v), \end{aligned} \quad (2.35)$$

where $H(u, v)$ is the 2D Fourier transform of $h(x, y)$ (see Sections 2.11 and 2.12). This shows that the Fourier transform of the LSF gives the values of the Fourier transform of the PSF along a line in the 2D Fourier plane (in this case, along the v axis).

In a manner similar to the discussion above, let us consider integrating the line function as follows:

$$\begin{aligned} f_e(x, y) &= \int_{\beta=-\infty}^y f_l(x, \beta) d\beta \\ &= \int_{\beta=-\infty}^y \delta(\beta) d\beta. \end{aligned} \quad (2.36)$$

The resulting function has the property

$$\forall x, f_e(x, y) = \begin{cases} 1 & \text{if } y > 0 \\ 0 & \text{if } y < 0, \end{cases} \quad (2.37)$$

which represents an edge or unit step function that is parallel to the x axis (see Figure 2.19). Thus, the edge or step function is obtained by integrating the line function. It follows that the ESF is given by

$$h_e(y) = \int_{\beta=-\infty}^y h_l(\beta) d\beta. \quad (2.38)$$

Conversely, the LSF is the derivative of the ESF:

$$h_l(y) = \frac{d}{dy} h_e(y). \quad (2.39)$$

Thus the ESF may be used to obtain the LSF, which may further be used to obtain the PSF and MTF as already explained. (Observe the use of the generalized delta function to derive the discontinuous line and edge functions in this section.)

In addition to the procedures and relationships described above, based upon the Fourier slice theorem (see Section 9.2 and Figure 9.2), it can be shown that the Fourier transform of a profile of the LSF is equal to the radial profile

of the Fourier transform of the PSF at the angle of placement of the line source. If the imaging system may be assumed to be isotropic in the plane of the line source, a single radial profile is adequate to reconstruct the complete 2D Fourier transform of the PSF. Then, an inverse 2D Fourier transform provides the PSF. This method, which is essentially the Fourier method of reconstruction from projections described in Section 9.2, was used by Hon et al. [132] and Boulfefel [86] to estimate the PSF of a SPECT system.

Example of application: In the work of Boulfefel [86], a line source was prepared using a plastic tube of internal radius 1 mm, filled with 1 mCi (milli Curie) of ^{99m}Tc . The phantom was imaged using a gamma camera at various source-to-collimator distances, using an energy window of width of 14 keV centered at 140 keV. Figure 2.21 shows a sample image of the line source. Figure 2.22 shows a sample profile of the LSF and the averaged profile obtained by averaging the 64 rows of the LSF image.




FIGURE 2.21

Nuclear medicine (planar) image of a line source obtained using a gamma camera. The size of the image is 64×64 pixels, with an effective width of 100 mm. The pixel size is 1.56 mm.

It is common practice to characterize an LSF or PSF with its full width at half the maximum (FWHM) value. Boulfefel observed that the FWHM of the LSF of the gamma cameras studied varied between 0.5 cm and 1.7 cm depending upon the radiopharmaceutical used, the source-to-collimator dis-




FIGURE 2.22

Sample profile (dotted line) and averaged profile (solid line) obtained from the image in Figure 2.21. Either profile may be taken to represent the LSF of the gamma camera.

tance, and the intervening medium. The LSF was used to estimate the PSF as explained above. The FWHM of the PSF of the SPECT system studied was observed to vary between 1.3 *cm* and 3.0 *cm*.

See Section 2.12 for illustrations of the ESF and LSF of a μ CT imaging system. See Chapter 10 for descriptions of methods for deblurring images.

2.10 Resolution

The spatial resolution of an imaging system or an image may be expressed in terms of the following:

- The sampling interval (in, for example, *mm* or μ *m*).
- The width of (a profile of) the PSF, usually FWHM (in *mm*).
- The size of the laser spot used to obtain the digital image by scanning an original film, or the size of the solid-state detector used to obtain the digital image (in μ *m*).
- The smallest visible object or separation between objects in the image (in *mm* or μ *m*).
- The finest grid pattern that remains visible in the image (in *lp/mm*).

The typical resolution limits of a few imaging systems are [6]:

- X-ray film: 25 – 100 *lp/mm*.
- screen-film combination: 5 – 10 *lp/mm*;
mammography: up to 20 *lp/mm*.
- CT: 0.7 *lp/mm*;
 μ CT: 50 *lp/mm* or 10 μ *m*;
- SPECT: < 0.1 *lp/mm*.

2.11 The Fourier Transform and Spectral Content

The Fourier transform is a linear, reversible transform that maps an image from the space domain to the frequency domain. Converting an image from the spatial to the frequency (Fourier) domain helps in assessing the spectral content and energy distribution over frequency bands. Sharp edges in the

image domain are associated with large proportions of high-frequency content. Oriented patterns in the space domain correspond to increased energy in bands of frequency in the spectral domain with the corresponding orientation. Simple geometric patterns such as rectangles and circles map to recognizable functions in the frequency domain, such as the sinc and Bessel functions, respectively. Transforming an image to the frequency domain assists in the application of frequency-domain filters to remove noise, enhance the image, or extract certain components that are better separated in the frequency domain than in the space domain.

The 2D Fourier transform of an image $f(x, y)$, denoted by $F(u, v)$, is given by [8, 9, 11, 131]

$$F(u, v) = \int_{x=-\infty}^{\infty} \int_{y=-\infty}^{\infty} f(x, y) \exp[-j 2\pi (ux + vy)] dx dy. \quad (2.40)$$

The variables u and v represent frequency in the horizontal and vertical directions, respectively. (The frequency variable in image analysis is often referred to as *spatial frequency* to avoid confusion with temporal frequency; we will, however, not use this terminology in this book.) Recall that the complex exponential is a combination of the 2D sine and cosine functions and is separable, as

$$\begin{aligned} & \exp[-j 2\pi (ux + vy)] \\ &= \exp(-j 2\pi ux) \exp(-j 2\pi vy) \\ &= [\cos(2\pi ux) - j \sin(2\pi ux)] [\cos(2\pi vy) - j \sin(2\pi vy)]. \end{aligned} \quad (2.41)$$

Images are typically functions of space; hence, the units of measurement in the image domain are m , cm , mm , μm , etc. In the 2D Fourier domain, the unit of frequency is *cycles/mm*, *cycles/m*, mm^{-1} , etc. Frequency is also expressed as lp/mm . If the distance to the viewer is taken into account, frequency could be expressed in terms of *cycles/degree* of the visual angle subtended at the viewer's eye. The unit *Hertz* is not used in 2D Fourier analysis.

In computing the Fourier transform, it is common to use the discrete Fourier transform (DFT) via the fast Fourier transform (FFT) algorithm. The 2D DFT of a digital image $f(m, n)$ of size $M \times N$ pixels is defined as

$$F(k, l) = \frac{1}{MN} \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} f(m, n) \exp \left[-j 2\pi \left(\frac{mk}{M} + \frac{nl}{N} \right) \right]. \quad (2.42)$$

For complete recovery of $f(m, n)$ from $F(k, l)$, the latter should be computed for $k = 0, 1, \dots, M - 1$, and $l = 0, 1, \dots, N - 1$, at the minimum [7, 8, 9]. Then, the inverse transform gives back the original image with no error or loss of information as

$$f(m, n) = \sum_{k=0}^{M-1} \sum_{l=0}^{N-1} F(k, l) \exp \left[+j 2\pi \left(\frac{mk}{M} + \frac{nl}{N} \right) \right], \quad (2.43)$$

for $m = 0, 1, \dots, M - 1$, and $n = 0, 1, \dots, N - 1$. This expression may be interpreted as resolving the given image into a weighted sum of mutually orthogonal exponential (or sinusoidal) basis functions. The eight sine functions, for $k = 0, 1, 2, \dots, 7$, that form the imaginary part of the basis functions of the 1D DFT for $M = 8$ are shown in Figure 2.23. Figures 2.24 and 2.25 show the first 64 cosine and sine basis functions (for $k, l = 0, 1, 2, \dots, 7$) that are the components of the 2D exponential function in Equation 2.43.





FIGURE 2.23

The first eight sine basis functions of the 1D DFT; $k = 0, 1, 2, \dots, 7$ from top to bottom. Each function was computed using 64 samples.

In order to use the FFT algorithm, it is common to pad the given image with zeros or some other appropriate background value and convert the image to a square of size $N \times N$ where N is an integral power of 2. Then, all indices in Equation 2.42 may be made to run from 0 to $N - 1$ as

$$F(k, l) = \frac{1}{N} \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} f(m, n) \exp \left[-j \frac{2\pi}{N} (mk + nl) \right], \quad (2.44)$$

**FIGURE 2.24**

The first 64 cosine basis functions of the 2D DFT. Each function was computed using a 64×64 matrix.




FIGURE 2.25

The first 64 sine basis functions of the 2D DFT. Each function was computed using a 64×64 matrix.

with $k = 0, 1, \dots, N - 1$, and $l = 0, 1, \dots, N - 1$. The inverse transform is given as

$$f(m, n) = \frac{1}{N} \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} F(k, l) \exp \left[+j \frac{2\pi}{N} (mk + nl) \right]. \quad (2.45)$$

In Equations 2.44 and 2.45, the normalization factor has been divided equally between the forward and inverse transforms to be $\frac{1}{N}$ for the sake of symmetry [8].

Example — the rectangle function and its Fourier transform: A 2D function with a rectangular base of size $X \times Y$ and height A is defined as

$$\begin{aligned} f(x, y) &= A \quad \text{if } 0 \leq x \leq X; 0 \leq y \leq Y \\ &= 0 \quad \text{otherwise.} \end{aligned} \quad (2.46)$$

The 1D version of the rectangle function is also known as the gate function. The 2D Fourier transform of the rectangle function above is given by

$$F(u, v) = AXY \left[\frac{\sin(\pi uX)}{\pi uX} \exp(-j\pi uX) \right] \left[\frac{\sin(\pi vY)}{\pi vY} \exp(-j\pi vY) \right]. \quad (2.47)$$

Observe that the Fourier transform of a real image is, in general, a complex function. However, an image with even symmetry about the origin will have a real Fourier transform. The $\exp[\]$ functions in Equation 2.47 indicate the phase components of the spectrum.

A related function that is commonly used is the rect function, defined as

$$\text{rect}(x, y) = \begin{cases} 1 & \text{if } |x| < \frac{1}{2}, |y| < \frac{1}{2} \\ 0 & \text{if } |x| > \frac{1}{2}, |y| > \frac{1}{2}. \end{cases} \quad (2.48)$$

The Fourier transform of the rect function is the sinc function:

$$\text{rect}(x, y) \Leftrightarrow \text{sinc}(u, v), \quad (2.49)$$

where

$$\text{sinc}(u, v) = \text{sinc}(u) \text{sinc}(v) = \frac{\sin(\pi u)}{\pi u} \frac{\sin(\pi v)}{\pi v}, \quad (2.50)$$

and \Leftrightarrow indicates that the two functions form a forward and inverse Fourier-transform pair.

Figure 2.26 shows three images with rectangular (square) objects and their Fourier log-magnitude spectra. Observe that the smaller the box, the greater the energy content in the higher-frequency areas of the spectrum. At the limits, we have the Fourier transform of an image of an infinitely large rectangle, that is, the transform of an image with a constant value of unity for all space, equal to $\delta(0, 0)$; and the Fourier transform of an image with an infinitesimally

small rectangle, that is, an impulse, equal to a constant of unity, representing a “white” spectrum. The frequency axes have been shifted such that $(u, v) = (0, 0)$ is at the center of the spectrum displayed. The frequency coordinates in this mode of display of image spectra are shown in Figure 2.27 (b). Figure 2.28 shows the log-magnitude spectrum in Figure 2.26 (f) with and without shifting; the shifted (or centered or folded) mode of display as in Figure 2.28 (b) is the preferred mode of display of 2D spectra.

The rectangle image in Figure 2.26 (e) as well as its magnitude spectrum are also shown as mesh plots in Figure 2.29. The mesh plot demonstrates more clearly the sinc nature of the spectrum.

Figure 2.30 shows three images with rectangular boxes oriented at 0° , 90° , and 135° , and their log-magnitude spectra. The sinc functions in the Fourier domain in Figure 2.30 are not symmetric in the u and v coordinates, as was the case in the spectra of the square boxes in Figure 2.26. The narrowing of the rectangle along a spatial axis results in the widening of the lobes of the sinc function and the presence of increased high-frequency energy along the corresponding frequency axis. The rotation of an image in the spatial domain results in a corresponding rotation in the Fourier domain.


Example — the circle function and its Fourier transform: Circular apertures and functions are encountered often in imaging and image processing. The circ function, which represents a circular disc or aperture, is defined as

$$\text{circ}(r) = \begin{cases} 1 & \text{if } r < 1 \\ 0 & \text{if } r > 1, \end{cases} \quad (2.51)$$


where $r = \sqrt{(x^2 + y^2)}$. The Fourier transform of $\text{circ}(r)$ may be shown to be $\frac{1}{\nu} J_1(2\pi\nu)$, where $\nu = \sqrt{(u^2 + v^2)}$ represents radial frequency in the 2D (u, v) plane, and J_1 is the first-order Bessel function of the first kind [3, 9].

Figure 2.31 shows an image of a circular disc and its log-magnitude spectrum. The disc image as well as its magnitude spectrum are also shown as mesh plots in Figure 2.32. Ignoring the effects due to the representation of the circular shape on a discrete grid, both the image and its spectrum are isotropic. Figure 2.33 shows two profiles of the log-magnitude spectrum in Figure 2.31 (b) taken along the central horizontal axis. The nature of the Bessel function is clearly seen in the 1D plots; the conjugate symmetry of the spectrum is also readily seen in the plot in Figure 2.33 (a). In displaying profiles of 2D system transfer functions, it is common to show only one half of the profile for positive frequencies, as in Figure 2.33 (b). If such a profile is shown, it is to be assumed that the system possesses axial or rotational symmetry; that is, the system is isotropic.


Examples of Fourier spectra of biomedical images: Figure 2.34 shows two TEM images of collagen fibers in rabbit ligament samples (in cross-section), and their Fourier spectra. The Bessel characteristics of the spectrum due to the circular shape of the objects in the image are clearly

**FIGURE 2.26**


(a) Rectangle image, with total size 128×128 pixels and a rectangle (square) of size 40×40 pixels. (b) Log-magnitude spectrum of the image in (a). (c) Rectangle size 20×20 pixels. (d) Log-magnitude spectrum of the image in (c). (e) Rectangle size 10×10 pixels. (f) Log-magnitude spectrum of the image in (e). The spectra have been scaled to map the range $[5, 12]$ to the display range $[0, 255]$. See also Figures 2.28 and 2.29.

**FIGURE 2.27**


Frequency coordinates in (a) the unshifted mode and (b) the shifted mode of display of image spectra. U and V represent the sampling frequencies along the two axes. Spectra of images with real values possess conjugate symmetry about $U/2$ and $V/2$. Spectra of sampled images are periodic, with the periods equal to U and V along the two axes. It is common practice to display one complete period of the shifted spectrum, including the conjugate symmetric parts, as in (b). See also Figure 2.28.

**FIGURE 2.28**

(a) Log-magnitude spectrum of the rectangle image in Figure 2.26 (e) without shifting. Most FFT routines provide spectral data in this format. (b) The spectrum in (a) shifted or folded such that $(u, v) = (0, 0)$ is at the center. It is common practice to display one complete period of the shifted spectrum, including the conjugate symmetric parts, as in (b). See also Figure 2.27.




(a)



(b)

FIGURE 2.29

(a) Mesh plot of the rectangle image in Figure 2.26 (e), with total size 128×128 pixels and a rectangle (square) of size 10×10 pixels. (b) Magnitude spectrum of the image in (a).

**FIGURE 2.30**

(a) Rectangle image, with total size 128×128 pixels and a rectangle of size 10×40 pixels. (b) Log-magnitude spectrum of the image in (a). (c) Rectangle size 40×10 pixels; this image may be considered to be that in (a) rotated by 90° . (d) Log-magnitude spectrum of the image in (c). (e) Image in (c) rotated by 45° using nearest-neighbor selection. (f) Log-magnitude spectrum of the image in (e). Spectra scaled to map $[5, 12]$ to the display range $[0, 255]$. See also Figure 8.1.




FIGURE 2.31



(a) Image of a circular disc. The radius of the disc is 10 pixels; the size of the image is 128×128 pixels. (b) Log-magnitude spectrum of the image in (a). See also Figures 2.32 and 2.33.

seen in Figure 2.34 (d). (Compare the examples in Figure 2.34 with those in Figure 2.31.)


Figure 2.35 shows two SEM images of collagen fibers as seen in freeze-fractured surfaces of rabbit ligament samples, and their Fourier spectra. The highly oriented and piecewise linear (rectangular) characteristics of the fibers in the normal sample in Figure 2.35 (a) are indicated by the concentrations of energy along radial lines at the corresponding angles in the spectrum in Figure 2.35 (b). The scar sample in Figure 2.35 (c) lacks directional preference, which is reflected in its spectrum in Figure 2.35 (d). (Compare the examples in Figure 2.35 with those in Figure 2.30.)

2.11.1 Important properties of the Fourier transform


The Fourier transform is a linear, reversible transform that maps an image from the space domain to the frequency domain. The spectrum of an image can provide useful information on the frequency content of the image, on the presence of oriented or directional elements, on the presence of specific image patterns, and on the presence of noise. A study of the spectrum of an image can assist in the development of filtering algorithms to remove noise, in the design of algorithms to enhance the image, and in the extraction of features for pattern recognition. Some of the important properties of the Fourier transform are described in the following paragraphs with illustrations as required [9, 8, 11]; both the discrete and continuous representations of functions are used as appropriate or convenient.

**FIGURE 2.32**

(a) Mesh plot of the circular disc in Figure 2.31 (a). The radius of the disc is 10 pixels; the size of the image is 128×128 pixels. (b) Magnitude spectrum of the image in (a).




(a)



(b)

FIGURE 2.33

(a) Profile of the log-magnitude spectrum in Figure 2.31 (b) along the central horizontal axis. (b) Profile in (a) shown only for positive frequencies. The frequency axis is indicated in samples; the true frequency values depend upon the sampling frequency.

**FIGURE 2.34**

(a) TEM image of collagen fibers in a normal rabbit ligament sample. (b) Log-magnitude spectrum of the image in (a). (c) TEM image of collagen fibers in a scar tissue sample. (d) Log-magnitude spectrum of the image in (c). See also Figure 1.5 and Section 1.4.




FIGURE 2.35

(a) SEM image of collagen fibers in a normal rabbit ligament sample. (b) Log-magnitude spectrum of the image in (a). (c) SEM image of collagen fibers in a scar tissue sample. (d) Log-magnitude spectrum of the image in (c). See also Figure 1.8 and Section 1.4.

1. The kernel function of the Fourier transform is separable and symmetric. This property facilitates the evaluation of the 2D DFT as a set of 1D row transforms, followed by a set of 1D column transforms. We have

$$F(k, l) = \frac{1}{N} \sum_{m=0}^{N-1} \exp\left(-j\frac{2\pi}{N}mk\right) \sum_{n=0}^{N-1} f(m, n) \exp\left(-j\frac{2\pi}{N}nl\right). \quad (2.52)$$

1D FFT routines may be used to obtain 2D and multidimensional Fourier transforms in the following manner:

$$F(m, l) = N \left[\frac{1}{N} \sum_{n=0}^{N-1} f(m, n) \exp\left(-j\frac{2\pi}{N}nl\right) \right], \quad (2.53)$$

$$F(k, l) = \frac{1}{N} \sum_{m=0}^{N-1} F(m, l) \exp\left(-j\frac{2\pi}{N}mk\right). \quad (2.54)$$

(Care should be taken to check if the factor $\frac{1}{N}$ is included in the forward or inverse 1D FFT routine, where required.)

2. The Fourier transform is an energy-conserving transform, that is,

$$\int_{x=-\infty}^{\infty} \int_{y=-\infty}^{\infty} |f(x, y)|^2 dx dy = \int_{u=-\infty}^{\infty} \int_{v=-\infty}^{\infty} |F(u, v)|^2 du dv. \quad (2.55)$$

This relationship is known as Parseval's theorem.

3. The inverse Fourier transform operation may be performed using the same FFT routine by taking the forward Fourier transform of the complex conjugate of the given function, and then taking the complex conjugate of the result.
4. The Fourier transform is a linear transform. The Fourier transform of the sum of two images is the sum of the Fourier transforms of the individual images.

Images are often corrupted by additive noise, such as

$$g(x, y) = f(x, y) + \eta(x, y). \quad (2.56)$$

Upon Fourier transformation, we have

$$G(u, v) = F(u, v) + \eta(u, v). \quad (2.57)$$

Most real-life images have a large portion of their energy concentrated around $(u, v) = (0, 0)$ in a low-frequency region; however, the presence

of edges, sharp features, and small-scale or fine details leads to increased strength of high-frequency components (see Figure 2.34). On the other hand, random noise has a spectrum that is equally spread all over the frequency space (that is, a flat, uniform, or “white” spectrum). Indiscriminate removal of high-frequency components could cause blurring of edges and the loss of the fine details in the image.

5. The DFT and its inverse are periodic signals:

$$F(k, l) = F(k \pm \alpha N, l) = F(k, l \pm \alpha N) = F(k \pm \alpha N, l \pm \beta N), \quad (2.58)$$

where α and β are integers.

6. The Fourier transform is conjugate-symmetric for images with real values:

$$F(-k, -l) = F^*(k, l). \quad (2.59)$$

It follows that $|F(-k, -l)| = |F(k, l)|$ and $\angle F(-k, -l) = -\angle F(k, l)$; that is, the magnitude spectrum is even symmetric and the phase spectrum is odd symmetric. The symmetry of the magnitude spectrum is illustrated by the examples in Figures 2.26 and 2.30.

7. A spatial shift or translation applied to an image leads to an additional linear phase component in its Fourier transform; the magnitude spectrum is unaffected. If $f(m, n) \Leftrightarrow F(k, l)$ are a Fourier-transform pair, we have

$$f(m - m_o, n - n_o) \Leftrightarrow F(k, l) \exp \left[-j \frac{2\pi}{N} (km_o + ln_o) \right], \quad (2.60)$$

where (m_o, n_o) is the shift applied in the space domain.

Conversely, we also have

$$f(m, n) \exp \left[j \frac{2\pi}{N} (k_o m + l_o n) \right] \Leftrightarrow F(k - k_o, l - l_o). \quad (2.61)$$

This property has important implications in the modulation of 1D signals for transmission and communication [1]; however, it does not have a similar application with 2D images.

8. $F(0, 0)$ gives the average value of the image; a scale factor may be required depending upon the definition of the DFT used.
9. For display purposes, $\log_{10}[1 + |F(k, l)|^2]$ is often used; the addition of unity (to avoid taking the log of zero), and the squaring may sometimes be dropped. It is also common to fold or shift the spectrum to bring the $(0, 0)$ frequency point (the “DC” point) to the center, and the folding frequency (half of the sampling frequency) components to the

edges. Figures 2.26, 2.27, and 2.28 illustrate shifted spectra and the corresponding frequency coordinates.

Folding of the spectrum could be achieved by multiplying the image $f(m, n)$ with $(-1)^{(m+n)}$ before the FFT is computed [8]. Because the indices m and n are integers, this amounts to merely changing the signs of alternate pixels. This outcome is related to the property in Equation 2.61 with $k_o = l_o = N/2$, which leads to

$$\exp\left[j\frac{2\pi}{N}(k_o m + l_o n)\right] = \exp[j\pi(m+n)] = (-1)^{(m+n)}, \quad (2.62)$$

and

$$f(m, n) (-1)^{(m+n)} \Leftrightarrow F(k - N/2, l - N/2). \quad (2.63)$$

10. Rotation of an image leads to a corresponding rotation of the Fourier spectrum.

$$f(m_1, n_1) \Leftrightarrow F(k_1, l_1), \quad (2.64)$$

where

$$m_1 = m \cos \theta + n \sin \theta; \quad n_1 = -m \sin \theta + n \cos \theta; \quad (2.65)$$

$$k_1 = k \cos \theta + l \sin \theta; \quad l_1 = -k \sin \theta + l \cos \theta. \quad (2.66)$$

This property is illustrated by the images and spectra in Figure 2.30, and is useful in the detection of directional or oriented patterns (see Chapter 8).

11. Scaling an image leads to an inverse scaling of its Fourier transform:

$$f(am, bn) \Leftrightarrow \frac{1}{|ab|} F\left(\frac{k}{a}, \frac{l}{b}\right), \quad (2.67)$$

where a and b are scalar scaling factors. The shrinking of an image leads to an expansion of its spectrum, with increased high-frequency content. On the contrary, if an image is enlarged, its spectrum is shrunk, with reduced high-frequency energy. The images and spectra in Figure 2.26 illustrate this property.

12. **Linear shift-invariant systems and convolution:** Most imaging systems may be modeled as linear and shift-invariant or position-invariant systems that are completely characterized by their PSFs. The output of such a system is given as the convolution of the input image with the PSF:

$$\begin{aligned} g(m, n) &= h(m, n) * f(m, n) \\ &= \sum_{\alpha=0}^{N-1} \sum_{\beta=0}^{N-1} h(\alpha, \beta) f(m - \alpha, n - \beta). \end{aligned} \quad (2.68)$$

Upon Fourier transformation, the convolution maps to the multiplication of the two spectra:

$$G(k, l) = H(k, l) F(k, l). \quad (2.69)$$

Thus, we have the important property

$$h(x, y) * f(x, y) \Leftrightarrow H(u, v) F(u, v), \quad (2.70)$$

expressed now in the continuous coordinates (x, y) and (u, v) . The characterization of imaging systems in the transform domain is discussed in Section 2.12.

It should be noted that the convolution \Leftrightarrow multiplication property with the DFT implies periodic or circular convolution; however, this type of convolution may be made to be equivalent to linear convolution by zero-padding. Details on this topic are presented in Section 3.5.3.

13. Multiplication of images in the space domain is equivalent to the convolution of their Fourier transforms:

$$f_1(x, y) f_2(x, y) \Leftrightarrow F_1(u, v) * F_2(u, v). \quad (2.71)$$

In medical imaging, some types of noise get multiplied with the image.

When a transparency, such as an X-ray image on film, is viewed using a light box, the resulting image $g(x, y)$ may be modeled as the product of the transparency or transmittance function $f(x, y)$ with the light source intensity field $s(x, y)$, giving $g(x, y) = f(x, y) s(x, y)$. If $s(x, y)$ is absolutely uniform with a value A , its Fourier transform will be an impulse: $S(u, v) = A\delta(u, v)$. The convolution of $F(u, v)$ with $A\delta(u, v)$ will have no effect on the spectrum except scaling by the constant A . If the source is not uniform, the viewed image will be a distorted version of the original; the corresponding convolution $G(u, v) = F(u, v) * S(u, v)$ will distort the spectrum $F(u, v)$ of the original image.

14. The correlation of two images $f(m, n)$ and $g(m, n)$ is given by the operation

$$\gamma_{f,g}(\alpha, \beta) = \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} f(m, n) g(m + \alpha, n + \beta). \quad (2.72)$$

Correlation is useful in the comparison of images where features that are common to the images may be present with a spatial shift (α, β) .

Upon Fourier transformation, we get the conjugate product of the spectra of the two images:

$$\Gamma_{f,g}(k, l) = F(k, l) G^*(k, l). \quad (2.73)$$

A related measure, known as the *correlation coefficient* and useful in template matching and image classification, is defined as

$$\gamma = \frac{\sum_{m=0}^{N-1} \sum_{n=0}^{N-1} f(m, n) g(m, n)}{\left[\sum_{m=0}^{N-1} \sum_{n=0}^{N-1} f^2(m, n) \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} g^2(m, n) \right]^{\frac{1}{2}}}. \quad (2.74)$$

Here, it is assumed that the two images f and g (or features thereof) are aligned and registered, and are of the same scale and orientation.

15. Differentiation of an image results in the extraction of edges and high-pass filtering:

$$\begin{aligned} \frac{\partial f(x, y)}{\partial x} &\Leftrightarrow j2\pi u F(u, v); \\ \frac{\partial f(x, y)}{\partial y} &\Leftrightarrow j2\pi v F(u, v). \end{aligned} \quad (2.75)$$

The gain of the operator increases linearly with frequency u or v .

When processing digital images, the derivatives are approximated by differences computed as

$$\begin{aligned} f'_y(m, n) &\approx f(m, n) - f(m - 1, n), \\ f'_x(m, n) &\approx f(m, n) - f(m, n - 1), \end{aligned} \quad (2.76)$$

(using matrix notation).

It should be noted that operators based upon differences could cause negative pixel values in the result. In order to display the result as an image, it will be necessary to map the full range of the pixel values, including the negative values, to the display range available. The magnitude of the result may also be displayed if the sign of the result is not important.

Examples: Figure 2.36 shows an image of a rectangle and its derivatives in the horizontal and vertical directions, as well as their log-magnitude spectra. The horizontal and vertical derivatives were obtained by convolving the image with $[-1, 1]$ and $[-1, 1]^T$, respectively. Figures 2.37 and 2.38 show similar sets of results for the image of a myocyte and an MR image of a knee. It is seen that the two derivatives extract edges in the corresponding directions; edges in the direction orthogonal to that of the operator are removed. The spectra show that the components in one direction are enhanced, whereas the components in the orthogonal direction are removed.

Observe that differentiation results in the removal of the intensity information from the image. Correspondingly, the values of the spectrum for $u = 0$ or $v = 0$ are set to zero.

16. The *Laplacian* of an image is defined as

$$\nabla^2 f(x, y) = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}. \quad (2.77)$$

In the Fourier domain, we get

$$\nabla^2 f(x, y) \Leftrightarrow -(2\pi)^2(u^2 + v^2)F(u, v). \quad (2.78)$$

The spectrum of the image is multiplied by the factor $(u^2 + v^2)$, which is isotropic and increases quadratically with frequency. Therefore, the high-frequency components are amplified by this operation. The Laplacian is an omnidirectional operator, and detects edges in all directions.

When processing digital images, the second derivatives may be approximated as follows: Taking the derivative of the expression for $f'_y(m, n)$ in Equation 2.76 for the second time, we get

$$\begin{aligned} f''_y(m, n) &\approx f(m, n) - f(m-1, n) - [f(m-1, n) - f(m-2, n)] \\ &= f(m, n) - 2f(m-1, n) + f(m-2, n) \end{aligned} \quad (2.79)$$

(using matrix notation). Causality is usually not a matter of concern in image processing, and it is often desirable to have operators use collections of pixels that are centered about the pixel being processed. Applying a shift of one pixel to the result above (specifically, adding 1 to the first index of each term) leads to

$$\begin{aligned} f''_y(m, n) &\approx f(m+1, n) - 2f(m, n) + f(m-1, n) \\ &= f(m-1, n) - 2f(m, n) + f(m+1, n). \end{aligned} \quad (2.80)$$

Similarly, we get

$$f''_x(m, n) \approx f(m, n-1) - 2f(m, n) + f(m, n+1). \quad (2.81)$$

The Laplacian could then be implemented as

$$f_L(m, n) = f(m-1, n) + f(m, n-1) - 4f(m, n) + f(m+1, n) + f(m, n+1). \quad (2.82)$$

This operation is achieved by convolving the image with the 3×3 mask or operator

$$\begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix}. \quad (2.83)$$

Examples: Figure 2.39 shows the Laplacian of the rectangle image in Figure 2.36 (a) and its log-magnitude spectrum. Similar results are shown in Figures 2.40 and 2.41 for the myocyte image in Figure 2.37 (a) and the knee MR image in Figure 2.38 (a). The Laplacian operator

has extracted all edges in all directions; correspondingly, high-frequency components in all directions in the spectrum have been strengthened. Observe that the images have lost gray-scale or intensity information; correspondingly, the $(u, v) = (0, 0)$ component has been removed from the spectra.

17. Integration of an image leads to smoothing or blurring, and lowpass filtering:

$$\int_{\alpha=-\infty}^x f(\alpha, y) d\alpha \Leftrightarrow \frac{1}{j2\pi u} F(u, v), \quad (2.84)$$

$$\int_{\beta=-\infty}^y f(x, \beta) d\beta \Leftrightarrow \frac{1}{j2\pi v} F(u, v). \quad (2.85)$$

The weighting factors that apply to $F(u, v)$ diminish with increasing frequency, and hence high-frequency components are attenuated by this operation.

The integration of an image from $-\infty$ to the current x or y position is seldom encountered in practice. Instead, it is common to encounter the integration of an image over a small region or aperture surrounding the current position, in the form

$$g(x, y) = \frac{1}{AB} \int_{\alpha=-A/2}^{A/2} \int_{\beta=-B/2}^{B/2} f(x + \alpha, y + \beta) d\alpha d\beta, \quad (2.86)$$

where the region of integration is a rectangle of size $A \times B$. The normalization factor $\frac{1}{AB}$ leads to the average intensity being computed over the area of integration. This operation may be interpreted as a moving-average (MA) filter.

In discrete terms, averaging over a 3×3 aperture or neighborhood is represented as

$$g(m, n) = \frac{1}{9} \sum_{\alpha=-1}^1 \sum_{\beta=-1}^1 f(m + \alpha, n + \beta). \quad (2.87)$$

This equation may be expanded as

$$g(m, n) = \frac{1}{9} \times \left[\begin{array}{ccc} f(m-1, n-1) + f(m-1, n) + f(m-1, n+1) \\ + f(m, n-1) + f(m, n) + f(m, n+1) \\ + f(m+1, n-1) + f(m+1, n) + f(m+1, n+1) \end{array} \right]. \quad (2.88)$$

The same operation is also achieved via convolution of the image $f(m, n)$ with the array

$$\frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}, \quad (2.89)$$

which may be viewed as the PSF of a filter. It follows that the corresponding effect in the frequency domain is multiplication of the Fourier transform of the image with a 2D sinc function.

Integration or averaging as above but only along the horizontal or vertical directions may be performed via convolution with the arrays $\frac{1}{3} [1, 1, 1]$ or $\frac{1}{3} [1, 1, 1]^T$, respectively.

Examples: Figure 2.42 shows an image of a rectangle with ideal edges, followed by the results of averaging along the horizontal and vertical directions via convolution with the arrays $\frac{1}{3} [1, 1, 1]$ and $\frac{1}{3} [1, 1, 1]^T$, respectively; the log-magnitude spectra of the images are also shown. Figure 2.43 shows the result of averaging the rectangle image using the 3×3 mask in Equation 2.89, as well as its spectrum. It is seen that averaging results in the smoothing of edges and a reduction in the strength of the high-frequency components in the direction(s) of averaging.

Similar results are shown in Figures 2.44 and 2.45 for an image of a myocyte, and in Figures 2.46 and 2.47 for an MR image of a knee joint. It is seen that minor details and artifacts in the images have been suppressed or removed by the averaging operation.

2.12 Modulation Transfer Function


Analysis of the characteristics of imaging systems, treated as 2D LSI systems, is easier in the frequency or Fourier domain. Taking the 2D Fourier transform of the convolution integral in Equation 2.32, we get

$$G(u, v) = H(u, v) F(u, v), \quad (2.90)$$

where $F(u, v)$, $G(u, v)$, and $H(u, v)$ are the 2D Fourier transforms of $f(x, y)$, $g(x, y)$, and $h(x, y)$, respectively.

The 2D frequency-domain function $H(u, v)$ is known as the *optical transfer function* (OTF), or simply the transfer function, of the imaging system. The OTF is, in general, a complex quantity. The magnitude of the OTF is known as the *modulation transfer function* (MTF). The OTF at each frequency coordinate gives the attenuation (or gain) for the corresponding frequency (u, v) as well as the phase introduced by the system (usually implying distortion).

The OTF of an imaging system may be estimated by compiling its responses to sinusoidal test patterns or gratings of various frequencies; it may also be computed from various spread functions, as described in Section 2.9. If the

**FIGURE 2.36**

(a) Image of a rectangular box. (c) Horizontal and (e) vertical derivatives of the image in (a), respectively. (b), (d), and (f): Log-magnitude spectra of the images in (a), (c), and (e), respectively. The images in (c) and (e) were obtained by mapping the range $[-200, 200]$ to the display range of $[0, 255]$. Negative differences appear in black, positive differences in white. The spectra show values in the range $[5, 12]$ mapped to $[0, 255]$.






FIGURE 2.37


(a) Image of a myocyte. (c) Horizontal and (e) vertical derivatives of the image in (a), respectively. (b), (d), and (f): Log-magnitude spectra of the images in (a), (c), and (e), respectively. Images in (c) and (e) were obtained by mapping the range $[-20, 20]$ to the display range of $[0, 255]$. The spectra show values in the range $[3, 12]$ mapped to $[0, 255]$.

**FIGURE 2.38**

(a) MR image of a knee. (c) Horizontal and (e) vertical derivatives of the image in (a), respectively. (b), (d), and (f): Log-magnitude spectra of the images in (a), (c), and (e), respectively. The images in (c) and (e) were obtained by mapping the range $[-50, 50]$ to the display range of $[0, 255]$. Negative differences appear in black, positive differences in white. The spectra show values in the range $[3, 12]$ mapped to $[0, 255]$.

**FIGURE 2.39**

(a) Laplacian of the rectangle image in Figure 2.36 (a). (b) Log-magnitude spectrum of the image in (a).

**FIGURE 2.40**

(a) Laplacian of the myocyte image in Figure 2.37 (a). (b) Log-magnitude spectrum of the image in (a).




FIGURE 2.41


(a) Laplacian of the MR image in Figure 2.38 (a). (b) Log-magnitude spectrum of the image in (a).

imaging system may be approximated as an LSI system, its response or output corresponding to any arbitrary input image may be determined with a knowledge of its PSF or OTF, as per Equation 2.32 or 2.90.


It should be noted that the widths of a PSF and the corresponding MTF bear an inverse relationship: the greater the blur, the wider the PSF, and the narrower the MTF (more high-frequency components are attenuated significantly). Therefore, resolution may also be expressed indirectly in the frequency domain as a point along the frequency axis beyond which the attenuation is significant. Furthermore, a larger area under the (normalized) MTF indicates a system with better resolution (more high-frequency components preserved) than a system with a smaller area under the MTF. Several MTF-based measures related to image quality are described in Section 2.15.

Example: Higashida et al. [133] compared the MTF of an image-intensifier system for DR with that of a screen-film system. A 2,100-line camera was used to obtain images in $2,048 \times 2,048$ matrices. The distance between the X-ray tube focal spot and the experimental table top in their studies was 100 cm; the distance between the camera (object plane) and the table top was 10 cm; the distance between the screen-film cassette and the table top was 5 cm. The magnification factors for the DR and screen-film systems are 1.28 and 1.17, respectively. Focal spots of nominal size 0.3 mm and 0.8 mm were used. MTFs were computed for different imaging conditions by capturing images of a narrow slit to obtain the LSF, and computing its Fourier spectrum. The effect of the MTF of the recording system was also included in computing the total MTF: the total MTF was computed as the product of the MTF due to the focal spot and the MTF of the detector system.

Figure 2.48 shows the MTFs for three imaging conditions. It is seen that the two MTFs of the DR system are poorer than that of the screen-film system

**FIGURE 2.42**

(a) Image of a rectangular box. Results of averaging using three pixels in the (c) horizontal and (e) vertical directions, respectively. (b), (d), and (f): Log-magnitude spectra of the images in (a), (c), and (e), respectively. The spectra show values in the range $[5, 12]$ mapped to $[0, 255]$.


**FIGURE 2.43**

(a) Result of 3×3 averaging of the rectangle image in Figure 2.42 (a). (b) Log-magnitude spectrum of the image in (a).


at spatial frequency beyond 1 *cycle/mm* in the object plane. The MTF of the DR system with a focal spot of 0.3 *mm* is slightly better than that of the same system with a focal spot of 0.8 *mm*. The poorer performance of the DR system was attributed to the pixel size being as large as 0.11 *mm* (the equivalent resolution being 4.5 *lp/mm*).

Higashida et al. also computed contrast-detail curves based upon experiments with images of square objects of varying thickness. The threshold object thickness was determined as that related to the lowest-contrast image where radiologists could visually detect the objects in images with a 50% confidence level. (*Note:* The background and object area remaining the same, the contrast of the object in an X-ray image increases as the object thickness is increased.) Figure 2.49 shows the contrast-detail curves for the screen-film system and the DR system, with the latter operated at the same dose as the former in one setting (labeled as iso-dose in the figure, with the focal spot being 0.8 *mm*), and at a low-dose setting with the focal spot being 0.3 *mm* in another setting. The performance in the detection of low-contrast signals with the screen-film system was comparable to that with the DR system at the same X-ray dose. The low-dose setting resulted in poorer detection capability with the DR system. In the study of Higashida et al., in spite of the poorer MTF, the DR system was found to be as effective as the screen-film system (at the same X-ray dose) in clinical studies in the application area of bone radiography.

Example: Figure 2.50 shows the MTF curve of an amorphous selenium detector system for direct digital mammography along with those of a screen-film system and an indirect digital imaging system. The amorphous selenium system has the best MTF of the three systems.

**FIGURE 2.44**

(a) Image of a myocyte. Results of averaging using three pixels in the (c) horizontal and (e) vertical directions, respectively. (b), (d), and (f): Log-magnitude spectra of the images in (a), (c), and (e), respectively. The spectra show values in the range $[3, 12]$ mapped to $[0, 255]$.

**FIGURE 2.45**

(a) Result of 3×3 averaging of the myocyte image in Figure 2.44 (a). (b) Log-magnitude spectrum of the image in (a).


Example: Pateyron et al. [134] developed a μ CT system for high-resolution 3D imaging of small bone samples using synchrotron radiation. In order to evaluate the resolution of the system, they imaged a sharp edge using the system, thereby obtaining the ESF of the system illustrated in Figure 2.51 (a). The derivative of the ESF in the directional orthogonal to that of the edge was then computed to obtain the LSF, shown in Figure 2.51 (b). The MTF of the system was then derived as explained earlier in this section, using the Fourier transform of the LSF, and is shown in Figure 2.51 (c). The value of the MTF at 55 lp/mm is 0.1. Using this information, Pateyron et al. estimated the spatial resolution of their μ CT system to be $(2 \times 55)^{-1} = 0.009 \text{ mm}$ or $9 \mu\text{m}$.

2.13 Signal-to-Noise Ratio


Noise is omnipresent! Some of the sources of random noise in biomedical imaging are scatter, photon-counting noise, and secondary radiation. Blemishes may be caused by scratches, defects, and nonuniformities in screens, crystals, and electronic detectors. It is common to assume that noise is additive, and to express the degraded image $g(x, y)$ as

$$g(x, y) = f(x, y) + \eta(x, y), \quad (2.91)$$


where $f(x, y)$ is the original image and $\eta(x, y)$ is the noise at (x, y) . Furthermore, it is common to assume that the noise process is statistically independent of (and hence, uncorrelated with) the image process. Then, we

**FIGURE 2.46**

(a) MR image of a knee. Results of averaging using three pixels in the (c) horizontal and (e) vertical directions, respectively. (b), (d), and (f): Log-magnitude spectra of the images in (a), (c), and (e), respectively. The spectra show values in the range $[3, 12]$ mapped to $[0, 255]$.

**FIGURE 2.47**

(a) Result of 3×3 averaging of the knee MR image in Figure 2.46 (a). (b) Log-magnitude spectrum of the image in (a).

**FIGURE 2.48**

MTFs of a DR system (II-TV = image-intensifier television) and a screen-film system at the same X-ray dose. FS = focal spot. Reproduced with permission from Y. Higashida, Y. Baba, M. Hatemura, A. Yoshida, T. Takada, and M. Takahashi, "Physical and clinical evaluation of a $2,048 \times 2,048$ -matrix image intensifier TV digital imaging system in bone radiography", *Academic Radiology*, 3(10):842–848. 1996. © Association of University Radiologists.




FIGURE 2.49

Contrast-detail curves of a DR system (II-TV = image-intensifier television) and a screen-film system. The DR system was operated at the same X-ray dose as the screen-film system (iso-dose) and at a low-dose setting. Reproduced with permission from Y. Higashida, Y. Baba, M. Hatemura, A. Yoshida, T. Takada, and M. Takahashi, "Physical and clinical evaluation of a $2,048 \times 2,048$ -matrix image intensifier TV digital imaging system in bone radiography", *Academic Radiology*, 3(10):842-848. 1996. © Association of University Radiologists.




FIGURE 2.50

MTF curves of an amorphous selenium (aSe) detector system for direct digital mammography, a screen-film system (S-F), and an indirect digital imaging system. $c/mm = \text{cycles}/mm$. Figure courtesy of J.E. Gray, Lorad, Danbury, CT.

have

$$\mu_g = \mu_f + \mu_\eta, \quad (2.92)$$

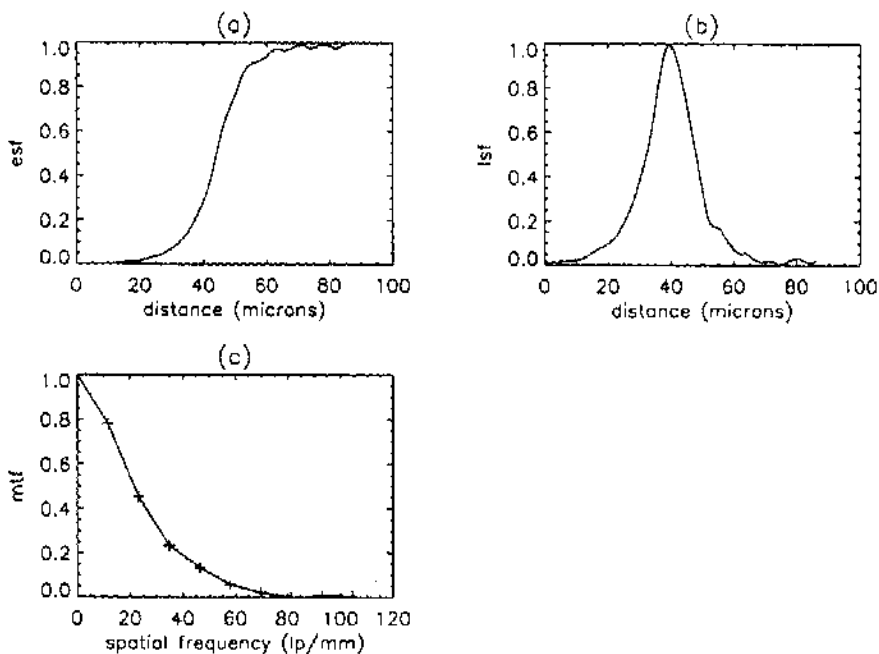
where μ represents the mean (average) of the process indicated by the subscript. In many cases, the mean of the noise process is zero. The variances (σ^2) of the processes are related as

$$\sigma_g^2 = \sigma_f^2 + \sigma_\eta^2. \quad (2.93)$$

SNR is a measure used to characterize objectively the relative strengths of the true image (signal) and the noise in an observed (noisy or degraded) image. Several definitions of SNR exist, and are used depending upon the information available and the feature of interest. A common definition of SNR is

$$\text{SNR}_1 = 10 \log_{10} \left[\frac{\sigma_f^2}{\sigma_\eta^2} \right] \text{ dB}. \quad (2.94)$$

The variance of noise may be estimated by computing the sample variance of pixels selected from areas of the given image that do not contain, or are not expected to contain, any image component. The variance of the true image as well as that of noise may be computed from the PDFs of the corresponding processes if they are known or if they can be estimated.

**FIGURE 2.51**

(a) Edge spread function, (b) line spread function, and (c) MTF of a μ CT system. 1 micron = 1 μ m. Reproduced with permission from M. Pateyron, F. Peyrin, A.M. Laval-Jeantet, P. Spanne, P. Cloetens, and G. Peix, "3D microtomography of cancellous bone samples using synchrotron radiation", *Proceedings of SPIE 2708: Medical Imaging 1996 - Physics of Medical Imaging*, Newport Beach, CA, pp 417-426. © SPIE.

In some applications, the variance of the image may not provide an appropriate indication of the useful range of variation present in the image. For this reason, another commonly used definition of SNR is based upon the dynamic range of the image, as

$$\text{SNR}_2 = 20 \log_{10} \left[\frac{f_{\max} - f_{\min}}{\sigma_{\eta}} \right] \text{ dB}. \quad (2.95)$$

Video signals in modern CRT monitors have SNR of the order of 60–70 dB with noninterlaced frame repetition rate in the range 70–80 frames per second.

Contrast-to-noise ratio (CNR) is a measure that combines the contrast or the visibility of an object and the SNR, and is defined as

$$\text{CNR} = \frac{\mu_f - \mu_b}{\sigma_b}, \quad (2.96)$$

where f is an ROI (assumed to be uniform, such as a disc being imaged using X rays), and b is a background region with no signal content (see Figure 2.7). Comparing this measure to the basic measure of simultaneous contrast in Equation 2.8, the difference lies in the denominator, where CNR uses the standard deviation. Whereas simultaneous contrast uses a background region that encircles the ROI, CNR could use a background region located elsewhere in the image. CNR is well suited to the analysis of X-ray imaging systems, where the density of an ROI on a film image depends upon the dose: the visibility of an object is dependent upon both the dose and the noise.

In a series of studies on image quality, Schade reported on image gradation, graininess, and sharpness in television and motion-picture systems [135]; on an optical and photoelectric analog of the eye [136]; and on the evaluation of photographic image quality and resolving power [137]. The sine-wave, edge-transition, and square-wave responses of imaging systems were discussed in detail. Schade presented a detailed analysis of the relationships between resolving power, contrast sensitivity, number of perceptible gray-scale steps, and granularity with the “three basic characteristics” of an imaging system: intensity transfer function, sine-wave response, and SNR. Schade also presented experimental setups and procedures with optical benches and equipment for photoelectric measurements and characterization of optical and imaging systems.

Burke and Snyder [138] reported on quality metrics of digital images as related to interpreter performance. Their test set included a collection of 250 transparencies of 10 digital images, each degraded by five levels of blurring and five levels of noise. Their work addressed the question “How can we measure the degree to which images are improved by digital processing?” The results obtained indicated that although the main effect of blur was not significant in their interpretation experiment (in terms of the extraction of the “essential elements of information”), the effect of noise was significant. However, in medical imaging applications such as SPECT, high levels of noise

are tolerated, but blurring of edges caused by filters used to suppress noise is not accepted.

Tapiovaara and Wagner [139] proposed a method to measure image quality in the context of the image information available for the performance of a specified detection or discrimination task by an observer. The method was applied to the analysis of fluoroscopy systems by Tapiovaara [140].

2.14 Error-based Measures

Notwithstanding several preceding works on image quality, Hall [141] stated (in 1981) that “A major problem which has plagued image processing has been the lack of an effective image quality measure.” In his paper on subjective evaluation of a perceptual quality metric, Hall discussed several image quality measures including the MSE, normalized MSE (NMSE), normalized error (NE), and Laplacian MSE (LMSE), and then defined a “perceptual MSE” or PMSE based on an HVS model. The measures are based upon the differences between a given test image $f(m, n)$ and its degraded version $g(m, n)$ after passage through the imaging system being evaluated, computed over the full image frame either directly or after some filter or transform operation, as follows:

$$\text{MSE} = \frac{1}{MN} \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} [f(m, n) - g(m, n)]^2; \quad (2.97)$$

$$\text{NMSE} = \frac{\sum_{m=0}^{M-1} \sum_{n=0}^{N-1} [f(m, n) - g(m, n)]^2}{\sum_{m=0}^{M-1} \sum_{n=0}^{N-1} [f(m, n)]^2}; \quad (2.98)$$

$$\text{NE} = \frac{\sum_{m=0}^{M-1} \sum_{n=0}^{N-1} |f(m, n) - g(m, n)|}{\sum_{m=0}^{M-1} \sum_{n=0}^{N-1} |f(m, n)|}; \quad (2.99)$$

$$\text{LMSE} = \frac{\sum_{m=1}^{M-2} \sum_{n=1}^{N-2} [f_L(m, n) - g_L(m, n)]^2}{\sum_{m=1}^{M-2} \sum_{n=1}^{N-2} [f_L(m, n)]^2}; \quad (2.100)$$

where the $M \times N$ images are defined over the range $m = 0, 1, 2, \dots, M - 1$, and $n = 0, 1, 2, \dots, N - 1$, and $f_L(m, n)$ is the Laplacian (second derivative) of $f(m, n)$ defined as in Equation 2.82 for $m = 1, 2, \dots, M - 2$, and $n = 1, 2, \dots, N - 2$. PMSE was defined in a manner similar to NMSE, but with each image replaced with the logarithm of the image convolved with a PSF representing the HVS. Hall’s results showed that PMSE correlated well with subjective ranking of images to greater than 99.9%, and performed better than NMSE or LMSE. It should be noted that the measures defined above assume the availability of a reference image for comparison in a *before-and-after* manner.

2.15 Application: Image Sharpness and Acutance

In the search for a single measure that could represent the combined effects of various imaging and display processes, several researchers proposed various measures under the general label of “acutance” [114, 115, 142, 143, 144, 145, 146]. The following paragraphs present a review of several such measures based upon the ESF and the MTF.

As an aside, it is important to note the distinction between acutance and acuity. Westheimer [147, 148] discussed the concepts of visual acuity and hyperacuity and their light-spread and frequency-domain descriptions. Acuity (as evaluated with Snellen letters or Landolt “C”s) tests the “minimum separable”, where the visual angle of a small feature is varied until a discrimination goal just can or cannot be achieved (a resolution task). On the other hand, hyperacuity (vernier or stereoscopic acuity) relates to spatial localization or discrimination.

The edge spread function: Higgins and Jones [115] discussed the nature and evaluation of the sharpness of photographic images, with particular attention to the importance of gradients. With the observation that the cones in the HVS, while operating in the mode of photopic vision (under high-intensity lighting), respond to temporal illuminance gradients, and that the eye moves to scan the field of vision, they argued that spatial luminance gradients in the visual field represent physical aspects of the object or scene that affect the perception of detail. Higgins and Jones conducted experiments with microdensitometric traces of knife edges recorded on various photographic materials, and found that the maximum gradient or average gradient measures along the knife-edge spread functions (KESF) failed to correlate with sharpness as judged by human observers.

Figure 2.20 illustrates an ideal sharp edge and a hypothetical KESF. Higgins and Jones proposed a measure of acutance based upon the mean-squared gradient across a KESF as

$$A = \frac{1}{f(b) - f(a)} \int_a^b \left[\frac{d}{dx} f(x) \right]^2 dx, \quad (2.101)$$

where $f(x)$ represents the intensity function along the edge, and a and b are the spatial limits of the (blurred) edge. Ten different photographic materials were evaluated with the measure of acutance, and the results indicated excellent correlation between acutance and subjective judgment of sharpness.

Wolfe and Eisen [149] reported on psychometric evaluation of the sharpness of photographic reproductions. They stated that resolving power, maximum gradient, and average gradient do not correlate well with sharpness, and that the variation of density across an edge is an obvious physical measurement to be investigated in order to obtain an objective correlate of sharpness. Perin [114] continued along these lines, and proposed an averaged measure of

acutance by averaging the mean-squared gradient measure of Higgins and Jones over many sections of the KESF, and further normalizing it with respect to the density difference across the knife edge. Perrin also discussed the relationship between the edge trace and the LSF.

MTF-based measures: Although Perrin [114] reported on an averaged acutance measure based on the mean-squared gradient measure of Higgins and Jones, he also remarked [150] that the sine-wave response better describes the behavior of an optical system than a single parameter (such as resolving power), and discussed the relationship between the sine-wave response and spread functions. The works of Schade and Perrin, perhaps, shifted interest from the spatial-gradient technique of Higgins and Jones to the frequency domain.

Frequency-domain measures related to image quality typically combine the areas under the MTF curves of the long chain of systems and processes involved from the initial stage of the camera lens, through the film and/or display device, to the final visual system of the viewer [146]. It is known from basic linear system theory that when a composite system includes a number of LSI systems with transfer functions $H_1(u, v)$, $H_2(u, v)$, \dots , $H_N(u, v)$ in series (cascade), the transfer function of the complete system is given by

$$H(u, v) = H_1(u, v) H_2(u, v) \cdots H_N(u, v) = \prod_{i=1}^N H_i(u, v). \quad (2.102)$$

Equivalently, we have the PSF of the net system given by

$$h(x, y) = h_1(x, y) * h_2(x, y) * \cdots * h_N(x, y). \quad (2.103)$$

Given that the high-frequency components in the spectrum of an image are associated with sharp edges in the image domain, it may be observed that the transfer function of an imaging or image processing system should possess large gains at high frequencies in order for the output image to retain the sharpness present in the input. This observation leads to the result that, over a given frequency range of interest, a system with larger gains at higher frequencies (and hence a sharper output image) will have a larger area under the normalized MTF than another system with lower gains (and hence poorer sharpness in the resulting image). By design, MTF-area-based measures represent the combined effect of all the systems between the image source and the viewer; they are independent of the actual image displayed.

Crane [142] started a series of definitions of acutance based on the MTFs of imaging system components. He discussed the need for objective correlates of the subjective property of image sharpness or crispness, and remarked that resolving power is misleading, and that the averaged squared gradient of edge profiles is dependable but cannot include the effects of all the components in a photographic system (camera to viewer). Crane proposed a single numerical rating based on the areas under the MTF curves of all the systems in the

chain from the camera to the viewer (for example, camera, negative, printer, intermediate processing systems, print film, projector, screen, and observer). He called the measure the *system modulation transfer acutance* (SMTA) and claimed that it could be readily comprehended, compared, and tabulated. He also recommended that the acutance measure proposed by Higgins and Jones [115] and Perrin [114] be called *image edge-profile acutance* (IEPA). Crane evaluated SMTA using 30 color films and motion-picture films, and found it to be a good tool.

Crane's work started another series of papers proposing modified definitions of MTF-based acutance measures for various applications: Gendron [146] proposed a "cascaded modulation transfer or CMT" measure of acutance (CMTA) to rectify certain deficiencies in SMTA. Crane wrote another paper on acutance and granulance [143] and defined "AMT acutance" (AMTA) based on the ratio of the MTF area of the complete imaging system including the human eye to that of the eye alone. He also presented measures of granulance based on root mean-squared (RMS) deviation from mean lightness in areas expected to be uniform, and discussed the relationships between acutance and granulance. CMTA was used by Kriss [145] to compare the system sharpness of continuous and discrete imaging systems. AMTA was used by Yip [144] to analyze the imaging characteristics of CRT multifformat printers.

Assuming the systems involved to be isotropic, the MTF is typically expressed as a 1D function of the radial unit of frequency $\nu = \sqrt{(u^2 + v^2)}$; see Figures 2.33 (b) and 2.51 (c). Let us represent the combined MTF of the complete chain of systems as $H_s(\nu)$. Some of the MTF-area-based measures are defined as follows:

$$A_1 = \int_0^{\nu_{\max}} [H_s(\nu) - H_e(\nu)] d\nu, \quad (2.104)$$

where $H_e(\nu)$ is the MTF threshold of the eye, ν represents the radial frequency at the eye of the observer, and ν_{\max} is given by the condition $H_s(\nu_{\max}) = H_e(\nu_{\max})$ [151]. In order to reduce the weighting on high-frequency components, another measure replaces the difference between the MTFs as above with their ratio, as [151]

$$A_2 = \int_0^{\infty} \frac{H_s(\nu)}{H_e(\nu)} d\nu. \quad (2.105)$$

AMTA was defined as [143, 144]

$$\text{AMTA} = 100 + 66 \log_{10} \left[\frac{\int_0^{\infty} H_s(\nu) H_e(\nu) d\nu}{\int_0^{\infty} H_e(\nu) d\nu} \right]. \quad (2.106)$$

The MTF of the eye was modeled as a Gaussian with standard deviation $\sigma = 13$ *cycles/degree*. AMTA values were interpreted as 100 : excellent, 90 : good, 80 : fair, and 70 : just passable [143].

Several authors have presented and discussed various other image quality criteria and measures that are worth mentioning here; whereas some are based on the MTF and hence have some common ground with acutance, others are based on different factors. Higgins [152] discussed various methods for analyzing photographic systems, including the effects of nonlinearity, LSFs, MTFs, granularity, and sharpness. Granger and Cupery [153] proposed a “subjective quality factor (SQF)” based upon the integral of the system MTF (including scaling effects to the retina) over a certain frequency range. Their results indicated a correlation of 0.988 between SQF and subjective ranking by observers.

Higgins [154] published a detailed review of various image quality criteria. Quality criteria as related to objective or subjective tone reproduction, sharpness, and graininess were described. Higgins reported on the results of tests evaluating various versions of MTF-based acutance and other measures with photographic materials having widely different MTFs, and recommended that MTF-based acutance measures are good when no graininess is present; SNR-based measures were found to be better when graininess was apparent. Task et al. [155] compared several television (TV) display image quality measures. Their tests included target recognition tasks and several FOMs such as limiting resolution, MTF area, threshold resolution, and gray-shade frequency product. They found MTF area to be the best measure among those evaluated.

Barten [151, 156] presented reviews of various image quality measures, and proposed the evaluation of image quality using the square-root integral (SQRI) method. The SQRI measure is based upon the ratio of the MTF of the display system to that of the eye, and can take into account the contrast sensitivity of the eye and various display parameters such as resolution, addressability, contrast, luminance, display size, and viewing distance. SQRI is defined as

$$\text{SQRI} = \frac{1}{\ln(2)} \int_0^{\nu_{\max}} \left[\frac{H_s(\nu)}{H_e(\nu)} \right]^{\frac{1}{2}} \frac{d\nu}{\nu}. \quad (2.107)$$

Here, ν_{\max} is the maximum frequency to be displayed. The SQRI measure overcomes some limitations in the SQF measure of Granger and Cupery [153]. Based upon good correlation between SQRI and perceived subjective image quality, Barten proposed SQRI as an “excellent universal measure of perceived image quality”.

Carlson and Cohen [157] proposed a psychophysical model for predicting the visibility of displayed information, combining the effects of MTF, noise, sampling, scene content, mean luminance, and display size. They noted that edge transitions are a significant feature of most scenes, and proposed “discriminable difference diagrams” of modulation transfer versus retinal frequency (in cycles per degree). Their work indicated that discriminable difference diagrams could be used to predict the visibility of MTF changes in magnitude but not in phase.

Several other measures of image quality based upon the HVS have been proposed by Saghri et al. [158], Nill and Bouzas [159], Lukas and Budrikis [160], and Budrikis [161].

Region-based measure of edge sharpness: Westerink and Roufs [162] proposed a local basis for perceptually relevant resolution measures. Their experiments included the presentation of a number of slides with complex scenes at variable resolution created by defocusing the lens of the projector, and at various widths. They showed that the width of the LSF correlates well with subjective quality, and remarked that MTF-based measures “do not reflect the fact that local aspects such as edges and contours play an important role in the quality sensation”.

Rangayyan and Elkadiki [116] discussed the importance of a local measure of quality, sharpness, or perceptibility of a region or feature of interest in a given image. The question asked was “Given two images of the same scene, which one permits better perception of a specific region or object in the image?” Such a situation may arise in medical imaging, where one may have an array of images of the same patient or phantom test object acquired using multiple imaging systems (different models or various imaging parameter settings on the same system). It would be of interest to determine which system or set of parameters provides the image where a specific object, such as a tumor, may be seen best. Whereas local luminance gradients are indeed reflected as changes at all frequencies in the MTF, such a global characteristic may dilute the desired difference in the situation mentioned above. Furthermore, MTF-based measures characterize the imaging and viewing systems in general, and are independent of the specific object or scene on hand.

Based upon the observations of Higgins and Jones [115], Wolfe and Eisen [149], Perrin [114], Carlson and Cohen [157], and Westerink and Roufs [162] on the importance of local luminance variations, gradients, contours, and edges (as reviewed above), Rangayyan and Elkadiki presented arguments in favor of a region-based measure of sharpness or acutance. They extended the measure of acutance defined by Higgins and Jones [115] and Perrin [114] to 2D regions by computing the mean-squared gradient across and around the contour of an object or ROI in the given image, and called the quantity “a region-based measure of image edge-profile or IEP acutance (IEPA)”. Figure 2.52 illustrates the basic principle involved in computing the gradient around an ROI, using normals (perpendiculars to the tangents) at every pixel on its boundary. Instead of the traditional difference defined as

$$f'(n) = f(n) - f(n - 1), \quad (2.108)$$

Rangayyan and Elkadiki (see Rangayyan et al. [163] for revised definitions) split the normal at each boundary pixel into a foreground part $f(n)$ and a background part $b(n)$ (see Figure 2.52), and defined an averaged gradient as

$$f_d(k) = \frac{1}{N} \sum_{n=1}^N \frac{f(n) - b(n)}{2n}, \quad (2.109)$$

where k is the index of the boundary pixel and N is the number of pairs of pixels (or differences) used along the normal. The averaged gradient values over all boundary pixels were then combined to obtain a single normalized value of acutance A for the entire region as

$$A = \frac{1}{d_{\max}} \left[\frac{1}{K} \sum_{k=1}^K f_d^2(k) \right]^{\frac{1}{2}}, \quad (2.110)$$

where K is the number of pixels along the boundary, and d_{\max} is the maximum possible gradient value used as a normalization factor. It was shown that the value of A was reduced by blurring and increased by sharpening of the ROI [116, 164]. Olabarriaga and Rangayyan [117] further showed that acutance as defined in Equation 2.110 is not affected significantly by noise, and that it correlates well with sharpness as judged by human observers.

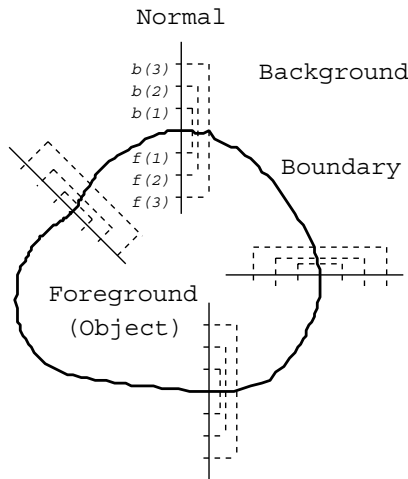


FIGURE 2.52

Computation of differences along the normals to a region in order to derive a measure of acutance. Four sample normals are illustrated, with three pairs of pixels being used to compute differences along each normal.

Example of application: In terms of their appearance on mammograms, most benign masses of the breast are well-circumscribed with sharp boundaries that delineate them from surrounding tissues. On the other hand, most malignant tumors possess fuzzy boundaries with slow and extended transition from a dense core region to the surrounding, less-dense tissues. Based upon this radiographic observation, Rangayyan et al. [163] hypothesized that the acutance measure should have higher values for benign masses than for ma-

lignant tumors. Acutance was computed using normals with variable length adapted to the complexity of the shape of the boundary of the mass being analyzed. The measure was tested using 39 mammograms, including 28 benign masses and 11 malignant tumors. Boundaries of the masses were drawn by a radiologist for this study. It was found that acutance could lead to the correct classification of all of the 11 malignant tumors and 26 out of the 28 benign masses, resulting in an overall accuracy of 94.9%.

Mudigonda et al. [165, 166] evaluated several versions of the acutance measure by defining the differences based upon successive pixel pairs and across-the-boundary pixel pairs. It was observed that the acutance measure is sensitive to the location of the reference boundary [163, 164, 165, 166]. See Sections 7.9.2 and 12.12 as well as Figure 12.4 for more details and illustrations related to acutance.

2.16 Remarks

We have reviewed several notions of image quality and information content in this chapter. We explored many methods and measures designed to characterize various image attributes associated with quality and information content. It should be observed that image quality considerations vary from one application of imaging to another, and that appropriate measures should be chosen after due assessment of the particular problem on hand. In medical diagnostic applications, emphasis is usually placed on the assessment of image quality in terms of its effect on the accuracy of diagnostic interpretation by human observers and specialists: methods related to this approach are discussed in Sections 4.11, 12.8, and 12.10. The use of measures of information content in the analysis of methods for image coding and data compression is described in Chapter 11.

2.17 Study Questions and Problems

(*Note:* Some of the questions may require background preparation with other sources on the basics of signals and systems as well as digital signal and image processing, such as Lathi [1], Oppenheim et al. [2], Oppenheim and Schaffer [7], Gonzalez and Woods [8], Pratt [10], Jain [12], Hall [9], and Rosenfeld and Kak [11].)

Selected data files related to some of the problems and exercises are available at the site

1. Explain the differences between spatial resolution and gray-scale resolution in a digitized image.
2. Give the typical units for the variables (x, y) and (u, v) used in the representation of images in the space and frequency domains.
3. How can a continuous (analog) image be recovered from its sampled (digitized) version? Describe the operations required in
 - (a) the space domain, and
 - (b) the frequency domain.
 What are the conditions to be met for exact recovery of the analog image?
4. Distinguish between gray-scale dynamic range and simultaneous contrast. Explain the effects of the former on the latter.
5. Draw schematic sketches of the histograms of the following types of images:
 - (a) A collection of objects of the same uniform gray level placed on a uniform background of a different gray level.
 - (b) A collection of relatively dark cells against a relatively bright background, with both having some intrinsic variability of gray levels.
 - (c) An under-exposed X-ray image.
 - (d) An over-exposed X-ray image.

Annotate the histograms with labels and comments.

6. Starting with the expression for the entropy of a continuous PDF, show that the entropy is maximized by a uniform PDF. (*Hint:* Treat this as a constrained optimization problem, with the constraint being that the integral of the PDF be equal to unity.)
7. Define two rectangular functions as

$$f_1(x, y) = \begin{cases} 1 & \text{if } 0 \leq x \leq X; 0 \leq y \leq Y \\ 0 & \text{otherwise,} \end{cases} \quad (2.111)$$

and

$$f_2(x, y) = \begin{cases} 1 & \text{if } |x| \leq \frac{X}{2}, |y| \leq \frac{Y}{2} \\ 0 & \text{otherwise.} \end{cases} \quad (2.112)$$

Starting from the definition of the 2D Fourier transform, derive the Fourier transforms $F_1(u, v)$ and $F_2(u, v)$ of the two functions; show all steps.

Explain the differences between the two functions in the spatial and frequency domains.

8. Using the continuous 2D convolution and Fourier transform expressions, prove that convolution in the space domain is equivalent to multiplication of the corresponding functions in the Fourier domain.
9. You are given three images of rectangular objects as follows:
 - (a) a horizontally placed rectangle with the horizontal side two times the vertical side;
 - (b) the rectangle in (a) rotated by 45° ; and
 - (c) the rectangle in (a) reduced in each dimension by a factor of two.

Draw schematic diagrams of the Fourier spectra of the three images. Explain the differences between the spectra.

10. Draw schematic diagrams of the Fourier magnitude spectra of images with
- a circle of radius R ;
 - a circle of radius $2R$; and
 - a circle of radius $R/2$.

The value of R is not relevant. Explain the differences between the three cases in both the space domain and the frequency domain.

11. (a) Derive the expression for the Fourier transform of $\frac{\partial f(x,y)}{\partial x}$ in terms of the Fourier transform of $f(x,y)$. Show and explain all steps. (Hint: Start with the definition of the inverse Fourier transform.)

Explain the effect of the differentiation operator in the space domain and the frequency domain.

(b) Based upon the result in (a), what is the Fourier transform of $\frac{\partial^2 f(x,y)}{\partial x^2}$? Explain.

(c) Based upon the result in (a), state the relationship between the Fourier transform of $[\frac{\partial f(x,y)}{\partial x}]^2$ and that of $f(x,y)$. State all properties that you use.

(d) Explain the differences between the operators in (a), (b), and (c) and their effects in both the space domain and the frequency domain.

12. Using the continuous 2D Fourier transform expression, prove that the inverse Fourier transform of a function $F(u,v)$ may be obtained by taking the forward Fourier transform of the complex conjugate of the given function [that is, taking the forward transform of $F^*(u,v)$], and then taking the complex conjugate of the result.

13. Starting with the 2D DFT expression, show how the 2D DFT may be computed as a series of 1D DFTs. Show and explain all steps.

14. An image of size $100\text{ mm} \times 100\text{ mm}$ is digitized into a matrix of size 200×200 pixels with uniform sampling and equal spacing between the samples in the horizontal and vertical directions. The spectrum of the image is computed using the FFT algorithm after padding the image to 256×256 pixels.

Draw a square to represent the array containing the spectrum and indicate the FFT array indices as well as the frequency coordinates in mm^{-1} at the four corners, at the mid-point of each side, and at the center of the square.

15. A system performs the operation $g(x,y) = f(x,y) - f(x-1,y)$. Derive the MTF of the system and explain its characteristics.

16. Using the continuous Fourier transform, derive the relationship between the Fourier transforms of an image $f(x,y)$ and its modified version given as $f_1(x,y) = f(x-x_1, y-y_1)$.

Explain the differences between the two images in the spatial and frequency domains.

17. The impulse response of a system is approximated by the 3×3 matrix

$$\begin{bmatrix} 1 & 2 & 1 \\ 2 & 3 & 2 \\ 1 & 2 & 1 \end{bmatrix}. \quad (2.113)$$

Derive the transfer function of the system and explain its characteristics.

18. The image in Equation 2.113 is processed by systems having the following impulse responses:

(a) $h(m, n) = [-1, 1]$;

(b) $h(m, n) = [-1, 1]^T$; and

(c) $h(m, n) =$ a 3×3 matrix with all elements equal to $\frac{1}{9}$.

Compute the output image in each case over a 3×3 array, assuming that the input is zero outside the array given in Equation 2.113.

19. The 5×5 image

$$f(m, n) = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 10 & 10 & 10 & 0 \\ 0 & 10 & 10 & 10 & 0 \\ 0 & 10 & 10 & 10 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (2.114)$$

is processed by two systems in cascade. The first system produces the output $g_1(m, n) = f(m, n) - f(m - 1, n)$. The second system produces the output $g_2(m, n) = g_1(m, n) - g_1(m, n - 1)$.

Compute the images g_1 and g_2 .

Does the sequence of application of the two operators affect the result? Why (not)?

Explain the effects of the two operators.

20. Derive the MTF for the Laplacian operator and explain its characteristics.

21. Write the expressions for the convolution and correlation of two images.

Explain the similarities and differences between the two.

What are the equivalent relationships in the Fourier domain?

Explain the effects of the operations in the spatial and frequency domains.

22. Consider two systems with the impulse responses

(a) $h_1(m, n) = [-1, 1]$; and

(b) $h_2(m, n) = [-1, 1]^T$.

What will be the effect of passing an image through the two systems in

(a) parallel (and adding the results of the individual systems), or

(b) series (cascade)?

Considering a test image made up of a bright square in the middle of a dark background, draw schematic diagrams of the outputs at each stage of the two systems mentioned above.

23. The squared gradient of an image $f(x, y)$ is defined as

$$g(x, y) = \left[\frac{\partial f(x, y)}{\partial x} \right]^2 + \left[\frac{\partial f(x, y)}{\partial y} \right]^2. \quad (2.115)$$

Derive the expression for $G(u, v)$, the Fourier transform of $g(x, y)$.

How does this operator differ from the Laplacian in the spatial and frequency domains?

24. Using mathematical expressions and operations as required, explain how a degraded image of an edge may be used to derive the MTF of an imaging system.

2.18 Laboratory Exercises and Projects

1. Prepare a phantom for X-ray imaging by attaching a few strips of metal (such as aluminum or copper) of various thickness to a plastic or plexiglass sheet. Ensure that the strips have straight edges. With the help of a qualified technologist, obtain X-ray images of the phantom at a few different kVp and mAs settings. Note the imaging parameters for each experiment.

Repeat the experiment with screens and films of different characteristics, with and without the grid (bucky), and with the grid being stationary. Study the contrast, noise, artifacts, and detail visibility in the resulting images.

Digitize the images for use in image processing experiments.

Scan across the edges of the metal strips and obtain the ESF. From this function, derive the LSF, PSF, and MTF of the imaging system for various conditions of imaging.

Measure the SNR and CNR of the various metal strips and study their dependence upon the imaging parameters.

2. Repeat the experiment above with wire meshes of different spacing in the range 1 – 10 lines per mm . Study the effects of the X-ray imaging and digitization parameters on the clarity and visibility of the mesh patterns.
3. Compute the Fourier spectra of several biomedical images with various objects and features of different size, shape, and orientation characteristics, as well as of varying quality in terms of noise and sharpness. Calibrate the spectra in terms of frequency in mm^{-1} or lp/mm . Explain the relationships between the spatial and frequency-domain characteristics of the images and their spectra.
4. Compute the histograms and log-magnitude Fourier spectra of at least ten test images that you have acquired.

Comment on the nature of the histograms and spectra.

Relate specific image features to specific components in the histograms and spectra.

Comment on the usefulness of the histograms and spectra in understanding the information content of images.

5. Create a test image of size 100×100 pixels, with a circle of diameter 30 pixels at its center. Let the value of the pixels inside the circle be 100, and those outside be 80.

Prepare three blurred versions of the test image by applying the 3×3 mean filter

- (i) once,
- (ii) three times, and
- (iii) five times successively.

To each of the three blurred images obtained as above, add three levels of

- (a) Gaussian noise, and
- (b) speckle noise.

Select the noise levels such that the edge of the circle becomes obscured in at least some of the images.

Compute the error measures MSE and NMSE between the original test image and each of the 18 degraded images obtained as above.

Study the effect of blurring and noise on the error measures and explain your findings.

6. From your collection of test images, select two images: one with strong edges of the objects or features present in the image, and the other with weaker definition of edges and features.

Compute the horizontal difference, vertical difference, and the Laplacian of the images. Find the minimum and maximum values in each result, and map appropriate ranges to the display range in order to visualize the results.

Study the results obtained and comment upon your findings in relation to the details present in the test images.

Removal of Artifacts

Noise is omnipresent! Biomedical images are often affected and corrupted by various types of noise and artifact. Any image, pattern, or signal other than that of interest could be termed as interference, artifact, or simply *noise*. The sources of noise could be physiological, the instrumentation used, or the environment of the experiment. The problems caused by artifacts in biomedical images are vast in scope and variety; their potential for degrading the performance of the most sophisticated image processing algorithms is high. The removal of artifacts without causing any distortion or loss of the desired information in the image of interest is often a significant challenge. The enormity of the problem of noise removal and its importance are reflected by the placement of this chapter as the first chapter on image processing techniques in this book.

This chapter starts with an introduction to the nature of the artifacts that are commonly encountered in biomedical images. Several illustrations of images corrupted by various types of artifacts are provided. Details of the design of filters spanning a broad range of approaches, from linear space-domain and frequency-domain fixed filters, to the optimal Wiener filter, and further on to nonlinear and adaptive filters, are then described. The chapter concludes with demonstrations of application of the filters described to a few biomedical images.

(*Note:* A good background in signal and system analysis [1, 2, 3, 167] as well as probability, random variables, and stochastic processes [3, 128, 168, 169, 170, 171, 172, 173] is required in order to follow the procedures and analyses described in this chapter.)

3.1 Characterization of Artifacts

3.1.1 Random noise

The term *random noise* refers to an interference that arises from a random process such as thermal noise in electronic devices and the counting of photons. A random process is characterized by the PDF representing the probabilities of occurrence of all possible values of a random variable. (See Papoulis [128]

or Bendat and Piersol [168] for background material on probability, random variables, and stochastic processes.)

Consider a random process η that is characterized by the PDF $p_\eta(\eta)$. The process could be a function of time as $\eta(t)$, or of space in 1D, 2D, or 3D as $\eta(x)$, $\eta(x, y)$, or $\eta(x, y, z)$; it could also be a spatio-temporal function as $\eta(x, y, z, t)$. The argument of the PDF represents the value that the random process can assume, which could be a voltage in the case of a function of time, or a gray level in the case of a 2D or 3D image. The use of the same symbol for the function and the value it can assume when dealing with PDFs is useful when dealing with several random processes.

The mean μ_η of the random process η is given by the first-order moment of the PDF, defined as

$$\mu_\eta = E[\eta] = \int_{-\infty}^{\infty} \eta p_\eta(\eta) d\eta, \quad (3.1)$$

where $E[\]$ represents the *statistical expectation operator*. It is common to assume the mean of a random noise process to be zero.

The mean-squared (MS) value of the random process η is given by the second-order moment of the PDF, defined as

$$E[\eta^2] = \int_{-\infty}^{\infty} \eta^2 p_\eta(\eta) d\eta. \quad (3.2)$$

The variance σ_η^2 of the process is defined as the second central moment:

$$\sigma_\eta^2 = E[(\eta - \mu_\eta)^2] = \int_{-\infty}^{\infty} (\eta - \mu_\eta)^2 p_\eta(\eta) d\eta. \quad (3.3)$$

The square root of the variance gives the standard deviation (SD) σ_η of the process. Note that $\sigma_\eta^2 = E[\eta^2] - \mu_\eta^2$. If the mean is zero, it follows that $\sigma_\eta^2 = E[\eta^2]$, that is, the variance and the MS values are the same.

Observe the use of the same symbol η to represent the random variable, the random process, and the random signal as a function of time or space. The subscript of the PDF or the statistical parameter derived indicates the random process of concern. The context of the discussion or expression should make the meaning of the symbol clear.

When the values of a random process η form a time series or a function of time, we have a random signal (or a stochastic process) $\eta(t)$; see Figure 3.1. When one such time series is observed, it is important to note that the entity represents but one single realization of the random process. An example of a random function of time is the current generated by a CCD detector element due to thermal noise when no light is falling on the detector (known as the *dark current*). The statistical measures described above then have physical meaning: the mean represents the DC component, the MS value represents the average power, and the square root of the mean-squared value (the root mean-squared or RMS value) gives the average noise magnitude. These measures

are useful in calculating the SNR, which is commonly defined as the ratio of the peak-to-peak amplitude range of the signal to the RMS value of the noise, or as the ratio of the average power of the desired signal to that of the noise. Special-purpose CCD detectors are cooled by circulating cold air, water, or liquid nitrogen to reduce thermal noise and improve the SNR.





FIGURE 3.1


A time series composed of random noise samples with a Gaussian PDF having $\mu = 0$ and $\sigma^2 = 0.01$. MS value = 0.01; RMS = 0.1. See also Figures 3.2 and 3.3.

When the values of a random process η form a 2D function of space, we have a noise image $\eta(x, y)$; see Figures 3.2 and 3.3. Several possibilities arise in this situation: We may have a single random process that generates random gray levels that are then placed at various locations in the (x, y) plane in some structured or random sequence. We may have an array of detectors with one detector per pixel of a digital image; the gray level generated by each detector may then be viewed as a distinct random process that is independent of those of the other detectors. A TV image generated by such a camera in the presence of no input image could be considered to be a noise process in (x, y, t) , that is, a function of space and time.

A biomedical image of interest $f(x, y)$ may also, for the sake of generality, be considered to be a realization of a random process f . Such a representation

**FIGURE 3.2**

An image composed of random noise samples with a Gaussian PDF having $\mu = 0$ and $\sigma^2 = 0.01$. MS value = 0.01; RMS = 0.1. The normalized pixel values in the range $[-0.5, 0.5]$ were linearly mapped to the display range $[0, 255]$. See also Figure 3.3.

**FIGURE 3.3**

Normalized histogram of the image in Figure 3.2. The samples were generated using a Gaussian process with $\mu = 0$ and $\sigma^2 = 0.01$. MS value = 0.01; RMS = 0.1. See also Figures 3.1 and 3.2.

allows for the statistical characterization of sample-to-sample or person-to-person variations in a collection of images of the same organ, system, or type. For example, although almost all CT images of the brain show the familiar cerebral structure, variations do exist from one person to another. A brain CT image may be represented as a random process that exhibits certain characteristics *on the average*. Statistical averages representing populations of images of a certain type are useful in designing filters, data compression techniques, and pattern classification procedures that are optimal for the specific type of images. However, it should be borne in mind that, in diagnostic applications, it is the *deviation from the normal* or the average that is present in the image on hand that is of critical importance.

When an image $f(x, y)$ is observed in the presence of random noise η , the detected image $g(x, y)$ may be treated as a realization of another random process g . In most cases, the noise is additive, and the observed image is expressed as

$$g(x, y) = f(x, y) + \eta(x, y). \quad (3.4)$$

Each of the random processes f , η , and g is characterized by its own PDF $p_f(f)$, $p_\eta(\eta)$, and $p_g(g)$, respectively.

In most practical applications, the random processes representing an image of interest and the noise affecting the image may be assumed to be *statistically independent processes*. Two random processes f and η are said to be statistically independent if their joint PDF $p_{f,\eta}(f, \eta)$ is equal to the product of their individual PDFs given as $p_f(f) p_\eta(\eta)$. It then follows that the first-order moment and second-order central moment of the processes in Equation 3.4 are related as

$$E[g] = \mu_g = \mu_f + \mu_\eta = \mu_f = E[f], \quad (3.5)$$

$$E[(g - \mu_g)^2] = \sigma_g^2 = \sigma_f^2 + \sigma_\eta^2, \quad (3.6)$$

where μ represents the mean and σ^2 represents the variance of the random process indicated by the subscript, and it is assumed that $\mu_\eta = 0$.

Ensemble averages: When the PDFs of the random processes of concern are not known, it is common to approximate the statistical expectation operation by averages computed using a collection or *ensemble* of sample observations of the random process. Such averages are known as *ensemble averages*. Suppose we have M observations of the random process f as functions of (x, y) : $f_1(x, y), f_2(x, y), \dots, f_M(x, y)$; see Figure 3.4. We may estimate the mean of the process at a particular spatial location (x_1, y_1) as

$$\mu_f(x_1, y_1) = \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{k=1}^M f_k(x_1, y_1). \quad (3.7)$$

The autocorrelation function (ACF) $\phi_f(x_1, x_1 + \alpha, y_1, y_1 + \beta)$ of the random process f is defined as

$$\phi_f(x_1, x_1 + \alpha, y_1, y_1 + \beta) = E[f(x_1, y_1), f(x_1 + \alpha, y_1 + \beta)], \quad (3.8)$$

which may be estimated as

$$\phi_f(x_1, x_1+\alpha, y_1, y_1+\beta, k) = \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{k=1}^M f_k(x_1, y_1) f_k(x_1+\alpha, y_1+\beta), \quad (3.9)$$

where α and β are spatial shift parameters. If the image $f(x, y)$ is complex, one of the versions of $f(x, y)$ in the products above should be conjugated; most biomedical images that are encountered in practice are real-valued functions, and this distinction is often ignored. The ACF indicates how the values of an image at a particular spatial location are statistically related to (or have characteristics in common with) the values of the same image at another shifted location. If the process is stationary, the ACF depends only upon the shift parameters, and may be expressed as $\phi_f(\alpha, \beta)$.




FIGURE 3.4
Ensemble and spatial averaging of images.

The three equations above may be applied to signals that are functions of time by replacing the spatial variables (x, y) with the temporal variable t , replacing the shift parameter α with τ to represent temporal delay, and making a few other related changes.

When $\mu_f(x_1, y_1)$ is computed for every spatial location or pixel, we get an average image that could be expressed as $\bar{f}(x, y)$. The image \bar{f} may be used to represent the random process f as a prototype. For practical use, such an average should be computed using sample observations that are of the same size, scale, orientation, etc. Similarly, the ACF may also be computed for all possible values of its indices to obtain an image.

Temporal and spatial averages: When we have a sample observation of a random process $f_k(t)$ as a function of time, it is possible to compute *time averages* or *temporal statistics* by integrating along the time axis [31]:

$$\mu_f(k) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-T/2}^{T/2} f_k(t) dt. \quad (3.10)$$

The integral would be replaced by a summation in the case of sampled or discrete-time signals. The time-averaged ACF $\phi_f(\tau, k)$ is given by

$$\phi_f(\tau, k) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-T/2}^{T/2} f_k(t) f_k(t + \tau) dt. \quad (3.11)$$

Similarly, given an observation of a random process as an image $f_k(x, y)$, we may compute averages by integrating over the spatial domain, to obtain spatial averages or spatial statistics; see Figure 3.4. The spatial mean of the image $f_k(x, y)$ is given by

$$\mu_f(k) = \frac{1}{A} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_k(x, y) dx dy, \quad (3.12)$$


where A is a normalization factor, such as the actual area of the image. Observe that the spatial mean above is a single-valued entity (a scalar). For a stationary process, the spatial ACF is given by

$$\phi_f(\alpha, \beta, k) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_k(x, y) f_k(x + \alpha, y + \beta) dx dy. \quad (3.13)$$

A suitable normalization factor, such as the total energy of the image [which is equal to $\phi_f(0, 0)$] may be included, if necessary. The sample index k becomes irrelevant if only one observation is available. In practice, the integrals change to summations over the space of the digital image available.

When we have a 2D image as a function of time, such as TV, video, fluoroscopy, and cine-angiography signals, we have a spatio-temporal signal that may be expressed as $f(x, y, t)$; see Figure 3.5. We may then compute statistics over a single frame $f(x, y, t_1)$ at the instant of time t_1 , which are known as intraframe statistics. We could also compute parameters through multiple frames over a certain period of time, which are called interframe statistics; the signal over a specific period of time may then be treated as a 3D dataset.

Random functions of time may thus be characterized in terms of ensemble and/or temporal statistics. Random functions of space may be represented

**FIGURE 3.5**

Spatial and temporal statistics of a video signal.

by their ensemble and/or spatial statistics. Figure 3.4 shows the distinction between ensemble and spatial averaging. Figure 3.5 illustrates the combined use of spatial and temporal statistics to analyze a video signal in (x, y, t) .

The mean does not play an important role in 1D signal analysis: it is usually assumed to be zero, and often subtracted out if it is not zero. However, the mean of an image represents its average intensity or density; removal of the mean leads to an image with only the edges and the fluctuations about the mean being depicted.

The ACF plays an important role in the characterization of random processes. The Fourier transform of the ACF is the power spectral density (PSD) function, which is useful in frequency-domain analysis. Statistical functions as above are useful in the analysis of the behavior of random processes, and in modeling, spectrum analysis, filter design, data compression, and data communication.

3.1.2 Examples of noise PDFs

As we have already seen, several types of noise sources are encountered in biomedical imaging. Depending upon the characteristics of the noise source and the phenomena involved in the generation of the signal and noise values, we encounter a few different types of PDFs, some of which are described in the following paragraphs [3, 128, 173].

Gaussian: The most commonly encountered and used noise PDF is the Gaussian or normal PDF, expressed as [3, 128]

$$p_x(x) = \frac{1}{\sqrt{2\pi}\sigma_x} \exp\left[-\frac{(x - \mu_x)^2}{2\sigma_x^2}\right]. \quad (3.14)$$

A Gaussian PDF is completely specified by its mean μ_x and variance σ_x^2 . Figure 3.6 shows three Gaussian PDFs with $\mu = 0, \sigma = 1$; $\mu = 0, \sigma = 2$; and $\mu = 3, \sigma = 1$. See also Figures 3.2 and 3.3.

When we have two jointly normal random processes x and y , the bivariate normal PDF is given by

$$p_{x,y}(x, y) = \frac{1}{\sqrt{4\pi^2(1 - \gamma^2)} \sigma_x \sigma_y} \times \exp\left\{-\frac{1}{2(1 - \gamma^2)} \left[\frac{(x - \mu_x)^2}{\sigma_x^2} - \frac{2\gamma(x - \mu_x)(y - \mu_y)}{\sigma_x \sigma_y} + \frac{(y - \mu_y)^2}{\sigma_y^2}\right]\right\}, \quad (3.15)$$

where γ is the correlation coefficient given by

$$\gamma = \frac{E[(x - \mu_x)(y - \mu_y)]}{\sigma_x \sigma_y}. \quad (3.16)$$

If $\gamma = 0$, the two processes are uncorrelated. The bivariate normal PDF then reduces to a product of two univariate Gaussians, which implies that the two processes are statistically independent.





FIGURE 3.6

Three Gaussian PDFs. Solid line: $\mu = 0$, $\sigma = 1$. Dashed line: $\mu = 0$, $\sigma = 2$. Dotted line: $\mu = 3$, $\sigma = 1$.

The importance of the Gaussian PDF in practice arises from a phenomenon that is expressed as the *central limit theorem* [3, 128]: The PDF of a random process that is the sum of several statistically independent random processes is equal to the cascaded convolution of their individual PDFs. When a large number of functions are convolved in cascade, the result tends toward a Gaussian-shaped function regardless of the forms of the individual functions. In practice, an image is typically affected by a series of independent sources of additive noise; the net noise PDF may then be assumed to be a Gaussian.

Uniform: All possible values of a uniformly distributed random process have equal probability of occurrence. The PDF of such a random process over the range (a, b) is a rectangle of height $\frac{1}{(b-a)}$ over the range (a, b) . The mean of the process is $\frac{(a+b)}{2}$, and the variance is $\frac{(b-a)^2}{12}$. Figure 3.7 shows two uniform PDFs corresponding to random processes with values spread over the ranges $(-10, 10)$ and $(-5, 5)$. The quantization of gray levels in an image to a finite number of integers leads to an error or noise that is uniformly distributed.

Poisson: The counting of discrete random events such as the number of photons emitted by a source or detected by a sensor in a given interval of time leads to a random variable with a Poisson PDF. The discrete nature of the packets of energy (that is, photons) and the statistical randomness in their emission and detection contribute to uncertainty, which is reflected as

**FIGURE 3.7**

Two uniform PDFs. Solid line: $\mu = 0$, range = $(-10, 10)$. Dashed line: $\mu = 0$, range = $(-5, 5)$.

quantum noise, photon noise, mottle, or Poisson noise in images. Shot noise in electronic devices may also be modeled as Poisson noise.

One of the formulations of the Poisson PDF is as follows: The probability that k photons are detected in a certain interval is given by

$$P(k) = \exp(-\mu) \frac{\mu^k}{k!}. \quad (3.17)$$

Here, μ is the mean of the process, which represents the average number of photons counted in the specified interval over many trials. The values of $P(k)$ for all (integer) k is the Poisson PDF. The variance of the Poisson PDF is equal to its mean.

The Poisson PDF tends toward the Gaussian PDF for large mean values. Figure 3.8 shows two Poisson PDFs along with the Gaussians for the same parameters; it is seen that the Poisson and Gaussian PDFs for $\mu = \sigma^2 = 20$ match each other well.




FIGURE 3.8

Two Poisson PDFs with the corresponding Gaussian PDFs superimposed. Bars with \times and dashed envelope: $\mu = \sigma^2 = 4$. Bars with \circ and solid envelope: $\mu = \sigma^2 = 20$.

Laplacian: The Laplacian PDF is given by the function

$$p_x(x) = \frac{1}{\sqrt{2} \sigma_x} \exp \left\{ -\frac{\sqrt{2} |x - \mu_x|}{\sigma_x} \right\}, \quad (3.18)$$

where μ_x and σ_x^2 are the mean and variance, respectively, of the process. Figure 3.9 shows two Laplacian PDFs. Error values in linear prediction have been observed to display Laplacian PDFs [174].




FIGURE 3.9

Two Laplacian PDFs with $\mu = 0$, $\sigma^2 = 1$ (solid) and $\mu = 0$, $\sigma^2 = 4$ (dashed).

Rayleigh: The Rayleigh PDF is given by the function

$$p_x(x) = \frac{2}{b} (x - a) \exp \left\{ -\frac{(x - a)^2}{b} \right\} u(x - a), \quad (3.19)$$

where $u(x)$ is the unit step function such that

$$u(x) = \begin{cases} 1 & \text{if } x \geq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (3.20)$$

The mean and variance of the Rayleigh PDF are determined by the parameters a and b as [173] $\mu_x = a + \sqrt{(\pi b/4)}$ and $\sigma_x^2 = b(4 - \pi)/4$.

Figure 3.10 shows a Rayleigh PDF with $a = 1$ and $b = 4$. The Rayleigh PDF has been used to model speckle noise [175].




FIGURE 3.10
Rayleigh PDF with $a = 1$ and $b = 4$.

3.1.3 Structured noise

Power-line interference at 50 *Hz* or 60 *Hz* is a common artifact in many biomedical signals. The typical waveform of the interference is known in advance and, hence, it could be considered to be a form of structured noise. It should, however, be noted that the phase of the interfering waveform will not usually be known. Furthermore, the interfering waveform may not be an exact sinusoid; this is indicated by the presence of harmonics of the fundamental 50 *Hz* or 60 *Hz* component. Notch and comb filters may be used to remove power-line artifact [31]. It is not common to encounter power-line interference in biomedical images.

A common form of structured noise in biomedical images is the grid artifact. On rare occasions, the grid used in X-ray imaging may not translate or oscillate as designed; the stationary grid then casts its shadow, which is superimposed upon the image of the subject. Stationary grids with parallel strips create an artifact that is a set of parallel lines which is also periodic.

Depending upon the strip width, strip orientation, and sampling resolution, the artifactual lines may appear as thin straight lines, or as relatively thick rectangular stripes with some internal gray-scale variation; see Figure 1.13 and Section 3.4.2.

In some imaging applications, grids and frames may be used to position the object or subject being imaged. The image of the grid or frame could serve a useful purpose in image registration and calibration, as well as in the modeling and removal of geometric distortion. The same patterns may be considered to be artifacts in other applications of image processing.

Labels indicating patient identification, patient positioning, imaging parameters, and other details are routinely used in X-ray imaging. Ideally, the image cast by such labels should be well removed from the image of the patient on the acquired image, although, occasionally, they may get connected or even overlap. Such items interfere in image analysis when the procedure is applied to the entire image. Preprocessing techniques are then required to recognize and remove such artifacts; see Section 5.9 for examples.

Surgical implants such as staples, pins, and screws create difficulties and artifacts in X-ray, MR, CT, and ultrasound imaging. The advantage with such artifacts is that the precise composition and geometry of the implants are known by design and the manufacturers' specifications. Methods may then be designed to remove each specific artifact.

3.1.4 Physiological interference

As we have already noted, the human body is a complex conglomeration of several systems and processes. Several physiological processes could be active at a given instant of time, each one affecting the system or process of interest in diverse ways. A patient or experimental subject may not be able to exercise control on all of his or her physiological processes and systems. The effect of systems or processes other than those of interest on the image being acquired may be termed as physiological interference; several examples are listed below.

- Effect of breathing on a chest X-ray image.
- Effect of breathing, peristalsis, and movement of material through the gastro-intestinal system on CT images of the abdomen.
- Effect of cardiovascular activity on CT images of the chest.
- Effect of pulsatile movement of arteries in subtraction angiography.

Physiological interference may not be characterized by any specific waveform, pattern, or spectral content, and is typically dynamic and nonstationary (varying with the level of the activity of relevance and hence with time; see Section 3.1.6 for a discussion on stationarity). Thus, simple, linear bandpass filters will usually not be effective in removing physiological interference.

Normal anatomical details such as the ribs in chest X-ray images and the skull in brain imaging may also be considered to be artifacts when other details in such images are of primary interest. Methods may need to be developed to remove their effects before the details of interest may be analyzed.

3.1.5 Other types of noise and artifact

Systematic errors are caused by several factors such as geometric distortion, miscalibration, nonlinear response of detectors, sampling, and quantization [3]. Such errors may be modeled from a knowledge of the corresponding parameters, which may be determined from specifications, measured experimentally, or derived mathematically.

A few other types of artifact that cannot be easily categorized into the groups discussed above are the following:

- Punctate or shot noise due to dust on the screen, film, or examination table.
- Scratches on film that could appear as intense line segments.
- Shot noise due to inactive elements in a detector array.
- Salt-and-pepper noise due to impulsive noise, leading to black or white pixels at the extreme ends of the pixel-value range.
- Film-grain noise due to scanning of films with high resolution.
- Punctate noise in chest X-ray or mammographic images caused by cosmetic powder or deodorant (which could masquerade as microcalcifications).
- Superimposed images of clothing accessories such as pins, hooks, buttons, and jewelry.

3.1.6 Stationary versus nonstationary processes

Random processes may be characterized in terms of their temporal/spatial and/or ensemble statistics. A random process is said to be *stationary in the strict sense* or *strongly stationary* if its statistics are not affected by a shift in the origin of time or space. In most practical applications, only the first-order and second-order averages are used. A random process is said to be *weakly stationary* or *stationary in the wide sense* if its mean is a constant and its ACF depends only upon the difference (or shift) in time or space. Then, we have $\mu_f(x_1, y_1) = \mu_f$ and $\phi_f(x_1, x_1 + \alpha, y_1, y_1 + \beta) = \phi_f(\alpha, \beta)$. The ACF is now a function of the shift parameters α and β only; the PSD of the process does not vary with space.

A stationary process is said to be *ergodic* if the temporal statistics computed are independent of the sample observed; that is, the same results are obtained with any sample observation $f_k(t)$. The time averages of the process are then independent of k : $\mu_f(k) = \mu_f$ and $\phi_f(\tau, k) = \phi_f(\tau)$. All ensemble statistics may be replaced by temporal statistics when analyzing ergodic processes. Ergodic processes are an important type of stationary random processes because their statistics may be computed from a single observation as a function of time. The same concept may be extended to functions of space as well, although the term “ergodic” is not commonly applied to spatial functions.

Signals or processes that do not meet the conditions described above may, in general, be called *nonstationary processes*. A nonstationary process possesses statistics that vary with time or space. The statistics of most images vary over space; indeed, such variations are the source of pictorial information. Most biomedical systems are dynamic systems and produce nonstationary signals and images. However, a physical or physiological system has limitations in the rate at which it can change its characteristics. This limitation facilitates the breaking of a signal into segments of short duration (typically a few tens of milliseconds), over which the statistics of interest may be assumed to remain constant [31]. The signal is then referred to as a *quasistationary process*; Techniques designed for stationary signals may then be extended and applied to nonstationary signals. Analysis of signals by this approach is known as *short-time analysis* [31, 176]. On the same token, the characteristics of the features in an image vary over relatively large scales of space; statistical parameters within small regions of space, within an object, or within an organ of a given type may be assumed to remain constant. The image may then be assumed to be *block-wise stationary*, which permits sectioned or block-by-block processing or *moving-window processing* using techniques designed for stationary processes [177, 178]. Figure 3.11 illustrates the notion of computing statistics within a moving window.

Certain systems, such as the cardiac system, normally perform rhythmic operations. Considering the dynamics of the cardiac system, it is obvious that the system is nonstationary. However, various phases of the cardiac cycle — as well as the related components of the associated electrocardiogram (ECG), phonocardiogram (PCG), and carotid pulse signals — repeat over time in an almost-periodic manner. A given phase of the process or signal possesses statistics that vary from those of the other phases; however, the statistics of a specific phase repeat cyclically. For example, the statistics of the PCG signal vary within the duration of a cardiac cycle, especially when murmurs are present, but repeat themselves at regular intervals over successive cardiac cycles. Such signals are referred to as *cyclo-stationary signals* [31]. The cyclical repetition of the process facilitates synchronized ensemble averaging (see Sections 3.2 and 3.10), using epochs or events extracted from an observation of the signal over many cycles.

The cyclical nature of cardiac activity may be exploited for synchronized averaging to reduce noise and improve the SNR of the ECG and PCG [31]. The




FIGURE 3.11

Block-by-block processing of an image. Statistics computed by using the pixels within the window shown with solid lines (3×3 pixels) are applicable to the pixel marked with the @ symbol. Statistics for use when processing the pixel marked with the # symbol (5×5 pixels) are computed by using the pixels within the window shown with dashed lines.

same technique may also be extended to imaging the heart: In gated blood-pool imaging, nuclear medicine images of the heart are acquired in several parts over short intervals of time. Images acquired at the same phases of the cardiac cycle — determined by using the ECG signal as a reference, trigger, or “gating” signal — are accumulated over several cardiac cycles. A sequence of such gated and averaged frames over a full cardiac cycle may then be played as a video or a movie to visualize the time-varying size and contents of the left ventricle. (See Section 3.10 for illustration of gated blood-pool imaging.)

3.1.7 Covariance and cross-correlation

When two random processes f and g need to be compared, we could compute the covariance between them as

$$\sigma_{fg} = E[(f - \mu_f)(g - \mu_g)] = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} (f - \mu_f)(g - \mu_g) p_{f,g}(f, g) df dg, \quad (3.21)$$

where $p_{f,g}(f, g)$ is the joint PDF of the two processes, and the image coordinates have been omitted for the sake of compact notation. The covariance parameter may be normalized to get the correlation coefficient, defined as

$$\rho_{fg} = \frac{\sigma_{fg}}{\sigma_f \sigma_g}, \quad (3.22)$$

with $-1 \leq \rho_{fg} \leq +1$. A high covariance indicates that the two processes have similar statistical variability or behavior. The processes f and g are said to be uncorrelated if $\rho_{fg} = 0$. Two processes that are statistically independent are also uncorrelated; the converse of this property is, in general, not true.

When dealing with random processes f and g that are functions of space, the cross-correlation function (CCF) between them is defined as

$$\phi_{fg}(\alpha, \beta) = E[f(x, y) g(x + \alpha, y + \beta)]. \quad (3.23)$$

In situations where the PDF is not available, the expressions above may be approximated by the corresponding ensemble or spatial statistics. Correlation functions are useful in analyzing the nature of variability and spectral bandwidth of images, as well as for the detection of objects by template matching.

3.1.8 Signal-dependent noise

Noise may be categorized as being independent of the signal of interest if no statistical parameter of any order of the noise process is a function of the signal. Although it is common to assume that the noise present in a signal or image is statistically independent of the true signal (or image) of interest, several cases exist in biomedical imaging where this assumption is not valid, and the noise is functionally related to or dependent upon the signal. The following paragraphs provide brief notes on a few types of signal-dependent noise encountered in biomedical imaging.

Poisson noise: Imaging systems that operate in low-light conditions, or in low-dose radiation conditions such as nuclear medicine imaging, are often affected by photon noise that can be modeled as a Poisson process [179, 180]; see Section 3.1.2. The probabilistic description of an observed image (pixel value) $g_o(m, n)$ under the conditions of a Poisson process is given by

$$P(g_o(m, n) | f(m, n), \lambda) = \frac{[\lambda f(m, n)]^{g_o(m, n)} \exp[-\lambda f(m, n)]}{g_o(m, n)!}, \quad (3.24)$$

where $f(m, n)$ is the undegraded pixel value (the observation in the absence of any noise), and λ is a proportionality factor. Because the mean of the degraded image g_o is given by

$$E[g_o(m, n)] = \lambda E[f(m, n)], \quad (3.25)$$

images corrupted with Poisson noise are usually normalized as

$$g(m, n) = \frac{g_o(m, n)}{\lambda}. \quad (3.26)$$

It has been shown [179] that, in this case, Poisson noise may be modeled as stationary noise uncorrelated with the signal and added to the signal as in Equation 3.4, with zero mean and variance given by

$$\sigma_{\eta}^2(m, n) = \frac{E[f(m, n)]}{\lambda} = \frac{E[g(m, n)]}{\lambda}. \quad (3.27)$$

Film-grain noise: The granular structure of film due to the silver-halide grains used contributes noise to the recorded image, which is known as film-grain noise. When images recorded on photographic film are digitized in order to be processed by a digital computer, film-grain noise is a significant source of degradation of the information. According to Froehlich et al. [181], the model for an image corrupted by film-grain noise is given by

$$g(m, n) = f(m, n) + \kappa \mathcal{F}[f(m, n)] \eta_1(m, n) + \eta_2(m, n), \quad (3.28)$$

where κ is a proportionality factor, $\mathcal{F}[\]$ is a mathematical function, and $\eta_1(m, n)$ and $\eta_2(m, n)$ are samples from two random processes independent of the signal. This model may be taken to represent a general imaging situation that includes signal-independent noise as well as signal-dependent noise, and the noise could be additive or multiplicative. Observe that the model reduces to the simple signal-independent additive noise model in Equation 3.4 if $\kappa = 0$.

Froehlich et al. [181] modeled film-grain noise with $\mathcal{F}[f(m, n)] = [f(m, n)]^p$, using $p = 0.5$. The two noise processes η_1 and η_2 were assumed to be Gaussian-distributed, uncorrelated, zero-mean random processes. According to this model, the noise that corrupts the image has two components: one that is signal-dependent through the factor $\kappa \sqrt{f(m, n)} \eta_1(m, n)$, and another that is signal-independent given by $\eta_2(m, n)$. Film-grain noise may be modeled as additive noise as in Equation 3.4, with $\eta(m, n) = \kappa \sqrt{f(m, n)} \eta_1(m, n) + \eta_2(m, n)$. It can be shown that $\eta(m, n)$ as above is stationary, has zero mean, and has its variance given by [182]

$$\sigma_\eta^2(m, n) = \kappa^2 E[g(m, n)] \sigma_{\eta_1}^2 + \sigma_{\eta_2}^2. \quad (3.29)$$

The fact that the mean of the corrupted image equals the mean of the noise-free image has been used in arriving at the relationship above.

Speckle noise: Speckle noise corrupts images that are obtained by coherent radiation, such as synthetic-aperture radar (SAR), ultrasound, laser, and sonar. When an object being scanned by a laser beam has random surface roughness with details of the order of the wavelength of the laser, the imaging system will be unable to resolve the details of the object's roughness; this results in speckle noise. The most widely used model for speckle noise is a multiplicative model, given as [175, 183, 184, 185, 186, 187]

$$g(m, n) = f(m, n) \eta_1(m, n), \quad (3.30)$$

where $\eta_1(m, n)$ is a stationary noise process that is assumed to be uncorrelated with the image. If the mean of the noise process μ_{η_1} is not equal to one, the noisy image may be normalized by dividing by μ_{η_1} such that, in the normalized image, the multiplicative noise has its mean equal to one. Depending upon the specific application, the distribution of the noise may be assumed to be exponential [175, 183, 186, 187], Gaussian [184], or Rayleigh [175].

The multiplicative model in Equation 3.30 may be converted to the additive model as in Equation 3.4 with $\eta(m, n)$ being zero-mean additive noise having

a space-variant, signal-dependent variance given by [188]

$$\sigma_{\eta}^2(m, n) = \frac{\sigma_{\eta}^2}{1 + \sigma_{\eta}^2} [\sigma_g^2(m, n) + \mu_g^2(m, n)]. \quad (3.31)$$

In the expression above, $\sigma_g^2(m, n)$ and $\mu_g(m, n)$ are the variance and the mean of the noisy image at the point (m, n) , respectively.

Transformation of signal-dependent noise to signal-independent noise: In the model used by Naderi and Sawchuk [182] and Arsenault et al. [189, 190], the signal-independent component of the noise as in Equation 3.28 is assumed to be zero. In this case, it has been shown [189, 190, 191] that by applying an appropriate transformation to the whole image, the noise can be made signal-independent. One of the transformations proposed is [189, 190]

$$T[g(m, n)] = \alpha \sqrt{g(m, n)}, \quad (3.32)$$

where α is an appropriate normalizing constant. It has been shown that the noise in the transformed image is additive, has a Gaussian distribution, is unbiased, and has a standard deviation that no longer depends on the signal but is given by $\frac{\alpha \kappa}{2}$.

3.2 Synchronized or Multiframe Averaging

In certain applications of imaging, if the object being imaged can remain free from motion or change of any kind (internal or external) over a long period of time compared to the time required to record an image, it becomes possible to acquire several frames of images of the object in precisely the same state or condition. Then, the frames may be averaged to reduce noise; this is known as multiframe averaging. The method may be extended to the imaging of dynamic systems whose movements follow a rhythm or cycle with phases that can be determined by another signal, such as the cardiac system whose phases of contraction are indicated by the ECG signal. Then, several image frames may be acquired at the same phase of the rhythmic movement over successive cycles, and averaged to reduce noise. Such a process is known as synchronized averaging. The process may be repeated or triggered at every phase of interest. (*Note:* A process as above in nuclear medicine imaging may be viewed simply as counting the photons emitted over a long period of time in total, albeit in a succession of short intervals gated to a particular phase of contraction of the heart. Ignoring the last step of division by the number of frames to obtain the average, the process simply accumulates the photon counts over the frames acquired.)

Synchronized averaging is a useful technique in the acquisition of several biomedical signals [31]. Observe that averaging as above is a form of ensemble averaging.

Let us represent a single image frame in a situation as above as

$$g_i(x, y) = f(x, y) + \eta_i(x, y), \quad (3.33)$$

where $g_i(x, y)$ is the i^{th} observed frame of the image $f(x, y)$, and $\eta_i(x, y)$ is the noise in the same frame. Let us assume that the noise process is independent of the signal source. Observe that the desired (original) image $f(x, y)$ is invariant from one frame to another. It follows that $\sigma_{g_i(x, y)}^2 = \sigma_{\eta_i(x, y)}^2$; that is, the variance at every pixel in the observed noisy image is equal to the corresponding variance of the noise process.

If M frames of the image are acquired and averaged, the averaged image is given by

$$\bar{g}(x, y) = \frac{1}{M} \sum_{i=1}^M g_i(x, y). \quad (3.34)$$

If the mean of the noise process is zero, we have $\sum_{i=1}^M \eta_i(x, y) \rightarrow 0$ as $M \rightarrow \infty$ (in practice, as the number of frames averaged increases to a large number). Then, it follows that [8]

$$E[\bar{g}(x, y)] = f(x, y), \quad (3.35)$$

and


$$\sigma_{\bar{g}(x, y)}^2 = \frac{1}{M} \sigma_{\eta(x, y)}^2. \quad (3.36)$$

Thus, the variance at every pixel in the averaged image is reduced by a factor of $\frac{1}{M}$ from that in a single frame; the SNR is improved by the factor \sqrt{M} .

The most important requirement in this procedure is that the frames being averaged be mutually synchronized, aligned, or registered. Any motion, change, or displacement between the frames will lead to smearing and distortion.

Example: Figure 3.12 (a) shows a test image with several geometrical objects placed at random. Images (b) and (c) show two examples of eight noisy frames of the test image that were obtained by adding Gaussian-distributed random noise samples. The results of averaging two, four, and eight noisy frames [including the two in (b) and (c)] are shown in parts (d), (e), and (f), respectively. It is seen that averaging using increasing numbers of frames of the noisy image leads to a reduction of the noise; the decreasing trend in the RMS values of the processed images (given in the caption of Figure 3.12) confirms the expected effect of averaging.

See Section 3.9 for an illustration of the application of multiframe averaging in confocal microscopy. See also Section 3.10 for details on gated blood-pool imaging in nuclear medicine.

**FIGURE 3.12**

(a) “Shapes”: a 128×128 test image with various geometrical objects placed at random. (b) Image in (a) with Gaussian noise added, with $\mu = 0$, $\sigma^2 = 0.01$ (normalized), RMS error = 19.32. (c) Second version of noisy image, RMS error = 19.54. Result of multiframe averaging using (d) the two frames in (b) and (c), RMS error = 15.30; (e) four frames, RMS error = 12.51; (f) eight frames, RMS error = 10.99.

3.3 Space-domain Local-statistics-based Filters

Consider the practical situation when we are given a single, noisy observation of an image of finite size. We do not have access to an ensemble of images to perform multiframe (synchronized) averaging, and spatial statistics computed over the entire image frame will lead to scalar values that do not assist in removing the noise and obtaining a cleaner image. Furthermore, we should also accommodate for nonstationarity of the image. In such situations, moving-window filtering using windows of small size such as 3×3 , 5×5 , or 7×7 pixels becomes a valuable option; rectangular windows as well as windows of other shapes may also be considered where appropriate. Various statistical parameters of the pixels within such a moving window may be computed, with the result being applied to the pixel in the output image at the same location where the window is placed (centered) on the input image; see Figure 3.13. Observe that only the pixel values in the given (input) image are used in the filtering process; the output is stored in a separate array. Figure 3.14 illustrates a few different neighborhood shapes that are commonly used in moving-window image filtering [192].




FIGURE 3.13

Moving-window filtering of an image. The size of the moving window in the illustration is 5×5 pixels. Statistics computed by using the pixels within the window are applied to the pixel at the same location in the output image. The moving window is shown for two pixel locations marked # and @.

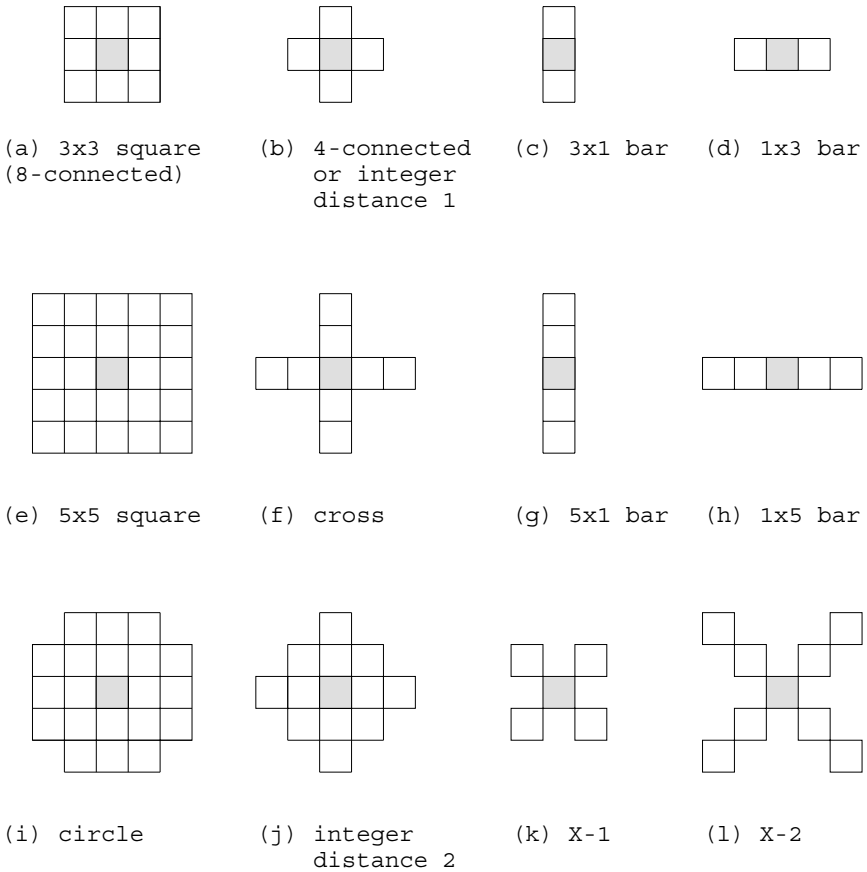


FIGURE 3.14

A few commonly used moving-window neighborhood shapes for image filtering. The result computed by using the pixels within a window is applied to the pixel at the location of its center, shown shaded, in the output image.

3.3.1 The mean filter

If we were to select the pixels in a small neighborhood around the pixel to be processed, the following assumptions may be made:

- the image component is relatively constant; that is, the image is quasistationary; and
- the only variations in the neighborhood are due to noise.

Further assumptions regarding the noise process that are typically made are that it is additive, is independent of the image, and has zero mean. Then, if we were to take the mean of the pixels in the neighborhood, the result will tend toward the true pixel value in the original, uncorrupted image. In essence, a spatial collection of pixels around the pixel being processed is substituted for an ensemble of pixels at the same location from multiple frames in the averaging process; that is, the image-generating process is assumed to be ergodic.

It is common to use a 3×3 or 8-connected neighborhood as in Figure 3.14 (a) for mean filtering. Then, the output of the filter $g(m, n)$ is given by

$$g(m, n) = \frac{1}{9} \sum_{\alpha=-1}^1 \sum_{\beta=-1}^1 f(m + \alpha, n + \beta), \quad (3.37)$$

where $f(m, n)$ is the input image. The summation above may be expanded as

$$g(m, n) = \frac{1}{9} \times \begin{bmatrix} f(m-1, n-1) + f(m-1, n) + f(m-1, n+1) \\ + f(m, n-1) + f(m, n) + f(m, n+1) \\ + f(m+1, n-1) + f(m+1, n) + f(m+1, n+1) \end{bmatrix}. \quad (3.38)$$

The same result is also achieved via convolution of the image $f(m, n)$ with the 3×3 array or mask

$$\frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}. \quad (3.39)$$

Note that the operation above cannot be directly applied at the edges of the input image array; it is common to extend the input array with a border of zero-valued pixels to permit filtering of the pixels at the edges. One may also elect not to process the pixels at the edges, or to replace them with the average of the available neighbors.

The mean filter can suppress Gaussian and uniformly distributed noise effectively in relatively homogeneous areas of an image. However, the operation leads to blurring at the edges of the objects in the image, and also to the loss of fine details and texture. Regardless, mean filtering is commonly employed

to remove noise and smooth images. The blurring of edges may be prevented to some extent by not applying the mean filter if the difference between the pixel being processed and the mean of its neighbors is greater than a certain threshold; this condition, however, makes the filter nonlinear.

3.3.2 The median filter


The median of a collection of samples is the value that splits the population in half: half the number of pixels in the collection will have values less than the median and half will have values greater than the median. In small populations of pixels under the constraint that the result be an integer, approximations will have to be made: the most common procedure rank-orders the pixels in a neighborhood containing an odd number of pixels, and the pixel value at the middle of the list is selected as the median. The procedure also permits the application of order-statistic filters [193]: the i^{th} element in a rank-ordered list of values is known as the i^{th} order statistic. The median filter is an order-statistic filter of order $N/2$ where N is the size of the filter; that is, the number of values used to derive the output.

The median filter is a nonlinear filter. Its success in filtering depends upon the number of the samples used to derive the output, as well as the spatial configuration of the neighborhood used to select the samples.


The median filter provides better noise removal than the mean filter without blurring, especially when the noise has a long-tailed PDF (resulting in outliers) and in the case of salt-and-pepper noise. However, the median filter could result in the clipping of corners and distortion of the shape of sharp-edged objects; median filtering with large neighborhoods could also result in the complete elimination of small objects. Neighborhoods that are not square in shape are often used for median filtering in order to limit the clipping of corners and other types of distortion of shape; see Figure 3.14.

Examples: Figure 3.15 (a) shows a 1D test signal with a rectangular pulse; part (b) of the same figure shows the test signal degraded with impulse (shot) noise. The results of filtering the noisy signal using the mean and median with filter length $N = 3$ are shown in plots (c) and (d), respectively, of Figure 3.15. The mean filter has blurred the edges of the pulse; it has also created artifacts in the form of small hills and valleys. The median filter has removed the noise without distorting the signal.

Figure 3.16 (a) shows a 1D test signal with two rectangular pulses; part (b) of the same figure shows the test signal degraded with uniformly distributed noise. The results of filtering the noisy signal using the mean and median with filter length $N = 5$ are shown in plots (c) and (d), respectively, of Figure 3.16. The mean filter has reduced the noise level, but has also blurred the edges of the pulses; in addition, the strength of the first, short pulse has been reduced. The median filter has removed the noise to some extent without distorting the edges of the long pulse; however, the short pulse has been obliterated.

**FIGURE 3.15**

(a) A 1D test signal with a rectangular pulse. (b) Degraded signal with impulse or shot noise. Result of filtering the degraded signal using (c) the mean and (d) the median operation with a sliding window of $N = 3$ samples.

**FIGURE 3.16**

(a) A 1D test signal with two rectangular pulses. (b) Degraded signal with uniformly distributed noise. Result of filtering the degraded signal using (c) the mean, and (d) the median operation with a sliding window of $N = 5$ samples.

Figure 3.17 shows the original test image “Shapes”, the test image degraded by the addition of Gaussian-distributed random noise with $\mu = 0$ and $\sigma^2 = 0.01$ (normalized), and the results of filtering the noisy image with the 3×3 and 5×5 mean and median filters. The RMS errors of the noisy and filtered images with respect to the test image are given in the figure caption. All of the filters except the 3×3 median have led to an increase in the RMS error. The blurring effect of the mean filter is readily seen in the results. Close observation of the result of 3×3 median filtering [Figure 3.17 (d)] shows that the filter has resulted in distortion of the shapes, in particular, clipping of the corners of the objects. The 5×5 median filter has led to the complete removal of small objects; see Figure 3.17 (f). Observe that the results of the 3×3 mean and 5×5 median filters have similar RMS error values; however, the blurring effect in the former case, and the distortion of shape information as well as the loss of small objects in the latter case need to be considered carefully.

Figure 3.18 gives a similar set of images with the noise being Poisson distributed. A comparable set with speckle noise is shown in Figure 3.19. Although the filters have reduced the noise to some extent, the distortions introduced have led to increased RMS errors for all of the results.

Figures 3.20 and 3.21 show two cases with salt-and-pepper noise, the density of pixels affected by noise being 0.05 and 0.1, respectively. The 3×3 median filter has given good results in both cases with the lowest RMS error and the least distortion. The 5×5 median filter has led to significant shape distortion and the loss of a few small features.

Figure 3.22 shows the normalized histograms of the Shapes test image and its degraded versions with Gaussian, Poisson, and speckle noise. It is evident that the signal-dependent Poisson noise and speckle noise have affected the histogram in a different manner compared to the signal-independent Gaussian noise.

Figure 3.23 shows the results of filtering the Peppers test image affected by Gaussian noise. Although the RMS errors of the filtered images are low compared to that of the noisy image, the filters have introduced a mottled appearance and fine texture in the smooth regions of the original image. Figure 3.24 shows the case with Poisson noise, where the 5×5 filters have provided visually good results, regardless of the RMS errors.

All of the filters have performed reasonably well in the presence of speckle noise, as illustrated in Figure 3.25, in terms of the reduction of RMS error. However, the visual quality of the images is poor.

Figures 3.26 and 3.27 show that the median filter has provided good results in filtering salt-and-pepper noise. Although the RMS values of the results of the mean filters are lower than those of the noisy images, visual inspection of the results indicates the undesirable effects of blurring and mottled appearance.


The RMS error (or the MSE) is commonly used to compare the results of various image processing operations; however, the examples presented above

illustrate the limitations in using the RMS error in comparing images with different types of artifact and distortion. In some of the results shown, an image with a higher RMS error may present better visual quality than another image with a lower RMS error. Visual inspection and analysis of the results by qualified users or experts in the domain of application is important. It is also important to test the proposed methods with phantoms or test images that demonstrate the characteristics that are relevant to the specific application being considered. Assessment of the advantages provided by the filtered results in further processing and analysis, such as the effects on diagnostic accuracy in the case of medical images, is another approach to evaluate the results of filtering.


3.3.3 Order-statistic filters

The class of order-statistic filters [193] is large, and includes several nonlinear filters that are useful in filtering different types of noise in images. The first step in order-statistic filtering is to rank-order, from the minimum to the maximum, the pixel values in an appropriate neighborhood of the pixel being processed. The i^{th} entry in the list is the output of the i^{th} order-statistic filter. A few order-statistic filters of particular interest are the following:


- *Min filter*: the first entry in the rank-ordered list, useful in removing high-valued impulse noise (isolated bright spots or “salt” noise).
- *Max filter*: the last entry in the rank-ordered list, useful in removing low-valued impulse noise (isolated dark spots or “pepper” noise).
- *Min/Max filter*: sequential application of the Min and Max filters, useful in removing salt-and-pepper noise.
- *Median filter*: the entry in the middle of the list. The median filter is the most popular and commonly used filter among the order-statistic filters; see Section 3.3.2 for detailed discussion and illustration of the median filter.
- *α -trimmed mean filter*: the mean of a reduced list where the first α and the last α of the list is rejected, with $0 \leq \alpha < 0.5$. Outliers, that is pixels with values very different from the rest of the pixels in the list, are rejected by the trimming process. A value close to 0.5 for α rejects the entire list except the median or a few values close to it, and the output is close to or equal to that of the median filter. The mean of the trimmed list provides a compromise between the generic mean and median filters.
- *L-filters*: a weighted combination of all of the elements in the rank-ordered list. The use of appropriate weights can provide outputs equal

**FIGURE 3.17**


(a) Shapes test image. (b) Image in (a) with Gaussian noise added, with $\mu = 0$, $\sigma^2 = 0.01$ (normalized), RMS error = 19.56. Result of filtering the noisy image in (b) using: (c) 3×3 mean, RMS error = 22.62; (d) 3×3 median, RMS error = 15.40; (e) 5×5 mean, RMS error = 28.08; (f) 5×5 median, RMS error = 22.35.

**FIGURE 3.18**


(a) Shapes test image. (b) Image in (a) with Poisson noise, RMS error = 5.00. Result of filtering the noisy image in (b) using: (c) 3×3 mean, RMS error = 19.40; (d) 3×3 median, RMS error = 13.19; (e) 5×5 mean, RMS error = 25.85; (f) 5×5 median, RMS error = 23.35.

**FIGURE 3.19**


(a) Shapes test image. (b) Image in (a) with speckle noise, with $\mu = 0$, $\sigma^2 = 0.04$ (normalized), RMS error = 12.28. Result of filtering the noisy image in (b) using: (c) 3×3 mean, RMS error = 20.30; (d) 3×3 median, RMS error = 15.66; (e) 5×5 mean, RMS error = 26.32; (f) 5×5 median, RMS error = 24.56.

**FIGURE 3.20**


(a) Shapes test image. (b) Image in (a) with salt-and-pepper noise added, with density = 0.05, RMS error = 40.99. Result of filtering the noisy image in (b) using: (c) 3×3 mean, RMS error = 24.85; (d) 3×3 median, RMS error = 14.59; (e) 5×5 mean, RMS error = 28.24; (f) 5×5 median, RMS error = 23.14.

**FIGURE 3.21**


(a) Shapes test image. (b) Image in (a) with salt-and-pepper noise added, with density = 0.1, RMS error = 56.32. Result of filtering the noisy image in (b) using: (c) 3×3 mean, RMS error = 29.87; (d) 3×3 median, RMS error = 15.42; (e) 5×5 mean, RMS error = 31.25; (f) 5×5 median, RMS error = 23.32.

**FIGURE 3.22**


Normalized histograms of (a) the Shapes test image and of the image with (b) Gaussian noise, (c) Poisson noise, and (d) speckle noise. The first histogram has been scaled to display the range of probability $(0, 0.05)$ only; the remaining histograms have been scaled to display the range $(0, 0.015)$ only in order to show the important details. The probability values of gray levels 0 and 255 have been clipped in some of the histograms. Each histogram represents the gray-level range of $[0, 255]$.

**FIGURE 3.23**


(a) “Peppers”: a 512×512 test image. (b) Image in (a) with Gaussian noise added, with $\mu = 0$, $\sigma^2 = 0.01$ (normalized), RMS error = 25.07. Result of filtering the noisy image in (b) using: (c) 3×3 mean, RMS error = 13.62; (d) 3×3 median, RMS error = 13.44; (e) 5×5 mean, RMS error = 16.17; (f) 5×5 median, RMS error = 13.47.

**FIGURE 3.24**


(a) Peppers test image. (b) Image in (a) with Poisson noise, RMS error = 10.94. Result of filtering the noisy image in (b) using: (c) 3×3 mean, RMS error = 11.22; (d) 3×3 median, RMS error = 8.56; (e) 5×5 mean, RMS error = 15.36; (f) 5×5 median, RMS error = 10.83.

**FIGURE 3.25**

(a) Peppers test image. (b) Image in (a) with speckle noise, with $\mu = 0$, $\sigma^2 = 0.04$ (normalized), RMS error = 26.08. Result of filtering the noisy image in (b) using: (c) 3×3 mean, RMS error = 13.68; (d) 3×3 median, RMS error = 15.73; (e) 5×5 mean, RMS error = 16.01; (f) 5×5 median, RMS error = 14.66.

**FIGURE 3.26**

(a) Peppers test image. (b) Image in (a) with salt-and-pepper noise added, with density = 0.05, RMS error = 30.64. Result of filtering the noisy image in (b) using: (c) 3×3 mean, RMS error = 15.17; (d) 3×3 median, RMS error = 7.38; (e) 5×5 mean, RMS error = 16.96; (f) 5×5 median, RMS error = 10.41.

**FIGURE 3.27**

(a) Peppers test image. (b) Image in (a) with salt-and-pepper noise added, with density = 0.1, RMS error = 43.74. Result of filtering the noisy image in (b) using: (c) 3×3 mean, RMS error = 18.98; (d) 3×3 median, RMS error = 8.62; (e) 5×5 mean, RMS error = 18.71; (f) 5×5 median, RMS error = 11.11.

to those of all of the filters listed above, and facilitate the design of several order-statistic-based nonlinear filters.

Order-statistic filters represent a family of nonlinear filters that have gained popularity in image processing due to their characteristics of removing several types of noise without blurring edges, and due to their simple implementation.

3.4 Frequency-domain Filters

Transforming an image from the space domain to the frequency domain using a transform such as the Fourier transform provides advantages in filtering and noise removal. Most images of natural beings, entities, and scenes vary slowly and smoothly across space, and are usually devoid of step-like changes. As a consequence, such images have most of their energy concentrated in small regions around $(u, v) = (0, 0)$ in their spectra. On the other hand, uncorrelated random noise fields have a uniform, flat, or “white” spectrum, with an almost-constant energy level across the entire frequency space. This leads to the common observation that the SNR of a noisy, natural image is higher in low-frequency regions than in high-frequency regions. It becomes evident, then, that such images may be improved in appearance by suppressing or removing their high-frequency components beyond a certain cut-off frequency. It should be recognized, however, that, in removing high-frequency components, along with the noise components, some desired image components will also be sacrificed. Furthermore, noise components in the low-frequency pass-band will continue to remain in the image.

The procedure for Fourier-domain filtering of an image $f(m, n)$ involves the following steps:

1. Compute the 2D Fourier transform $F(k, l)$ of the image. This may require padding the image with zeros to increase its size to an $N \times N$ array, with N being an integral power of 2, if an FFT algorithm is to be used.
2. Design or select an appropriate 2D filter transfer function $H(k, l)$.
3. Obtain the filtered image (in the Fourier domain) as

$$G(k, l) = H(k, l) F(k, l). \quad (3.40)$$

It is common to define $H(k, l)$ as a real function, thereby affecting only the magnitude of the input image spectrum; the phase remains unchanged. Depending upon the definition or computation of $H(k, l)$, the spectrum $F(k, l)$ may have to be centered or folded (see Figures 2.26, 2.27, and 2.28).

4. Compute the inverse Fourier transform of $G(k, l)$. If $F(k, l)$ was folded prior to filtering, it must be unfolded prior to the inverse transformation.
5. If the input image was zero-padded, trim the resulting image $g(m, n)$.

Although the discussion above concentrates on the removal of noise, it should be noted that frequency-domain filtering permits the removal of several types of artifacts, as well as the selection of the desired frequency components in the image in various manners. The transfer function $H(k, l)$ may be specified in several ways in the frequency domain; phase corrections may be introduced in a separate step. It should be recognized that, while dealing with real-valued images, $H(k, l)$ should maintain (conjugate) symmetry in the frequency plane; it is common to use isotropic filter functions.

3.4.1 Removal of high-frequency noise


Lowpass filters are useful in removing high-frequency noise, under the assumption that the noise is additive and that the Fourier components of the original image past a certain frequency cutoff are negligible. The so-called *ideal* lowpass filter is defined in the 2D Fourier space as

$$H(u, v) = \begin{cases} 1 & \text{if } D(u, v) \leq D_0 \\ 0 & \text{otherwise.} \end{cases} \quad (3.41)$$

where $D(u, v) = \sqrt{u^2 + v^2}$ is the distance of the frequency component at (u, v) from the DC point $(u, v) = (0, 0)$, with the spectrum being positioned such that the DC is at its center (see Figures 2.26, 2.27, and 2.28). [Note: The coordinates (u, v) and (k, l) are used interchangeably to represent the 2D frequency space.] D_0 is the cutoff frequency, beyond which all components of the Fourier transform of the given image are set to zero. Figure 3.28 (a) shows the ideal lowpass filter function. Figure 3.29 shows profiles of the ideal and Butterworth lowpass filters.

Example: Figures 3.30 (a) and (b) show the Shapes test image and a noisy version with Gaussian noise added. Figure 3.31 shows the Fourier magnitude spectra of the original and noisy versions of the Shapes image. The sharp edges in the image have resulted in significant high-frequency energy in the spectrum of the image shown in Figure 3.31 (a); in addition, the presence of strong features at certain angles in the image has resulted in the concentration of energy in the corresponding angular bands in the spectrum. The spectrum of the noisy image, shown in Figure 3.31 (b), demonstrates the presence of additional and increased high-frequency components.

Figure 3.30 (c) shows the noisy image filtered using an ideal lowpass filter with the normalized cutoff $D_0 = 0.4$ times the highest frequency along the u or v axis [see Figure 3.28 (a)]. A glaring artifact is readily seen in the result of the filter: while the noise has been reduced, faint echoes of the edges present in the image have appeared in the result. This is due to the

**FIGURE 3.28**

(a) The magnitude transfer function of an ideal lowpass filter. The cutoff frequency D_0 is 0.4 times the maximum frequency (that is, 0.2 times the sampling frequency). (b) The magnitude transfer function of a Butterworth lowpass filter, with normalized cutoff $D_0 = 0.4$ and order $n = 2$. The $(u, v) = (0, 0)$ point is at the center. The gain is proportional to the brightness (white represents 1.0 and black represents 0.0.)




FIGURE 3.29

Profiles of the magnitude transfer functions of an ideal lowpass filter (solid line) and a Butterworth lowpass filter (dashed line), with normalized cutoff $D_0 = 0.4$ and order $n = 2$.

fact that the inverse Fourier transform of the circular ideal filter defined in Equation 3.41 is a Bessel function (see Figure 2.31 for the circle – Bessel Fourier transform pair). Multiplication of the Fourier transform of the image with the circle function is equivalent to convolution of the image in the space domain with the corresponding Bessel function. The ripples or lobes of the Bessel function lead to echoes of strong edges, an artifact known as the *ringing artifact*. The example illustrates that the “ideal” filter’s abrupt transition from the passband to the stopband is, after all, not a desirable characteristic.

The Butterworth lowpass filter: Prevention of the ringing artifacts encountered with the ideal lowpass filter requires that the transition from the passband to the stopband (and vice-versa in the case of highpass filters) be smooth. The Butterworth filter is a commonly used frequency-domain filter due to its simplicity of design and the property of a maximally flat magnitude response in the passband. For a 1D Butterworth lowpass filter of order n , the first $2n - 1$ derivatives of the squared magnitude response are zero at $\omega = 0$, where ω represents the radian frequency. The Butterworth filter response is monotonic in the passband as well as in the stopband. (See Rangayyan [31] for details and illustrations of 1D the Butterworth filter.)


In 2D, the Butterworth lowpass filter is defined as [8]

$$H(u, v) = \frac{1}{1 + (\sqrt{2} - 1) \left[\frac{D(u, v)}{D_0} \right]^{2n}}, \quad (3.42)$$

where n is the order of the filter, $D(u, v) = \sqrt{u^2 + v^2}$, and D_0 is the half-power 2D radial cutoff frequency [the scale factor in the denominator leads to the gain of the filter being $\frac{1}{\sqrt{2}}$ at $D(u, v) = D_0$]. The filter’s transition from the passband to the stopband becomes steeper as the order n is increased. Figures 3.28 (b) and 3.29 illustrate the magnitude (gain) of the Butterworth lowpass filter with the normalized cutoff $D_0 = 0.4$ and order $n = 2$.

Example: The result of filtering the noisy Shapes image in Figure 3.30 (b) with the Butterworth lowpass filter as above is shown in Figure 3.30 (d). It is seen that noise has been suppressed in the filtered image without causing the ringing artifact; however, the filter has caused some blurring of the edges of the objects in the image.

Example: Figure 3.32 (a) shows a test image of a clock. The image was contaminated by a significant amount of noise, suspected to be due to poor shielding of the video-signal cable between the camera and the digitizing frame buffer. Part (b) of the figure shows the log-magnitude spectrum of the image. The sinc components due to the horizontal and vertical lines in the image are readily seen on the axes of the spectrum. Radial concentrations of energy are also seen in the spectrum, related to the directional (or oriented) components of the image. Part (c) of the figure shows the result of application of the ideal lowpass filter with cutoff $D_0 = 0.4$. The strong ringing artifacts caused by the filter render the image useless, although some noise reduction has

**FIGURE 3.30**

(a) The Shapes test image. (b) The test image with Gaussian noise having a normalized variance of 0.01 added. (c) The result of ideal lowpass filtering the noisy image, with normalized cutoff $D_0 = 0.4$; see Figure 3.28. (d) The result of filtering with a Butterworth lowpass filter having $D_0 = 0.4$ and order $n = 2$. See also Figure 3.31.




FIGURE 3.31

The centered (folded) Fourier log-magnitude spectrum of (a) the Shapes images in Figure 3.30 (a) and (b) the noisy Shapes image in Figure 3.30 (b).


been achieved. Part (d) of the figure shows the result of filtering using a Butterworth lowpass filter, with $D_0 = 0.4$ and order $n = 2$. The noise in the image has been suppressed well without causing any artifact (except blurring or smoothing).

See Section 3.9 for illustration of the application of frequency-domain low-pass filters to an image obtained with a confocal microscope.


3.4.2 Removal of periodic artifacts

Periodic components in images give rise to impulse-like and periodic concentrations of energy in their Fourier spectra. This characteristic facilitates the removal of periodic artifacts through selective band-reject, notch, or comb filtering [31].

Example: Figure 3.33 (a) shows a part of an image of a mammographic phantom acquired with the grid (bucky) remaining fixed. The white objects simulate calcifications found in mammograms. The projections of the grid strips have suppressed the details in the phantom. Figure 3.33 (b) shows the log-magnitude Fourier spectrum of the image, where the periodic concentrations of energy along the v axis as well as along the corresponding horizontal strips are related to the grid lines. Figure 3.33 (d) shows the spectrum with selected regions corresponding to the artifactual components set to zero. Figure 3.33 (c) shows the corresponding filtered image, where the grid lines have been almost completely removed.

**FIGURE 3.32**

(a) Clock test image (101×101 pixels). (b) Log-magnitude spectrum of the image. (c) Result of the ideal lowpass filter, $D_0 = 0.4$. (d) Result of the Butterworth lowpass filter, with $D_0 = 0.4$ and order $n = 2$.

**FIGURE 3.33**

(a) Part of an image of a mammographic phantom with grid artifact; see also Figure 3.34. (b) Log-magnitude Fourier spectrum of the image in (a). (c) Filtered image. (d) Filtered version of the spectrum in (b). Phantom image courtesy of L.J. Hahn, Foothills Hospital, Calgary.

Figure 3.34 (a) shows a corresponding image of the phantom acquired with the bucky moving in the recommended manner; the image is free of the grid-line artifact. Figure 3.34 (b) shows the corresponding log-magnitude Fourier spectrum, which is also free of the artifactual components that are seen in the spectrum in Figure 3.33 (b). It should be noted that removing the artifactual components as indicated by the spectrum in Figure 3.33 (d) leads to the loss of the frequency-domain components of the desired image in the same regions, which could lead to some distortion in the filtered image.




FIGURE 3.34

(a) Part of an image of a mammographic phantom with no grid artifact; compare with the image in Figure 3.33 (a). (b) Log-magnitude Fourier spectrum of the image in (a). Phantom image courtesy of L.J. Hahn, Foothills Hospital, Calgary.

3.5 Matrix Representation of Image Processing

We have thus far expressed images as 2D arrays, and transform and image processing procedures with operations such as addition, multiplication, integration, and convolution of arrays. Such expressions define the images and the related entities on a point-by-point, pixel-by-pixel, or single-element basis. The design of optimal filters and statistical estimation procedures require the application of operators such as differentiation and statistical expectation

to expressions involving images and image processing operations. The array or single-element form of representation of images and operations does not lend easily to such procedures. It would be convenient if we could represent a whole image as a single algebraic entity, and if image processing operations could be expressed using basic algebraic operations. In this section, we shall see that matrix representation of images, convolution, filters, and transforms facilitates efficient and compact expression of image processing, optimization, and estimation procedures.

3.5.1 Matrix representation of images

A sampled image may be represented by a matrix as

$$\mathbf{f} = \{f(m, n) : m = 0, 1, 2, \dots, M - 1; n = 0, 1, 2, \dots, N - 1\}. \quad (3.43)$$

The matrix has M rows, each with N elements; the matrix has N columns. Matrix methods may then be used in the analysis of images and in the derivation of optimal filters. (See Section 2.3.3 for a discussion on array and matrix representation of images.)


In treating images as matrices and mathematical entities, it should be recognized always that images are not merely arrays of numbers: certain constraints are imposed on the image matrix due to the physical properties of the image. Some of the constraints to be noted are [9]:

- Nonnegativity and upper bound: $f_{\min} \leq f(m, n) \leq f_{\max}$, where f_{\min} and f_{\max} are the minimum and maximum values, respectively, imposed by the characteristics of the object being represented by the image as well as by the image quantization process. Usually, $f_{\min} = 0$, and the pixel values are constrained to be nonnegative.
- Finite energy: $E_f = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} f^2(m, n) \leq E_{\max}$, where E_{\max} is a finite limit on the total energy of the image.
- Smoothness: Most images that represent real-life entities cannot change their characteristics abruptly. The difference between a pixel and the average of its immediate neighbors is usually bounded by a limit S as

$$f(m, n) - \frac{1}{8} \left(\begin{array}{ccc} f(m-1, n-1) + f(m-1, n) + f(m-1, n+1) \\ + f(m, n-1) & & + f(m, n+1) \\ + f(m+1, n-1) + f(m+1, n) + f(m+1, n+1) \end{array} \right) \leq S. \quad (3.44)$$

An $M \times N$ matrix may be converted to a vector by row ordering as

$$\mathbf{f} = [\mathbf{f}^1, \mathbf{f}^2, \dots, \mathbf{f}^M]^T \quad (3.45)$$

**FIGURE 3.35**

(a) Matrix representation of a 3×3 image. Vector representation of the image in (a) by: (b) row ordering and (c) column ordering or stacking.

where $\mathbf{f}^m = [f(m, 1), f(m, 2), \dots, f(m, N)]^T$ is the m^{th} row vector. Column ordering may also be performed; Figure 3.35 illustrates the conversion of an image to a vector in schematic form.

Using the vector notation, we get the energy of the image as

$$E = \mathbf{f}^T \mathbf{f} = \mathbf{f} \cdot \mathbf{f} = \sum_{i=1}^{MN} \mathbf{f}^2(i), \quad (3.46)$$

which is the inner product or dot product of the vector with itself. The energy of the image may also be computed using the outer product as

$$E = \text{Tr}[\mathbf{f} \mathbf{f}^T], \quad (3.47)$$

where $\text{Tr}[\]$ represents the trace (the sum of the main diagonal elements) of the resulting $MN \times MN$ matrix.

If the image elements are considered to be random variables, images may be treated as samples of stochastic processes, and characterized by their statistical properties as follows:

- Mean $\bar{\mathbf{f}} = E[\mathbf{f}]$, which is an $MN \times 1$ matrix or vector.

- Covariance $\boldsymbol{\sigma} = E[(\mathbf{f} - \bar{\mathbf{f}})(\mathbf{f} - \bar{\mathbf{f}})^T] = E[\mathbf{f} \mathbf{f}^T] - \bar{\mathbf{f}} \bar{\mathbf{f}}^T$, which is an $MN \times MN$ matrix given by

$$\boldsymbol{\sigma} = \begin{bmatrix} \sigma_{11}^2 & \sigma_{12} & \cdots & \sigma_{1P} \\ \sigma_{21} & \sigma_{22}^2 & \cdots & \sigma_{2P} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{P1} & \sigma_{P2} & \cdots & \sigma_{PP}^2 \end{bmatrix}, \quad (3.48)$$

where $P = MN$ and the matrix elements are

$$\sigma_{pq} = E[\{\mathbf{f}(p) - \bar{\mathbf{f}}(p)\}\{\mathbf{f}(q) - \bar{\mathbf{f}}(q)\}], \quad p, q = 1, 2, \dots, P, \quad (3.49)$$

representing the covariance between the p^{th} and q^{th} elements of the image vector. $\boldsymbol{\sigma}$ is symmetric: $\sigma_{pq} = \sigma_{qp}$. The diagonal terms σ_{pp}^2 are the variances of the elements of the image vector.

- Autocorrelation or scatter matrix $\boldsymbol{\phi} = E[\mathbf{f} \mathbf{f}^T]$, which is an $MN \times MN$ matrix. The normalized autocorrelation coefficients are defined as $\rho_{pq} = \sigma_{pq}/(\sigma_{pp}\sigma_{qq})$. Then, $-1 \leq \rho_{pq} \leq 1$. The normalized autocorrelation matrix is given by

$$\boldsymbol{\rho} = \begin{bmatrix} 1 & \rho_{12} & \cdots & \rho_{1P} \\ \rho_{21} & 1 & \cdots & \rho_{2P} \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{P1} & \rho_{P2} & \cdots & 1 \end{bmatrix}. \quad (3.50)$$

The absolute scale of variation is retained in the diagonal standard deviation matrix:

$$\mathbf{D} = \begin{bmatrix} \sigma_{11} & 0 & \cdots & 0 \\ 0 & \sigma_{22} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma_{PP} \end{bmatrix}. \quad (3.51)$$

Then, $\boldsymbol{\sigma} = \mathbf{D} \boldsymbol{\rho} \mathbf{D}$.

Two image vectors \mathbf{f} and \mathbf{g} are

- Uncorrelated if $E[\mathbf{f} \mathbf{g}^T] = E[\mathbf{f}] E[\mathbf{g}^T]$. Then, the cross-covariance matrix $\boldsymbol{\sigma}_{fg}$ is a diagonal matrix, and the cross-correlation $\boldsymbol{\phi}_{fg} = \mathbf{I}$, the identity matrix.
- Orthogonal if $E[\mathbf{f} \mathbf{g}^T] = \mathbf{0}$.
- Statistically independent if $p(\mathbf{f}, \mathbf{g}) = p(\mathbf{f}) p(\mathbf{g})$. Then, \mathbf{f} and \mathbf{g} are uncorrelated.

The representation of image and noise processes as above facilitates the application of optimization techniques and the design of optimal filters.

3.5.2 Matrix representation of transforms

1D transforms: In signal analysis, it is often useful to represent a signal $f(t)$ over the interval t_0 to $t_0 + T$ by an expansion of the form

$$f(t) = \sum_{k=0}^{\infty} a_k \varphi_k(t), \quad (3.52)$$

where the functions $\varphi_k(t)$ are mutually orthogonal, that is,

$$\int_{t_0}^{t_0+T} \varphi_k(t) \varphi_l^*(t) dt = \begin{cases} C & \text{if } k = l \\ 0 & \text{if } k \neq l. \end{cases} \quad (3.53)$$

The functions are said to be orthonormal if $C = 1$.

The coefficients a_k may be obtained as

$$a_k = \frac{1}{C} \int_{t_0}^{t_0+T} f(t) \varphi_k^*(t) dt, \quad (3.54)$$

$k = 0, 1, 2, \dots, \infty$; that is, a_k is the projection of $f(t)$ on to $\varphi_k(t)$.

The set of functions $\{\varphi_k(t)\}$ is said to be complete or closed if there exists no square-integrable function $f(t)$ for which

$$\int_{t_0}^{t_0+T} f(t) \varphi_k^*(t) dt = 0, \quad \forall k; \quad (3.55)$$

that is, the function $f(t)$ is orthogonal to all of the members of the set $\{\varphi_k(t)\}$. If such a function exists, it should be a member of the set in order for the set to be closed or complete.

When the set $\{\varphi_k(t)\}$ is complete, it is said to be an *orthogonal basis*, and may be used for accurate representation of signals. For example, the Fourier series representation of periodic signals is based upon the use of an infinite set of sine and cosine functions of frequencies that are integral multiples (harmonics) of the fundamental frequency of the given signal. In order for such a transformation to exist, the functions $f(t)$ and $\varphi_k(t)$ must be square-integrable.

With a 1D sampled signal expressed as an $N \times 1$ vector or column matrix, we may represent transforms using an $N \times N$ matrix \mathbf{L} as

$$\mathbf{F} = \mathbf{L} \mathbf{f} \quad \text{and} \quad \mathbf{f} = \mathbf{L}^{*T} \mathbf{F}, \quad (3.56)$$

with $\mathbf{L} \mathbf{L}^{*T} = \mathbf{I}$. The matrix operations are equivalent to

$$F(k) = \sum_{n=0}^{N-1} L(k, n) f(n) \quad \text{and} \quad f(n) = \sum_{k=0}^{N-1} L^*(n, k) F(k), \quad (3.57)$$

for k or n going $0, 1, \dots, N-1$, respectively. For the DFT, we need to define the matrix \mathbf{L} with its elements given by $L(k, n) = \exp(-j \frac{2\pi}{N} kn)$. With

the notation $W_N = \exp\left(-j\frac{2\pi}{N}\right)$, we have $L(k, n) = W_N^{kn}$. Then, the following representations of the DFT in array or series format and matrix multiplication format are equivalent:

$$F(k) = \sum_{n=0}^{N-1} f(n) \exp\left(-j\frac{2\pi}{N}kn\right), \quad k = 0, 1, \dots, N-1. \quad (3.58)$$

$$\begin{bmatrix} F(0) \\ F(1) \\ F(2) \\ \vdots \\ F(N-2) \\ F(N-1) \end{bmatrix} = \begin{bmatrix} W_N^0 & W_N^0 & W_N^0 & \cdots & W_N^0 & W_N^0 \\ W_N^0 & W_N^{1 \times 1} & W_N^{1 \times 2} & \cdots & W_N^{1 \times (N-2)} & W_N^{1 \times (N-1)} \\ W_N^0 & W_N^{2 \times 1} & W_N^{2 \times 2} & \cdots & W_N^{2 \times (N-2)} & W_N^{2 \times (N-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ W_N^0 & W_N^{(N-2) \times 1} & W_N^{(N-2) \times 2} & \cdots & W_N^{(N-2) \times (N-2)} & W_N^{(N-2) \times (N-1)} \\ W_N^0 & W_N^{(N-1) \times 1} & W_N^{(N-1) \times 2} & \cdots & W_N^{(N-1) \times (N-2)} & W_N^{(N-1) \times (N-1)} \end{bmatrix} \begin{bmatrix} f(0) \\ f(1) \\ f(2) \\ \vdots \\ f(N-2) \\ f(N-1) \end{bmatrix}. \quad (3.59)$$

However, because of the periodicity of the exponential function, we have $W_N^{m(N-1)} = W_N^{-m} = W_N^{N-m}$ for any integer m ; note that $W_N^N = 1$. Thus, for a given N , there are only N distinct functions W_N^k , $k = 0, 1, 2, \dots, N-1$. Figure 3.36 illustrates the vectors (or phasors) representing W_N^k , $k = 0, 1, 2, \dots, 7$. This property of the exponential function reduces the \mathbf{W} matrix to one with only N distinct values, leading to the relationship

$$\begin{bmatrix} F(0) \\ F(1) \\ F(2) \\ \vdots \\ F(N-2) \\ F(N-1) \end{bmatrix} = \begin{bmatrix} W_N^0 & W_N^0 & W_N^0 & \cdots & W_N^0 & W_N^0 \\ W_N^0 & W_N^1 & W_N^2 & \cdots & W_N^{(N-2)} & W_N^{(N-1)} \\ W_N^0 & W_N^2 & W_N^4 & \cdots & W_N^{(N-4)} & W_N^{(N-2)} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ W_N^0 & W_N^{(N-2)} & W_N^{(N-4)} & \cdots & W_N^4 & W_N^2 \\ W_N^0 & W_N^{(N-1)} & W_N^{(N-2)} & \cdots & W_N^2 & W_N^1 \end{bmatrix} \begin{bmatrix} f(0) \\ f(1) \\ f(2) \\ \vdots \\ f(N-2) \\ f(N-1) \end{bmatrix}. \quad (3.60)$$

For $N = 8$, we get

$$\begin{bmatrix} F(0) \\ F(1) \\ F(2) \\ F(3) \\ F(4) \\ F(5) \\ F(6) \\ F(7) \end{bmatrix} = \begin{bmatrix} W^0 & W^0 & W^0 & W^0 & W^0 & W^0 & W^0 & W^0 \\ W^0 & W^1 & W^2 & W^3 & W^4 & W^5 & W^6 & W^7 \\ W^0 & W^2 & W^4 & W^6 & W^0 & W^2 & W^4 & W^6 \\ W^0 & W^3 & W^6 & W^1 & W^4 & W^7 & W^2 & W^5 \\ W^0 & W^4 & W^0 & W^4 & W^0 & W^4 & W^0 & W^4 \\ W^0 & W^5 & W^2 & W^7 & W^4 & W^1 & W^6 & W^3 \\ W^0 & W^6 & W^4 & W^2 & W^0 & W^6 & W^4 & W^2 \\ W^0 & W^7 & W^6 & W^5 & W^4 & W^3 & W^2 & W^1 \end{bmatrix} \begin{bmatrix} f(0) \\ f(1) \\ f(2) \\ f(3) \\ f(4) \\ f(5) \\ f(6) \\ f(7) \end{bmatrix}, \quad (3.61)$$

where the subscript $N = 8$ to W has been suppressed in order to show the structure of the matrix clearly.




FIGURE 3.36

Vectors (or phasors) representing the $N = 8$ roots of unity, or W_8^k , $k = 0, 1, 2, \dots, 7$, where $W_8 = \exp(-j\frac{2\pi}{8})$. Based upon a similar figure by Hall [9].

2D transforms: The transformation of an $N \times N$ image $f(m, n)$, $m = 0, 1, 2, \dots, N - 1$; $n = 0, 1, 2, \dots, N - 1$, may be expressed in a generic representation as:

$$F(k, l) = \frac{1}{N} \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} f(m, n) \varphi(m, n, k, l), \quad (3.62)$$

$$f(m, n) = \frac{1}{N} \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} F(k, l) \psi(m, n, k, l), \quad (3.63)$$

where $\varphi(m, n, k, l)$ is the forward transform kernel and $\psi(m, n, k, l)$ is the inverse transform kernel. The kernel is said to be separable if $\varphi(m, n, k, l) = \varphi_1(m, k) \varphi_2(n, l)$, and symmetric in addition if φ_1 and φ_2 are functionally equal. Then, the 2D transform may be computed in two simpler steps of 1D

row transforms followed by 1D column transforms (or vice versa) as follows:

$$F_1(m, l) = \sum_{n=0}^{N-1} f(m, n) \varphi(n, l); \quad m, l = 0, 1, \dots, N-1; \quad (3.64)$$

$$F(k, l) = \sum_{m=0}^{N-1} F_1(m, l) \varphi(m, k); \quad k, l = 0, 1, \dots, N-1. \quad (3.65)$$

In the case of the 2D Fourier transform, we have the kernel

$$\begin{aligned} \varphi(m, n, k, l) &= \exp \left[-j \frac{2\pi}{N} (mk + nl) \right] \\ &= \exp \left[-j \frac{2\pi}{N} mk \right] \exp \left[-j \frac{2\pi}{N} nl \right]. \end{aligned} \quad (3.66)$$

The kernel is separable and symmetric. The 2D DFT may be expressed as

$$\mathbf{F} = \mathbf{W} \mathbf{f} \mathbf{W}, \quad (3.67)$$

where \mathbf{f} is the $N \times N$ image matrix, and \mathbf{W} is a symmetric $N \times N$ matrix with $W_N^{km} = \exp \left[-j \frac{2\pi}{N} km \right]$; however, due to the periodicity of the W_N function, the matrix \mathbf{W} has only N distinct values, as shown in Equations 3.60 and 3.61.

The DFT matrix \mathbf{W} is symmetric, with its rows and columns being mutually orthogonal:

$$\sum_{m=0}^{N-1} W_N^{mk} W_N^{ml*} = \begin{cases} N & k = l \\ 0 & k \neq l \end{cases}. \quad (3.68)$$

Then, $\mathbf{W}^{-1} = \frac{1}{N} \mathbf{W}^*$, which leads to

$$\mathbf{f} = \frac{1}{N^2} \mathbf{W}^* \mathbf{F} \mathbf{W}. \quad (3.69)$$

A number of transforms such as the Fourier, Walsh–Hadamard, and discrete cosine may be expressed as $\mathbf{F} = \mathbf{A} \mathbf{f} \mathbf{A}$, with the matrix \mathbf{A} constructed using the relevant basis functions. The transform matrices may be decomposed into products of matrices with fewer nonzero elements, reducing redundancy and computational requirements. The DFT matrix may be factored into a product of $2 \ln N$ sparse and diagonal matrices, leading to the FFT algorithm [9, 194].

The Walsh–Hadamard Transform: The orthonormal, complete set of 1D Walsh functions defined over the interval $0 \leq x \leq 1$ is given by the iterative relationships [8, 9]

$$\varphi_n(x) = \begin{cases} \varphi_{\lfloor \frac{n}{2} \rfloor}(2x), & x < \frac{1}{2}, \\ \varphi_{\lfloor \frac{n}{2} \rfloor}(2x - 1), & x \geq \frac{1}{2}, n \text{ odd}, \\ -\varphi_{\lfloor \frac{n}{2} \rfloor}(2x - 1), & x \geq \frac{1}{2}, n \text{ even}, \end{cases} \quad (3.70)$$

where $\left[\frac{n}{2}\right]$ is the integral part of $\frac{n}{2}$,

$$\varphi_0(x) = 1, \quad (3.71)$$

and

$$\varphi_1(x) = \begin{cases} 1 & x < \frac{1}{2} \\ -1 & x \geq \frac{1}{2}. \end{cases} \quad (3.72)$$

The n^{th} function φ_n is generated by compression of the function $\varphi_{\left[\frac{n}{2}\right]}$ into its first half and $\pm\varphi_{\left[\frac{n}{2}\right]}$ into its second half. Figure 3.37 shows the first eight Walsh functions sampled with 100 samples over the interval $(0,1)$, obtained using the definition in Equation 3.70. (*Note:* The ordering of the Walsh functions varies with the formulation of the functions.)




FIGURE 3.37

The first eight Walsh functions sampled with 100 samples over the interval $(0,1)$.

Walsh functions are ordered by the number of zero-crossings in the interval $(0,1)$, called sequency. If the Walsh functions with the number of zero-crossings $\leq (2^n - 1)$ are sampled with $N = 2^n$ uniformly spaced points, we

get a square matrix representation that is orthogonal and has its rows ordered with increasing number of zero-crossings; for $N = 8$ we have

$$\mathbf{A} = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 \\ 1 & 1 & -1 & -1 & -1 & -1 & 1 & 1 \\ 1 & 1 & -1 & -1 & 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 & 1 & -1 & -1 & 1 \\ 1 & -1 & -1 & 1 & -1 & 1 & 1 & -1 \\ 1 & -1 & 1 & -1 & -1 & 1 & -1 & 1 \\ 1 & -1 & 1 & -1 & 1 & -1 & 1 & -1 \end{bmatrix}. \quad (3.73)$$

(Considering sequency ordering as above, the functions for $n = 4$ and $n = 5$ in Figure 3.37 need to be swapped.) The Walsh transform of a 1D signal \mathbf{f} may then be expressed as $\mathbf{F} = \mathbf{A} \mathbf{f}$.

Another formulation of the Walsh transform defines the forward and inverse kernel function as [8]

$$\varphi(n, k) = \frac{1}{N} \prod_{p=0}^{P-1} (-1)^{[b_p(n)b_{P-1-p}(k)]}, \quad (3.74)$$

for 1D signals, and

$$\varphi(m, n, k, l) = \frac{1}{N} \prod_{p=0}^{P-1} (-1)^{[b_p(m)b_{P-1-p}(k) + b_p(n)b_{P-1-p}(l)]}, \quad (3.75)$$

for 2D signals, where $b_p(m)$ is the p^{th} bit in the P -bit binary representation of m .

The major advantage of the Walsh transform is that the kernel has integral values of $+1$ and -1 only. As a result, the transform involves only addition and subtraction of the input image pixels. Furthermore, the operations involved in the forward and inverse transformation are identical.

Except for the ordering of rows, the discrete Walsh matrices are equivalent to the Hadamard matrices of rank 2^n , which are constructed as follows [8]:

$$\mathbf{A}_2 = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \quad (3.76)$$

$$\mathbf{A}_{2N} = \begin{bmatrix} \mathbf{A}_N & \mathbf{A}_N \\ \mathbf{A}_N & -\mathbf{A}_N \end{bmatrix}. \quad (3.77)$$

Then, by defining $\mathbf{A} = \frac{1}{\sqrt{N}} \mathbf{A}_N$, the Walsh-Hadamard transform (WHT) of a 2D function may be expressed as

$$\mathbf{F} = \mathbf{A} \mathbf{f} \mathbf{A}, \text{ and } \mathbf{f} = \mathbf{A} \mathbf{F} \mathbf{A}, \quad (3.78)$$

where all matrices are of size $N \times N$.

Figure 3.38 shows the 2D Walsh–Hadamard basis functions (matrices) for $k, l = 0, 1, 2, \dots, 7$. The functions were generated by using the 2D equivalents of the recursive definition given for the 1D Walsh functions in Equation 3.70. Note that the ordering of the functions could vary from one definition to another.




FIGURE 3.38

The first 64 Walsh–Hadamard 2D basis functions. Black represents a pixel value of -1 , and white $+1$. Each function was computed as an 8×8 matrix.

The computational and representational simplicity of the WHT is evident from the expressions above. The WHT has applications in image coding, image sequency filtering, and feature extraction for pattern recognition.

3.5.3 Matrix representation of convolution

With images represented as vectors, linear system operations (convolution and filtering) may be represented as matrix-vector multiplications [8, 9, 195]. For the sake of simplicity, let us first consider the 1D LSI system. Let us

assume the system to be causal, and to have an infinite impulse response (IIR). Then, we have the input-output relationship given by the linear convolution operation

$$g(n) = \sum_{\alpha=0}^n f(\alpha) h(n - \alpha). \quad (3.79)$$

If the input is given over N samples, we could represent the output over the interval $[0, N]$ as

$$\mathbf{g} = \mathbf{h} \mathbf{f}, \quad (3.80)$$

or, in expanded form,

$$\begin{bmatrix} g(0) \\ g(1) \\ g(2) \\ \vdots \\ g(N) \end{bmatrix} = \begin{bmatrix} h(0) & 0 & \cdots & \cdots & 0 \\ h(1) & h(0) & 0 & \cdots & 0 \\ h(2) & h(1) & h(0) & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ h(N) & h(N-1) & h(N-2) & \cdots & h(0) \end{bmatrix} \begin{bmatrix} f(0) \\ f(1) \\ f(2) \\ \vdots \\ f(N) \end{bmatrix}. \quad (3.81)$$

The matrix \mathbf{h} is a Toeplitz-like matrix, which leads to computational advantages. (*Note:* A Toeplitz matrix is a square matrix whose elements are equal along every diagonal.) There will be zeros in the lower-left portion of \mathbf{h} if $h(n)$ has fewer samples than $f(n)$ and $g(n)$: \mathbf{h} is then said to be banded.

If the impulse response of the filter is of a finite duration of $M + 1$ samples, that is, the filter is of the finite impulse response (FIR) type, it is common to use the noncausal, moving-window representation of convolution as

$$\sum_{\alpha=n-\frac{M}{2}}^{n+\frac{M}{2}} f(\alpha) h(n - \alpha), \quad (3.82)$$

or

$$g(n) = \sum_{\alpha=0}^M f\left(\alpha + n - \frac{M}{2}\right) h\left(\frac{M}{2} - \alpha\right). \quad (3.83)$$

This relationship may also be expressed as $\mathbf{g} = \mathbf{h} \mathbf{f}$, with the matrix and vectors constructed as in Figure 3.39. The matrix \mathbf{h} is now banded and Toeplitz-like. Furthermore, each row of \mathbf{h} , except the first, is a right-shifted version of the preceding row. (It has been assumed that $M < N$.)

Another representation of convolution is the periodic or circular form, where all of the signals are assumed to be of finite duration and periodic, with the period being equal to N samples. The shifting required in convolution is interpreted as periodic shifting: samples that go out of the frame of N samples at one end will reappear at the other end [7]. The periodic convolution operation is expressed as

$$g_p(n) = \sum_{\alpha=0}^{N-1} f_p(\alpha) h_p([n - \alpha] \bmod N), \quad (3.84)$$

$$\begin{bmatrix} g(0) \\ g(1) \\ \vdots \\ g(N) \end{bmatrix} = \begin{bmatrix} h(\frac{M}{2}) & h(\frac{M}{2} - 1) & \cdots & h(0) & \cdots & h(1 - \frac{M}{2}) & h(-\frac{M}{2}) & 0 & \cdots & 0 \\ 0 & h(\frac{M}{2}) & h(\frac{M}{2} - 1) & \cdots & h(0) & \cdots & h(1 - \frac{M}{2}) & h(-\frac{M}{2}) & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \cdots & \ddots & \cdots & \ddots & \ddots & \vdots \\ 0 & \cdots & h(\frac{M}{2}) & h(\frac{M}{2} - 1) & \cdots & h(0) & \cdots & \cdots & h(1 - \frac{M}{2}) & h(-\frac{M}{2}) \end{bmatrix} \begin{bmatrix} f(-\frac{M}{2}) \\ f(1 - \frac{M}{2}) \\ \vdots \\ f(N + \frac{M}{2}) \end{bmatrix}$$

FIGURE 3.39

Construction of the matrix and vectors for convolution in the case of an FIR filter.

where the subscript p indicates the periodic version or interpretation of the signals. Note that, whereas the result of periodic convolution of two periodic signals of period N samples each is another periodic signal of the same period of N samples, the result of their linear convolution would be of duration $2N - 1$ samples. However, periodic convolution may be used to achieve the same result as linear convolution by padding the signals with zeros so as to increase their duration to at least $2N - 1$ samples, and then considering the extended signals to be periodic with the period of $2N - 1$ samples. It should be observed that implementing convolution by taking the inverse DFT of the product of the DFTs of the two signals will result in periodic convolution of the signals; zero padding as above will be necessary in such a case if linear convolution is desired. The operation performed by a causal, stable, LSI system is *linear* convolution.

Periodic convolution may also be expressed as $\mathbf{g} = \mathbf{h} \mathbf{f}$, with the matrix and vectors constructed as

$$\begin{bmatrix} g(0) \\ g(1) \\ \vdots \\ g(N-2) \\ g(N-1) \end{bmatrix} = \begin{bmatrix} h(0) & h(N-1) & \cdots & h(2) & h(1) \\ h(1) & h(0) & \cdots & h(3) & h(2) \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ h(N-2) & h(N-3) & \cdots & h(0) & h(N-1) \\ h(N-1) & h(N-2) & \cdots & h(1) & h(0) \end{bmatrix} \begin{bmatrix} f(0) \\ f(1) \\ \vdots \\ f(N-2) \\ f(N-1) \end{bmatrix}. \quad (3.85)$$

(The subscript p has been dropped for the sake of concise notation.) Each row of the matrix \mathbf{h} is a right-circular shift of the previous row, and the matrix is square: such a matrix is known as a *circulant matrix*.

3.5.4 Illustrations of convolution

Convolution is an important operation in the processing of signals and images. We encounter convolution when filtering a signal or an image with an LSI system: the output is the convolution of the input with the impulse response of the system.

Convolution in 1D: In 1D, for causal signals and systems, we have the linear convolution operation given by

$$g(n) = \sum_{k=0}^n f(k) h(n-k). \quad (3.86)$$

As expressed above, the signal h needs to be reversed with respect to the index of summation k , and shifted by the interval n at which the output sample is to be computed. The index n may be run over a certain range of interest or over all time for which the result $g(n)$ exists. For each value of n , the reversed and shifted version of h is multiplied with the signal f on a point-by-point basis and summed. The multiplication and summation operation, together, are comparable to the dot product operation, performed over the nonzero overlapping parts of the two signals.

Example: Consider the signal $f(n) = [4, 1, 3, 1]$, defined for $n = 0, 1, 2, 3$. Let the signal be processed by an LSI system with the impulse response $h(n) = [3, 2, 1]$, for $n = 0, 1, 2$. The signal and system are assumed to be causal, that is, the values of $f(n)$ and $h(n)$ are zero for $n < 0$; furthermore, it is assumed that f and h are zero beyond the last sample provided.

The operation of convolution is illustrated in Figure 3.40. Observe the reversal and shifting of h . Observe also that the result g has more samples (or, is of longer duration) than either f or h : the result of linear convolution of two signals with N_1 and N_2 samples will have $N_1 + N_2 - 1$ samples (not including trailing zero-valued samples).

In the matrix notation of Equation 3.81, the convolution example above is expressed as

$$\begin{bmatrix} 3 & 0 & 0 & 0 & 1 & 2 \\ 2 & 3 & 0 & 0 & 0 & 1 \\ 1 & 2 & 3 & 0 & 0 & 0 \\ 0 & 1 & 2 & 3 & 0 & 0 \\ 0 & 0 & 1 & 2 & 3 & 0 \\ 0 & 0 & 0 & 1 & 2 & 3 \end{bmatrix} \begin{bmatrix} 4 \\ 1 \\ 3 \\ 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 12 \\ 11 \\ 15 \\ 10 \\ 5 \\ 1 \end{bmatrix}. \quad (3.87)$$

The result is identical to that shown in Figure 3.40.

Convolution in 2D: The output of an LSI imaging or image processing system is given as the convolution of the input image with the PSF:

$$g(m, n) = \sum_{\alpha=0}^{N-1} \sum_{\beta=0}^{N-1} f(\alpha, \beta) h(m - \alpha, n - \beta). \quad (3.88)$$

In this expression, for the sake of generality, the range of summation is allowed to span the full spatial range of the resulting output image. When the filter PSF is of a much smaller spatial extent than the input image, it becomes convenient to locate the origin $(0, 0)$ at the center of the PSF, and use positive and negative indices to represent the omnidirectional (and noncausal) nature of the PSF. Then, 2D convolution may be expressed as

$$g(m, n) = \sum_{\alpha=-M}^M \sum_{\beta=-M}^M f(\alpha, \beta) h(m - \alpha, n - \beta), \quad (3.89)$$

where the size of the PSF is assumed to be odd, given by $(2M + 1) \times (2M + 1)$. In this format, 2D convolution may be interpreted as a mask operation performed on the input image: the PSF is reversed (flipped or reflected) about both of its axes, placed on top of the input image at the coordinate where the output value is to be computed, a point-by-point multiplication is performed of the overlapping areas of the two functions, and the resulting products are added. The operation needs to be performed at every spatial location for which the output exists, by dragging and placing the reversed PSF at every

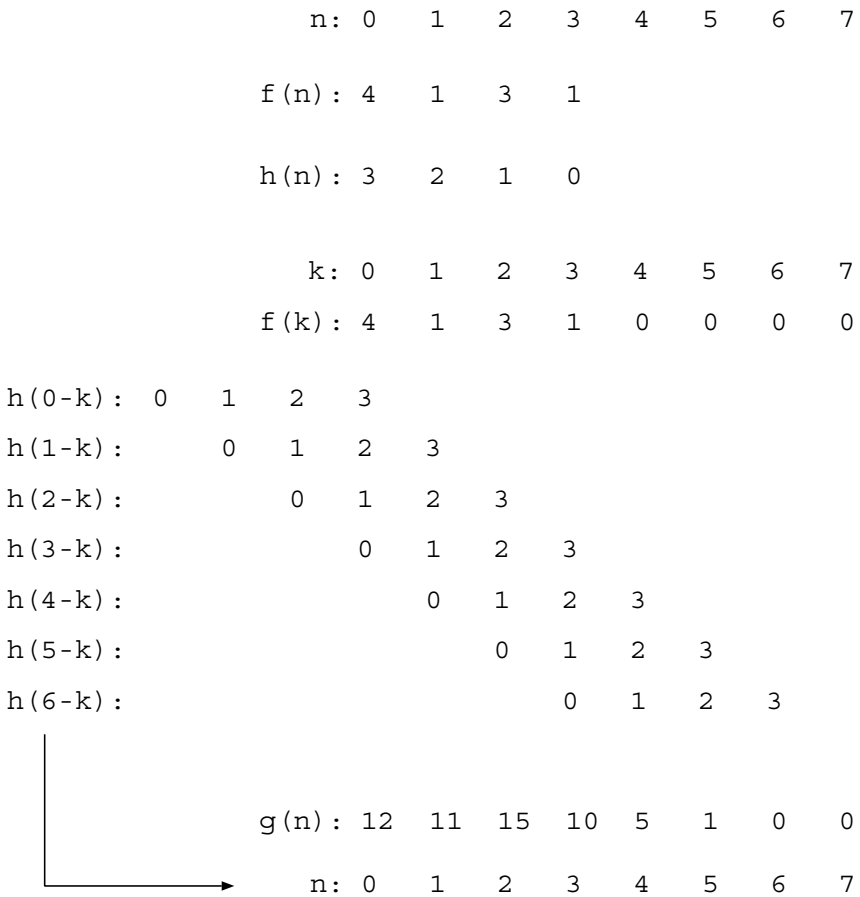


FIGURE 3.40

Illustration of the linear convolution of two 1D signals. Observe the reversal of $h(n)$, shown as $h(0 - k)$, and the shifting of the reversed signal, shown as $h(1 - k)$, $h(2 - k)$, etc.

pixel of the input image. Figure 3.41 illustrates linear 2D convolution performed as described above. Note that, in the case of a PSF having symmetry about both of its axes, the reversal step has no effect and is not required.

Matrix representation of 2D convolution is described in Section 3.5.6.

3.5.5 Diagonalization of a circulant matrix

An important property of a circulant matrix is that it is diagonalized by the DFT [8, 9, 196]. Consider the general circulant matrix

$$\mathbf{C} = \begin{bmatrix} C(0) & C(1) & C(2) & \cdots & C(N-2) & C(N-1) \\ C(N-1) & C(0) & C(1) & \cdots & C(N-3) & C(N-2) \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ C(2) & C(3) & C(4) & \cdots & C(0) & C(1) \\ C(1) & C(2) & C(3) & \cdots & C(N-1) & C(0) \end{bmatrix}. \quad (3.90)$$

Let $W = \exp(j\frac{2\pi}{N})$, which leads to $W^{kN} = 1$ for any integer k . Then, W^k , $k = 0, 1, 2, \dots, N-1$, are the N distinct roots of unity. Now, consider

$$\lambda(k) = C(0) + C(1)W^k + C(2)W^{2k} + \cdots + C(N-1)W^{(N-1)k}. \quad (3.91)$$

It follows that

$$\lambda(k)W^k = C(N-1) + C(0)W^k + C(1)W^{2k} + \cdots + C(N-2)W^{(N-1)k},$$

$$\lambda(k)W^{2k} = C(N-2) + C(N-1)W^k + C(0)W^{2k} + \cdots + C(N-3)W^{(N-1)k},$$

and so on to

$$\lambda(k)W^{(N-1)k} = C(1) + C(2)W^k + C(3)W^{2k} + \cdots + C(0)W^{(N-1)k}.$$

This series of relationships may be expressed in compact form as

$$\lambda(k) \mathbf{W}(k) = \mathbf{C} \mathbf{W}(k), \quad (3.92)$$

where

$$\mathbf{W}(k) = \left[1, W^k, W^{2k}, \dots, W^{(N-1)k} \right]^T. \quad (3.93)$$

Therefore, $\lambda(k)$ is an *eigenvalue* and $\mathbf{W}(k)$ is an *eigenvector* of the circulant matrix \mathbf{C} . Because there are N values W^k , $k = 0, 1, \dots, N-1$, that are distinct, there are N distinct eigenvectors $\mathbf{W}(k)$, which may be written as the $N \times N$ matrix

$$\mathbf{W} = [\mathbf{W}(0) \ \mathbf{W}(1) \ \cdots \ \mathbf{W}(N-1)], \quad (3.94)$$

which is related to the DFT. [The $(n, k)^{\text{th}}$ element of \mathbf{W} is $\exp(j\frac{2\pi}{N}nk)$. Due to the orthogonality of the complex exponential functions, the $(n, k)^{\text{th}}$ element

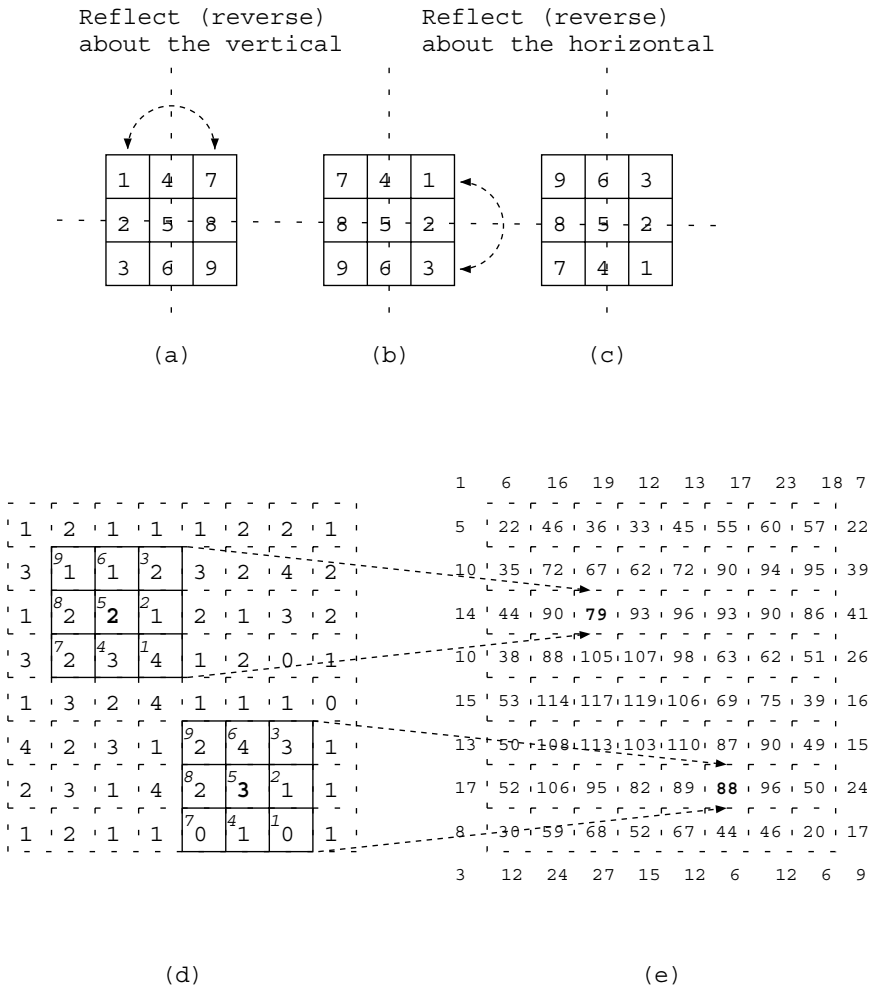


FIGURE 3.41

Illustration of the linear convolution of two 2D functions. Observe the reversal of the PSF in parts (a) – (c), and the shifting of the reversed PSF as a mask placed on the image to be filtered, in part (d). The shifted mask is shown for two pixel locations. Observe that the result needs to be written in a different array. The result, of size 10×10 and shown in part (e), has two rows and two columns more than the input image (of size 8×8).

of \mathbf{W}^{-1} is $\frac{1}{N} \exp(-j\frac{2\pi}{N}nk)$. We then have $\mathbf{W} \mathbf{W}^{-1} = \mathbf{W}^{-1} \mathbf{W} = \mathbf{I}$, where \mathbf{I} is the $N \times N$ identity matrix. The columns of \mathbf{W} are linearly independent.]

The eigenvalue relationship may be written as

$$\mathbf{W} \mathbf{\Lambda} = \mathbf{C} \mathbf{W}, \quad (3.95)$$

where all the terms are $N \times N$ matrices, and $\mathbf{\Lambda}$ is a diagonal matrix whose elements are equal to $\lambda(k)$, $k = 0, 1, \dots, N-1$. The expression above may be modified to

$$\mathbf{C} = \mathbf{W} \mathbf{\Lambda} \mathbf{W}^{-1}. \quad (3.96)$$

Thus, we see that a circulant matrix is diagonalized by the DFT operator \mathbf{W} .

Returning to the relationships of periodic convolution, because \mathbf{h} is circulant, we have

$$\mathbf{h} = \mathbf{W} \mathbf{D}_h \mathbf{W}^{-1}, \quad (3.97)$$

where \mathbf{D}_h is a diagonal matrix (corresponding to $\mathbf{\Lambda}$ in the preceding discussion). The elements of \mathbf{D}_h are given by multiplying the first row of the matrix \mathbf{h} in Equation 3.85 with W^{nk} , $n = 0, 1, 2, \dots, N-1$, as in Equation 3.91:

$$H(k) = h(0) + h(N-1)W^k + h(N-2)W^{2k} + \dots + h(1)W^{(N-1)k}, \quad (3.98)$$

which is a DFT relationship; that is, $H(k)$ is the DFT of $h(n)$. [Note: The series of h above represents $h(N-n)$, $n = 0, 1, 2, \dots, N-1$, which is equal to $h(-n)$ due to periodicity. The series of W values represents $\exp(+j\frac{2\pi}{N}nk)$, $n = 0, 1, 2, \dots, N-1$. The expression may be converted to the usual forward DFT form by substituting $-n = m$.]

It follows that the result of the convolution operation is given by

$$\mathbf{g} = \mathbf{W} \mathbf{D}_h \mathbf{W}^{-1} \mathbf{f}. \quad (3.99)$$

The interpretation of the matrix relationships above is as follows: $\mathbf{W}^{-1} \mathbf{f}$ is the (forward) DFT of \mathbf{f} (with a scale factor of $\frac{1}{N}$). The multiplication of this expression by \mathbf{D}_h corresponds to point-by-point transform-domain filtering with the DFT of \mathbf{h} . The multiplication by \mathbf{W} corresponds to the inverse DFT (except for the scale factor $\frac{1}{N}$). We now have the following equivalent relationships that represent convolution:

$$\begin{aligned} g(n) &= h(n) * f(n) \\ G(k) &= H(k) F(k) \\ \mathbf{g} &= \mathbf{h} \mathbf{f} \\ \mathbf{g} &= \mathbf{W} \mathbf{D}_h \mathbf{W}^{-1} \mathbf{f}. \end{aligned} \quad (3.100)$$

Note: The representation of the Fourier transform operator above is different from that in Equation 3.67.

3.5.6 Block-circulant matrix representation of a 2D filter

In the preceding sections, we saw how a 1D LSI filter may be represented by the matrix relationship $\mathbf{g} = \mathbf{h} \mathbf{f}$, with the special case of periodic convolution being represented as above with \mathbf{h} being a circulant matrix. Now, let us consider the matrix representation of 2D filters. Let \mathbf{f} represent an original image or the input image to a 2D filter, and let \mathbf{g} represent the corresponding filtered image, with the 2D arrays having been converted to vectors or column matrices by row ordering (see Figure 3.35). Let us assume that all of the images have been padded with zeros and extended to $M \times N$ arrays, with M and N being large such that circular convolution yields a result equivalent to that of linear convolution. The images may be considered to be periodic with the period $M \times N$. The matrices \mathbf{f} and \mathbf{g} are of size $MN \times 1$.

The 2D periodic convolution expression in array form is given by

$$g(m, n) = \sum_{\alpha=0}^{M-1} \sum_{\beta=0}^{N-1} f(\alpha, \beta) h([m - \alpha] \bmod M, [n - \beta] \bmod N), \quad (3.101)$$

for $m = 0, 1, 2, \dots, M-1$, and $n = 0, 1, 2, \dots, N-1$. The result is also periodic with the period $M \times N$.

In order for the expression $\mathbf{g} = \mathbf{h} \mathbf{f}$ to represent 2D convolution, we need to construct the matrix \mathbf{h} as follows [8, 9]:

$$\mathbf{h} = \begin{bmatrix} \mathbf{h}_0 & \mathbf{h}_{M-1} & \mathbf{h}_{M-2} & \cdots & \mathbf{h}_2 & \mathbf{h}_1 \\ \mathbf{h}_1 & \mathbf{h}_0 & \mathbf{h}_{M-1} & \cdots & \mathbf{h}_3 & \mathbf{h}_2 \\ \mathbf{h}_2 & \mathbf{h}_1 & \mathbf{h}_0 & \cdots & \mathbf{h}_4 & \mathbf{h}_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{h}_{M-1} & \mathbf{h}_{M-2} & \mathbf{h}_{M-3} & \cdots & \mathbf{h}_1 & \mathbf{h}_0 \end{bmatrix}, \quad (3.102)$$

where the submatrices are given by

$$\mathbf{h}_m = \begin{bmatrix} h(m, 0) & h(m, N-1) & h(m, N-2) & \cdots & h(m, 1) \\ h(m, 1) & h(m, 0) & h(m, N-1) & \cdots & h(m, 2) \\ h(m, 2) & h(m, 1) & h(m, 0) & \cdots & h(m, 3) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ h(m, N-1) & h(m, N-2) & h(m, N-3) & \cdots & h(m, 0) \end{bmatrix}. \quad (3.103)$$

The matrix \mathbf{h} is of size $MN \times MN$. Each $N \times N$ submatrix \mathbf{h}_m is a circulant matrix. The submatrices of \mathbf{h} are subscripted in a circular manner: \mathbf{h} is known as a *block-circulant* matrix.

Example: Let us consider filtering the image $f(m, n)$ given by

$$f(m, n) = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 2 & 3 & 0 \\ 0 & 6 & 5 & 4 & 0 \\ 0 & 7 & 8 & 9 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Although the image has nonzero pixels over only a 3×3 region, it has been padded with zeros to the extent of a 5×5 array to allow for the result of convolution to be larger without wrap-around errors due to periodic convolution.

The 3×3 subtracting Laplacian operator, also extended to a 5×5 array, is given by

$$h(m, n) = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & -4 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

However, this form of the operator has its origin at the center of the array, whereas the origin of the image $f(m, n)$ as above would be at the top-left corner in matrix-indexing order. Therefore, we need to rewrite $h(m, n)$ as follows:

$$h(m, n) = \begin{bmatrix} -4 & 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Observe that due to the assumption of periodicity, the values of $h(m, n)$ corresponding to negative indices now appear on the opposite ends of the matrix.

The matrices corresponding to the relationship $\mathbf{g} = \mathbf{h} \mathbf{f}$ are given in Figure 3.42. The resulting image $g(m, n)$ in array format is

$$g(m, n) = \begin{bmatrix} 0 & 1 & 2 & 3 & 0 \\ 1 & 4 & 1 & -6 & 3 \\ 6 & -11 & 0 & 1 & 4 \\ 7 & -14 & -11 & -24 & 9 \\ 0 & 7 & 8 & 9 & 0 \end{bmatrix}.$$

Diagonalization of a block-circulant matrix: Let us define the following functions that are related to the 2D DFT [8]: $w_M(k, m) = \exp[j\frac{2\pi}{M}km]$, and $w_N(l, n) = \exp[j\frac{2\pi}{N}ln]$. Let us define a matrix \mathbf{W} of size $MN \times MN$, containing M^2 partitions each of size $N \times N$. The $(k, m)^{\text{th}}$ partition of \mathbf{W} is

$$\mathbf{W}(k, m) = w_M(k, m) \mathbf{W}_N, \quad (3.104)$$

for $k, m = 0, 1, 2, \dots, M-1$, where \mathbf{W}_N is an $N \times N$ matrix with its elements given by $w_N(l, n)$ for $l, n = 0, 1, 2, \dots, N-1$.

Now, \mathbf{W}^{-1} is also a matrix of size $MN \times MN$, with M^2 partitions of size $N \times N$. The $(k, m)^{\text{th}}$ partition of \mathbf{W}^{-1} is

$$\mathbf{W}^{-1}(k, m) = \frac{1}{M} w_M^{-1}(k, m) \mathbf{W}_N^{-1}, \quad (3.105)$$

where $w_M^{-1}(k, m) = \exp[-j\frac{2\pi}{M}km]$, for $k, m = 0, 1, 2, \dots, M-1$. The matrix \mathbf{W}_N^{-1} has its elements given by $\frac{1}{N}w_N^{-1}(l, n)$, where $w_N^{-1}(l, n) = \exp[-j\frac{2\pi}{N}ln]$, for $l, n = 0, 1, 2, \dots, N-1$. The definitions above lead to $\mathbf{W}\mathbf{W}^{-1} = \mathbf{W}^{-1}\mathbf{W} = \mathbf{I}$, where \mathbf{I} is the $MN \times MN$ identity matrix. If \mathbf{h} is a block-circulant matrix, it can be shown [196] that $\mathbf{h} = \mathbf{W}\mathbf{D}_h\mathbf{W}^{-1}$ or $\mathbf{D}_h = \mathbf{W}^{-1}\mathbf{h}\mathbf{W}$, where \mathbf{D}_h is a diagonal matrix whose elements are related to the DFT of $h(m, n)$, that is, to $H(k, l)$.

Similar to the 1D case expressed by the relationships in Equation 3.100, we have the following equivalent relationships that represent 2D convolution:

$$\begin{aligned} g(m, n) &= h(m, n) * f(m, n) \\ G(k, l) &= H(k, l) F(k, l) \\ \mathbf{g} &= \mathbf{h} \mathbf{f} \\ \mathbf{g} &= \mathbf{W} \mathbf{D}_h \mathbf{W}^{-1} \mathbf{f}. \end{aligned} \quad (3.106)$$

Note: Considering an $N \times N$ image $f(m, n)$, the $N^2 \times N^2$ DFT matrix \mathbf{W} above is different from the $N \times N$ DFT matrix \mathbf{W} in Equation 3.67. The image matrix \mathbf{f} in Equation 3.67 is of size $N \times N$, whereas the image is represented as an $N^2 \times 1$ vector in Equation 3.106.

Differentiation of functions of matrices: The major advantage of expressing images and image processing operations in matrix form as above is that mathematical procedures for optimization and estimation may be applied with ease. For example, we have the following derivatives: given the vectors \mathbf{f} and \mathbf{g} and a symmetric matrix \mathbf{W} ,

$$\frac{\partial}{\partial \mathbf{f}}(\mathbf{f}^T \mathbf{g}) = \frac{\partial}{\partial \mathbf{f}}(\mathbf{g}^T \mathbf{f}) = \mathbf{g}, \quad (3.107)$$

and

$$\frac{\partial}{\partial \mathbf{f}}(\mathbf{f}^T \mathbf{W} \mathbf{f}) = 2 \mathbf{W} \mathbf{f}. \quad (3.108)$$

Several derivations in Sections 3.6.1 and 3.7.1, as well as in Chapters 10 and 11, demonstrate how optimization of filters may be performed using matrix representation of images and image processing operations as above.

3.6 Optimal Filtering

The field of image processing includes several procedures that may be characterized as *ad hoc* methods: procedures that have been designed to address

a particular problem and observed to yield good results in certain specific applications or scenarios. Often, the conditions under which such methods perform well are not known or understood, and the application of the methods to other images or situations may not lead to useful results. Certain mathematical models and procedures permit the design of optimal filters — filters that are derived through an optimization procedure that minimizes a cost function under specific conditions. Such procedures state explicitly the necessary conditions, and the behavior of the filter is predictable. However, as we shall see in the following paragraphs, the application of optimization methods and the resultant optimal filters require specific knowledge of the image and noise processes.

3.6.1 The Wiener filter

The Wiener filter is a linear filter designed to minimize the MSE between the output of the filter and the undegraded, unknown, original image. The filter output is an optimal estimate of the original, undegraded image in the MSE sense, and hence is known as the linear minimum mean squared-error (LMMSE) or the least-mean-square (LMS) estimate.

Considering the degradation of an image \mathbf{f} by additive noise $\boldsymbol{\eta}$ that is independent of the image process, we have the degraded image given by

$$\mathbf{g} = \mathbf{f} + \boldsymbol{\eta}. \quad (3.109)$$

(Degradation including the PSF matrix \mathbf{h} is considered in Chapter 10.)

The Wiener estimation problem may be stated as follows [8, 9, 197, 198]: determine a linear estimate $\tilde{\mathbf{f}} = \mathbf{L}\mathbf{g}$ of \mathbf{f} from the given image \mathbf{g} , where \mathbf{L} is the linear filter or transform operator to be designed. Recall that \mathbf{f} and \mathbf{g} are $N^2 \times 1$ matrices formed by row or column ordering of the corresponding $N \times N$ images, and that \mathbf{L} is an $N^2 \times N^2$ matrix.

The optimization criterion used to design the Wiener filter is to minimize the MSE, given by

$$\varepsilon^2 = E \left[\|\mathbf{f} - \tilde{\mathbf{f}}\|^2 \right]. \quad (3.110)$$

Let us express the MSE as the trace of the outer product matrix of the error vector:

$$\varepsilon^2 = E \left[\text{Tr} \left\{ (\mathbf{f} - \tilde{\mathbf{f}})(\mathbf{f} - \tilde{\mathbf{f}})^T \right\} \right]. \quad (3.111)$$

We have the following expressions that result from, or are related to, the above:

$$(\mathbf{f} - \tilde{\mathbf{f}})(\mathbf{f} - \tilde{\mathbf{f}})^T = \mathbf{f}\mathbf{f}^T - \mathbf{f}\tilde{\mathbf{f}}^T - \tilde{\mathbf{f}}\mathbf{f}^T + \tilde{\mathbf{f}}\tilde{\mathbf{f}}^T; \quad (3.112)$$

$$\tilde{\mathbf{f}}^T = \mathbf{g}^T \mathbf{L}^T = (\mathbf{f}^T + \boldsymbol{\eta}^T) \mathbf{L}^T; \quad (3.113)$$

$$\mathbf{f}\tilde{\mathbf{f}}^T = \mathbf{f}\mathbf{f}^T \mathbf{L}^T + \mathbf{f}\boldsymbol{\eta}^T \mathbf{L}^T; \quad (3.114)$$

$$\tilde{\mathbf{f}}\mathbf{f}^T = \mathbf{L}\mathbf{f}\mathbf{f}^T + \mathbf{L}\boldsymbol{\eta}\mathbf{f}^T; \quad (3.115)$$

$$\tilde{\mathbf{f}} \tilde{\mathbf{f}}^T = \mathbf{L} (\mathbf{f} \mathbf{f}^T + \mathbf{f} \boldsymbol{\eta}^T + \boldsymbol{\eta} \mathbf{f}^T + \boldsymbol{\eta} \boldsymbol{\eta}^T) \mathbf{L}^T. \quad (3.116)$$

Because the trace of a sum of matrices is equal to the sum of their traces, the E and Tr operators may be interchanged in order. Applying the $E[\]$ operator to the expressions above, we get the following expressions:

- $E[\mathbf{f} \mathbf{f}^T] = \boldsymbol{\phi}_f$, which is the autocorrelation matrix of the image;
- $E[\mathbf{f} \tilde{\mathbf{f}}^T] = \boldsymbol{\phi}_f \mathbf{L}^T$;
- $E[\mathbf{f} \boldsymbol{\eta}^T] = 0$, because \mathbf{f} and $\boldsymbol{\eta}$ are assumed to be statistically independent;
- $E[\tilde{\mathbf{f}} \mathbf{f}^T] = \mathbf{L} \boldsymbol{\phi}_f$;
- $E[\tilde{\mathbf{f}} \tilde{\mathbf{f}}^T] = \mathbf{L} \boldsymbol{\phi}_f \mathbf{L}^T + \mathbf{L} \boldsymbol{\phi}_\eta \mathbf{L}^T$;
- $E[\boldsymbol{\eta} \boldsymbol{\eta}^T] = \boldsymbol{\phi}_\eta$, which is the noise autocorrelation matrix.

Now, the MSE may be written as

$$\begin{aligned} \varepsilon^2 &= Tr [\boldsymbol{\phi}_f - \boldsymbol{\phi}_f \mathbf{L}^T - \mathbf{L} \boldsymbol{\phi}_f + \mathbf{L} \boldsymbol{\phi}_f \mathbf{L}^T + \mathbf{L} \boldsymbol{\phi}_\eta \mathbf{L}^T] \\ &= Tr [\boldsymbol{\phi}_f - 2 \boldsymbol{\phi}_f \mathbf{L}^T + \mathbf{L} \boldsymbol{\phi}_f \mathbf{L}^T + \mathbf{L} \boldsymbol{\phi}_\eta \mathbf{L}^T]. \end{aligned} \quad (3.117)$$

(Note: $Tr [\boldsymbol{\phi}_f \mathbf{L}^T] = Tr [\mathbf{L} \boldsymbol{\phi}_f]$ because $\boldsymbol{\phi}_f$ is symmetric.) At this point, the MSE is no longer a function of the images \mathbf{f} , \mathbf{g} , or $\boldsymbol{\eta}$, but depends only on the statistical characteristics of \mathbf{f} and $\boldsymbol{\eta}$, and on \mathbf{L} .

To obtain the optimal filter operator \mathbf{L} , we may now differentiate the expression above with respect to \mathbf{L} , equate it to zero, and solve the resulting expression as follows:

$$\frac{\partial \varepsilon^2}{\partial \mathbf{L}} = -2 \boldsymbol{\phi}_f + 2 \mathbf{L} \boldsymbol{\phi}_f + 2 \mathbf{L} \boldsymbol{\phi}_\eta = 0. \quad (3.118)$$

The optimal filter is given by

$$\mathbf{L}_{\text{Wiener}} = \boldsymbol{\phi}_f (\boldsymbol{\phi}_f + \boldsymbol{\phi}_\eta)^{-1}. \quad (3.119)$$

The filtered image is given by

$$\tilde{\mathbf{f}} = \boldsymbol{\phi}_f (\boldsymbol{\phi}_f + \boldsymbol{\phi}_\eta)^{-1} \mathbf{g}. \quad (3.120)$$

Implementation of the Wiener filter: Consider the matrix $\boldsymbol{\phi}_f + \boldsymbol{\phi}_\eta$ that needs to be inverted in Equation 3.119. The matrix would be of size $N^2 \times N^2$ for $N \times N$ images; hence, inversion of the matrix as such would be impractical when N is large. Inversion becomes easier if the matrix can be written as the product of a diagonal matrix and a unitary matrix.

Now, $\boldsymbol{\phi}_\eta$ is a diagonal matrix if η is an uncorrelated random (white) noise process. In most real images, correlation between pixels reduces as the spatial

shift or distance between the pixel positions considered increases: ϕ_f is then banded with several zeros, and may be approximated by a block-circulant matrix. Then, we can write $\phi_f = \mathbf{W} \Phi_f \mathbf{W}^{-1}$, and $\phi_\eta = \mathbf{W} \Phi_\eta \mathbf{W}^{-1}$, where Φ represents the diagonal matrix corresponding to ϕ resulting from the Fourier transform operation by \mathbf{W} (see Section 3.5.5). This step leads to the Wiener filter output being expressed as

$$\tilde{\mathbf{f}} = \mathbf{W} \Phi_f (\Phi_f + \Phi_\eta)^{-1} \mathbf{W}^{-1} \mathbf{g}. \quad (3.121)$$

The following interpretation of the entities involved in the Wiener filter reduces the expression above to a more familiar form:

- $\mathbf{W}^{-1} \mathbf{g}$ is equivalent to $G(k, l)$, the Fourier transform of $g(m, n)$.
- $\Phi_f = \mathbf{W}^{-1} \phi_f \mathbf{W}$ is equivalent to $S_f(k, l)$, the PSD of $f(m, n)$.
- $\Phi_\eta = \mathbf{W}^{-1} \phi_\eta \mathbf{W}$ is equivalent to $S_\eta(k, l)$, the noise PSD.

Then, we get the Wiener estimate in the Fourier domain as

$$\begin{aligned} \tilde{F}(k, l) &= \left[\frac{S_f(k, l)}{S_f(k, l) + S_\eta(k, l)} \right] G(k, l) \\ &= \left[\frac{1}{1 + \frac{S_\eta(k, l)}{S_f(k, l)}} \right] G(k, l). \end{aligned} \quad (3.122)$$

(For other derivations of the Wiener filter, see Wiener [198], Lim [199], and Rangayyan [31].)

Observe that the Wiener filter transfer function depends upon the PSD of the original signal and noise processes; the dependence is upon the second-order statistics of the processes rather than upon single realizations or observations of the processes. The design of the Wiener filter as above requires the estimation of the PSDs or the development of models thereof. Although the original signal itself is unknown, it is often possible in practice to estimate its PSD from the average spectra of several artifact-free observations of images of the same class or type. Noise statistics, such as variance, may be estimated from signal-free parts of noisy observations of the image.

The gain of the Wiener filter varies from one frequency sample to another in accordance with the SNR [expressed as a function of frequency in the 2D (u, v) or (k, l) space]: the gain is high wherever the signal component $S_f(k, l)$ is strong as compared to the noise component $S_\eta(k, l)$, that is, wherever the SNR is high; the gain is low wherever the SNR is low. The gain is equal to unity if the noise PSD is zero. However, it should be noted that the Wiener filter (as above) is not spatially adaptive: its characteristics remain the same for the entire image. For this reason, while suppressing noise, the Wiener filter is likely to blur sharp features and edges that may exist in the image (and share the high-frequency spectral regions with noise).

Frequency-domain implementation of the Wiener filter as in Equation 3.122 obviates the need for the inversion of large correlation matrices as in Equation 3.120. The use of the FFT algorithm facilitates fast computation of the Fourier transforms of the image data.

Example: Figure 3.43 (a) shows the original Shapes test image; part (b) shows the test image with Gaussian-distributed noise added ($\mu = 0$, normalized $\sigma^2 = 0.01$). The log-magnitude spectrum of the noisy image is shown in part (c) of the figure. In order to implement the Wiener filter as in Equation 3.122, the true image PSD was modeled using a Laplacian function with $\sigma = 5$ pixels in the Fourier domain, represented using a 128×128 array. The noise PSD was modeled by a uniform function having its total energy (area under the function) equal to 0.5 times that of the Laplacian PSD model. The Wiener filter transfer function is illustrated in Figure 3.43 (d). The output of the Wiener filter is shown in part (e) of the figure: while the noise in the uniform areas of the image has been suppressed, the edges of the objects in the image have been severely blurred. The blurring of edges has resulted in an increase of the RMS error from 19.56 for the noisy image to 52.89 for the Wiener filter output. It is clear that the assumption of stationarity is not appropriate for the test image; the use of a fixed filter, albeit optimal in the MSE sense, has led to a result that is not desirable. Furthermore, the design of appropriate signal and noise PSD models is difficult in practice; inappropriate models could lead to poor performance, as illustrated by the present example.


The implicit assumption made in deriving the Wiener filter is that the noise and image processes are second-order stationary processes; that is, their mean and variance do not vary from one image region to another. This also leads to the assumption that the entire image may be characterized by a single frequency spectrum or PSD. Most real-life images do not satisfy these assumptions to the fullest extent, which calls for the design of spatially or locally adaptive filters. See Section 3.7.1 for details on locally adaptive optimal filters.

3.7 Adaptive Filters

3.7.1 The local LMMSE filter

Lee [200] developed a class of adaptive, local-statistics-based filters to obtain the LMMSE estimate of the original image from a degraded version. The degradation model represents an image $f(m, n)$ corrupted by additive noise $\eta(m, n)$ as

$$g(m, n) = f(m, n) + \eta(m, n), \quad \forall m, n. \quad (3.123)$$

**FIGURE 3.43**

(a) Shapes test image. (b) Image in (a) with Gaussian-distributed noise added, with $\mu = 0$, normalized $\sigma^2 = 0.01$; RMS error = 19.56. (c) Log-magnitude spectrum of the image in (b). (d) Gain of the Wiener filter in the frequency domain, that is, the magnitude transfer function. Result of filtering the noisy image in (b) using: (e) the Wiener filter as in (d), RMS error = 52.89; (f) the local LMMSE filter with a 5×5 window, RMS error = 13.78.

This model is equivalent to that in Equation 3.109, but has been shown in the 2D array form with the indices (m, n) to indicate the pixel location in order to demonstrate the locally adaptive nature of the filter to be derived. The original image f is considered to be a realization of a *nonstationary random field*, characterized by spatially varying moments (mean, standard deviation, etc.). The noise process η may be either signal-independent or signal-dependent, and could be nonstationary as well.

The LMMSE approach computes at *every spatial location* (m, n) an estimate $\tilde{f}(m, n)$ of the original image value $f(m, n)$ by applying a *linear operator* to the available corrupted image value $g(m, n)$. Scalars $a(m, n)$ and $b(m, n)$ are sought such that the value $f(m, n)$ computed as

$$\tilde{f}(m, n) = a(m, n) g(m, n) + b(m, n) \quad (3.124)$$

minimizes the local MSE

$$\varepsilon^2(m, n) = \overline{[\tilde{f}(m, n) - f(m, n)]^2}, \quad (3.125)$$

where the bar above the expression indicates some form of averaging (statistical expectation, ensemble averaging, or spatial averaging). We have the local MSE in expanded form as

$$\varepsilon^2(m, n) = \overline{[a(m, n) g(m, n) + b(m, n) - f(m, n)]^2}. \quad (3.126)$$

The values $a(m, n)$ and $b(m, n)$ that minimize $\varepsilon^2(m, n)$ are computed by taking the partial derivatives of $\varepsilon^2(m, n)$ with respect to $a(m, n)$ and $b(m, n)$, setting them to zero, and solving the resulting equation, as follows:

$$\frac{\partial \varepsilon^2(m, n)}{\partial b(m, n)} = \overline{2\{a(m, n) g(m, n) + b(m, n) - f(m, n)\}} = 0, \quad (3.127)$$

which leads to

$$b(m, n) = \overline{f(m, n)} - a(m, n) \overline{g(m, n)}. \quad (3.128)$$

Now, replacing $b(m, n)$ in Equation 3.126 with its value given by Equation 3.128, we get the local MSE as

$$\varepsilon^2(m, n) = \overline{[a(m, n)\{g(m, n) - \overline{g(m, n)}\} - \{f(m, n) - \overline{f(m, n)}\}]^2}. \quad (3.129)$$

Differentiating this expression with respect to $a(m, n)$ and setting the result to zero, we get:

$$\overline{[a(m, n)\{g(m, n) - \overline{g(m, n)}\} - \{f(m, n) - \overline{f(m, n)}\}]\{g(m, n) - \overline{g(m, n)}\}} = 0. \quad (3.130)$$

Now, we may allow $\overline{[g(m, n) - \overline{g(m, n)}]^2} = \sigma_g^2(m, n)$ to represent the local variance of g , and $\overline{[g(m, n) - \overline{g(m, n)}][f(m, n) - \overline{f(m, n)}]} = \sigma_{fg}(m, n)$, the local covariance between f and g . This leads to

$$a(m, n) = \frac{\sigma_{fg}(m, n)}{\sigma_g^2(m, n)}. \quad (3.131)$$

With $a(m, n)$ given by Equation 3.131 and $b(m, n)$ given by Equation 3.128, the LMMSE estimate formula in Equation 3.124 becomes

$$\tilde{f}(m, n) = \bar{f}(m, n) + \frac{\sigma_{fg}(m, n)}{\sigma_g^2(m, n)} [g(m, n) - \bar{g}(m, n)]. \quad (3.132)$$

Because the true statistics of both the original and the corrupted image as well as their joint statistics are usually unknown in a practical situation, Lee proposed to estimate them locally in a spatial neighborhood of the pixel (m, n) being processed, leading to the *local* LMMSE (that is, the LLMMSE) estimate. Using a rectangular window of size $(2P + 1) \times (2Q + 1)$ centered at the pixel (m, n) being processed, we get local estimates of the mean and variance of the noisy image g as

$$\mu_g(m, n) = \frac{1}{(2P + 1)(2Q + 1)} \sum_{p=-P}^P \sum_{q=-Q}^Q g(m + p, n + q), \quad (3.133)$$

and

$$\sigma_g^2(m, n) = \frac{1}{(2P + 1)(2Q + 1)} \sum_{p=-P}^P \sum_{q=-Q}^Q [g(m + p, n + q) - \mu_g(m, n)]^2. \quad (3.134)$$

It should be noted that the parameters are expressed as functions of space and are space-variant entities. The LLMMSE estimate is then approximated by the following pixel-by-pixel operation:

$$\tilde{f}(m, n) = \mu_g(m, n) + \left[\frac{\sigma_g^2(m, n) - \sigma_\eta^2(m, n)}{\sigma_g^2(m, n)} \right] [g(m, n) - \mu_g(m, n)]. \quad (3.135)$$

(For other derivations of the LLMMSE filter, also known as the Wiener filter, see Lim [199].)

Comparing Equation 3.132 with 3.135, observe that $\bar{f}(m, n)$ is approximated by $\mu_g(m, n)$, and that $\sigma_{fg}(m, n)$ is estimated by the difference between the local variance of the degraded image and that of the noise process. The LLMMSE filter is rendered spatially adaptive — and nonlinear — by the space-variant estimation of the statistical parameters used.

In deriving Equation 3.135 from Equation 3.132, the assumption that the noise is uncorrelated with the image is taken into account. The variance of the noise $\sigma_\eta^2(m, n)$ is constant over the whole image if the noise is assumed to be signal-independent, but varies if the noise is signal-dependent; in the latter case $\sigma_\eta^2(m, n)$ should be estimated locally with a knowledge of the type of the noise that corrupts the image. Lee's filter was originally derived to deal with signal-independent additive noise and signal-dependent multiplicative noise, but may be adapted to other types of signal-dependent noise, such as Poisson or film-grain noise [201].

The interpretation of Equation 3.135 is as follows: if the processing window overlaps a uniform region in which any variation is due mostly to the noise, the second term will be small. Thus, the LLMMSE estimate is equal to the local mean of the noisy image; the noise is thereby reduced. If, on the contrary, there is an edge in the processing window, the variance of the noisy image is larger than the variance of the noise. The LLMMSE estimate in this case is closer to the actual noisy value $g(m, n)$, and the edge does not get blurred. The filter provides good noise attenuation over uniform areas, but poor noise filtering near edges.

Aghdasi et al. [202] proposed detailed models of degradation of mammograms, including Poisson noise, film-grain noise, and blurring by several components along the chain of image acquisition systems. They applied the local LMMSE filter as above to remove noise, and compared its performance with a Bayesian filter. The parametric Wiener filter (see Section 10.1.3) was then applied to deblur the noise-free mammograms.

Matrix representation: A matrix or vectorial version of Equation 3.132 may also be derived as follows: The LMMSE estimate is expressed as

$$\tilde{\mathbf{f}} = \mathbf{A}\mathbf{g} + \mathbf{b}. \quad (3.136)$$

The MSE between the estimate and the unknown original image may be expressed as

$$\varepsilon^2 = E \left[\text{Tr} \left\{ (\mathbf{f} - \tilde{\mathbf{f}})(\mathbf{f} - \tilde{\mathbf{f}})^T \right\} \right]. \quad (3.137)$$

Substituting the expression in Equation 3.136, we get

$$\begin{aligned} \varepsilon^2 &= E \left[\text{Tr} \left\{ (\mathbf{f} - \mathbf{A}\mathbf{g} - \mathbf{b})(\mathbf{f} - \mathbf{A}\mathbf{g} - \mathbf{b})^T \right\} \right] \\ &= E \left[\text{Tr} \left\{ \mathbf{f}\mathbf{f}^T - \mathbf{f}\mathbf{g}^T\mathbf{A}^T - \mathbf{f}\mathbf{b}^T - \mathbf{A}\mathbf{g}\mathbf{f}^T + \mathbf{A}\mathbf{g}\mathbf{g}^T\mathbf{A}^T + \mathbf{A}\mathbf{g}\mathbf{b}^T \right. \right. \\ &\quad \left. \left. - \mathbf{b}\mathbf{f}^T + \mathbf{b}\mathbf{g}^T\mathbf{A}^T + \mathbf{b}\mathbf{b}^T \right\} \right]. \end{aligned} \quad (3.138)$$

Differentiating the expression above with respect to \mathbf{b} and setting it to zero, we get

$$E \left[\text{Tr} \left\{ -\mathbf{f} + \mathbf{A}\mathbf{g} - \mathbf{f} + \mathbf{A}\mathbf{g} + 2\mathbf{b} \right\} \right] = 0, \quad (3.139)$$

solving which we get

$$\mathbf{b} = \bar{\mathbf{f}} - \mathbf{A}\bar{\mathbf{g}}, \quad (3.140)$$

where $\bar{\cdot}$ indicates some form of averaging.

Using the expression derived for \mathbf{b} above, we get the following:

$$\begin{aligned} \mathbf{f} - \tilde{\mathbf{f}} &= \mathbf{f} - \mathbf{A}\mathbf{g} - \mathbf{b} \\ &= \mathbf{f} - \mathbf{A}\mathbf{g} - \bar{\mathbf{f}} + \mathbf{A}\bar{\mathbf{g}} \\ &= (\mathbf{f} - \bar{\mathbf{f}}) - \mathbf{A}(\mathbf{g} - \bar{\mathbf{g}}) \\ &= \mathbf{f}_1 - \mathbf{A}\mathbf{g}_1, \end{aligned} \quad (3.141)$$

where $\mathbf{f}_1 = \mathbf{f} - \bar{\mathbf{f}}$ and $\mathbf{g}_1 = \mathbf{g} - \bar{\mathbf{g}}$ for the sake of compactness in further derivation. Now, the MSE becomes

$$\begin{aligned}\varepsilon^2 &= E [Tr \{(\mathbf{f}_1 - \mathbf{A}\mathbf{g}_1)(\mathbf{f}_1 - \mathbf{A}\mathbf{g}_1)^T\}] \\ &= E [Tr \{\mathbf{f}_1\mathbf{f}_1^T - \mathbf{f}_1\mathbf{g}_1^T\mathbf{A}^T - \mathbf{A}\mathbf{g}_1\mathbf{f}_1^T + \mathbf{A}\mathbf{g}_1\mathbf{g}_1^T\mathbf{A}^T\}].\end{aligned}\quad (3.142)$$

Differentiating the expression above with respect to \mathbf{A} and setting it to zero, we get

$$E [-2\mathbf{f}_1\mathbf{g}_1^T + 2\mathbf{A}\mathbf{g}_1\mathbf{g}_1^T] = 0. \quad (3.143)$$

Now, $E [\mathbf{g}_1\mathbf{g}_1^T] = E [(\mathbf{g} - \bar{\mathbf{g}})(\mathbf{g} - \bar{\mathbf{g}})^T] = \boldsymbol{\sigma}_g$, the covariance matrix of \mathbf{g} . Similarly, $E [\mathbf{f}_1\mathbf{g}_1^T] = \boldsymbol{\sigma}_{fg}$, the cross-covariance matrix of \mathbf{f} and \mathbf{g} . Thus, we get $\mathbf{A} = \boldsymbol{\sigma}_{fg}\boldsymbol{\sigma}_g^{-1}$. Finally, we obtain the LMMSE estimate as

$$\hat{\mathbf{f}} = \bar{\mathbf{g}} + \boldsymbol{\sigma}_{fg}\boldsymbol{\sigma}_g^{-1}(\mathbf{g} - \bar{\mathbf{g}}). \quad (3.144)$$

This expression reduces to Equation 3.135 when local statistics are substituted for the expectation-based statistical parameters.

Due to the similarity of the optimization criterion used, the LLMMSE filter as above is also referred to as the Wiener filter in some publications [199, 203]. The use of local statistics overcomes the limitations of the Wiener filter due to the assumption of stationarity; the procedure also removes the need to invert large matrices. Nonlinearity, if it is of concern, is the price paid in order to gain these advantages.

Example: Figure 3.43 (f) shows the result of application of the LLMMSE filter as in Equation 3.135, using a 5×5 window, to the noisy test image in part (b) of the same figure. The noise in the uniform background regions in the image, as well as within the geometric objects with uniform gray levels, has been suppressed well by the filter. However, the noise on and around the edges of the objects has not been removed. Although the filter has led to a reduction in the RMS error, the leftover noise around edges has led to a poor appearance of the result.

Refined LLMMSE filter: In a refined version of the LLMMSE filter [204], if the local signal variance $\sigma_g^2(m, n)$ is high, it is assumed that the processing window is overlapping an edge. With the further assumption that the edge is straight, which is reasonable for small windows, the direction of the edge is computed using a gradient operator with eight possible directions for the edge. According to the direction of the edge detected, the processing window is split into two sub-areas, each of which is assumed to be uniform; see Figure 3.44. Then, the statistics computed within the sub-area that holds the pixel being processed are used in Equation 3.135 to estimate the output for the current pixel. This step reduces the noise present in the neighborhood of edges without blurring the edges. Over uniform areas where $\sigma_g^2(m, n)$ has reasonably small values, the statistics are computed over the whole area overlapped by the processing window.

Results of application of the refined LLMMSE filter are presented in Section 3.8.




FIGURE 3.44

Splitting of a 7×7 neighborhood for adaptive filtering based upon the direction of a local edge detected within the neighborhood in a refined version of the local LMMSE filter [204]. One of the eight cases shown is selected according to the direction of the gradient within the 7×7 neighborhood. Pixels in the partition containing the pixel being processed are used to compute the local statistics and the output of the filter. Based upon a similar figure in J.S. Lee, “Refined filtering of image noise using local statistics”, *Computer Graphics and Image Processing*, 15:380–389, 1981.

3.7.2 The noise-updating repeated Wiener filter

The noise-updating repeated Wiener (NURW) filter was introduced by Jiang and Sawchuk [179] to deal with signal-independent additive, signal-dependent Poisson, and multiplicative noise. The image is treated as a random field that is nonstationary in mean and variance [188]. The NURW filter consists of an iterative application of the LLMMSE filter. After each iteration, the variance of the noise is updated for use in the LLMMSE estimate formula of Equation 3.135 in the next iteration as

$$\sigma_{\eta}^{2\text{new}}(m, n) = \left[1 - \frac{\sigma_{\eta}^2(m, n)}{\sigma_g^2(m, n)} + \frac{1}{(2P+1)(2Q+1)} \frac{\sigma_{\eta}^2(m, n)}{\sigma_g^2(m, n)} \right]^2 \sigma_{\eta}^2(m, n) + \left[\frac{1}{(2P+1)(2Q+1)} \frac{\sigma_{\eta}^2(m, n)}{\sigma_g^2(m, n)} \right]^2 \underbrace{\sum_{p=-P}^{+P} \sum_{q=-Q}^{+Q} \sigma_{\eta}^2(m+p, n+q)}_{(p,q) \neq (0,0)}. \quad (3.145)$$

By iterating the filter, noise is substantially reduced even in areas near edges. In order to avoid the blurring of edges, a different (smaller) processing window size may be chosen for each iteration. Jiang and Sawchuk [179] demonstrated

the use of the NURW filter to improve the quality of images degraded with additive, multiplicative, and Poisson noise.

Results of application of the NURW filter are presented in Section 3.8.

3.7.3 The adaptive 2D LMS filter

Noise reduction is usually accomplished by a filtering procedure optimized with respect to an error measure; the most widely used error measure is the MSE. The Wiener filter is a classical solution to this problem. However, the Wiener filter is designed under the assumption of stationary as well as statistically independent signal and noise PDF models. This premise is unlikely to hold true for images that contain large gray-level fluctuations such as edges; furthermore, it does not hold true for signal-dependent noise. Recent methods of circumventing this problem have taken into account the nonstationarity of the given image. One example is the method developed by Chan and Lim [205] which takes into account the image nonstationarity by varying the filter parameters according to the changes in the characteristics or statistics of the image. Another example is the adaptive 2D LMS algorithm developed by Hadhoud and Thomas [206], which is described next.

The 2D LMS method is an example of a fixed-window Wiener filter in which the filter coefficients vary depending upon the image characteristics. The algorithm is based on the method of steepest descent, and tracks the variations in the local statistics of the given image, thereby adapting to different image features. The advantage of this algorithm is that it does not require any *a priori* information about the image, the noise statistics, or their correlation properties. Also, it does not require any averaging, differentiation, or matrix operations.

The 2D LMS algorithm is derived by defining a causal FIR filter $w_l(p, q)$ whose region of support (ROS) is $P \times P$ (P typically being 3) such that

$$\tilde{f}(m, n) = \sum_{p=0}^{P-1} \sum_{q=0}^{P-1} w_l(p, q) g(m-p, n-q), \quad (3.146)$$

where $\tilde{f}(m, n)$ is the estimate of the original pixel value $f(m, n)$; $g(m, n)$ is the noise-corrupted input image; and l marks the current position of the filter in the image, which is given by $l = mM + n$ for the pixel position (m, n) in an $M \times N$ image, and will take values from 0 to $MN - 1$.

The filter coefficients $w_{l+1}(p, q)$ for the pixel position $l + 1$ are determined by minimizing the MSE between the desired pixel value $f(m, n)$ and the estimated pixel value $\tilde{f}(m, n)$ at the present pixel location l , using the method of steepest descent. The filter coefficients $w_{l+1}(m, n)$ are estimated as the present coefficients $w_l(p, q)$ plus a change proportional to the negative gradient of the error power (MSE), expressed as

$$w_{l+1}(p, q) = w_l(p, q) - \mu \nabla [e_l^2], \quad (3.147)$$

where μ is a scalar multiplier controlling the rate of convergence and filter stability; e_l is the error signal, defined as the difference between the desired signal $f(m, n)$ and the estimate $\tilde{f}(m, n)$; and ∇ is a gradient operator [with respect to $w_l(p, q)$] applied to the error power e_l^2 at l . Because the original image $f(m, n)$ is unknown, and the only image on hand is the noise-corrupted image $g(m, n)$, an approximation to the original image $d(m, n)$ is used, and the error is estimated as

$$e_l = d(m, n) - \tilde{f}(m, n). \quad (3.148)$$

The technique used by Hadhoud and Thomas [206] to obtain $d(m, n)$ was to estimate it from the input image $g(m, n)$ by decorrelation, the decorrelation operator being the 2D delay operator of (1, 1) samples. This allows the correlation between $d(m, n)$ and $g(m, n)$ to be similar to the correlation between $f(m, n)$ and $g(m, n)$, and, in turn, makes $d(m, n)$ correlated to $f(m, n)$ to some extent. Evaluation of Equation 3.147 using Equation 3.146 and Equation 3.148 gives

$$w_{l+1}(p, q) = w_l(p, q) + 2 \mu e_l g(m - p, n - q), \quad (3.149)$$

which is a recursive equation defining the filter coefficients at the pixel position $l + 1$ in terms of those at the position l .

Implementation of the 2D LMS filter: Equation 3.146 and Equation 3.149 give the 2D LMS filter and the filter weight updating algorithms, respectively. Convergence of the algorithm does not depend upon the initial conditions; it converges for any arbitrary initial value, and hence provides good nonstationary performance. In comparing the 2D LMS filter with the nonadaptive LMS algorithm, the second term of Equation 3.149 would not be included in the latter, and the filter coefficients would not change from pixel to pixel under the assumption that the image is stationary. This would put a constraint on the initial coefficient values, because they would be the values used for the whole image, and thus, different initial values would result in different filtered outputs. Although the initial conditions do not affect the convergence of the 2D LMS filter, the choice of the convergence factor μ depends on the particular application, and involves a trade-off between the rate of convergence, the ability to track nonstationarity, and steady-state MSE.

Example: Figures 3.45 (a) and (b) show a test image and its noisy version, the latter obtained by adding zero-mean Gaussian noise with $\sigma^2 = 256$ to the former. Part (c) of the figure shows the result of application of the 2D LMS filter to the noisy image. In implementing Equation 3.146 and Equation 3.149, the initial weights of the filter, $w_0(p, q)$, were estimated by processing 10 lines of the given image starting with zero weights. The weights obtained after processing the 10 lines were then used as the initial conditions, and processing was restarted at $(m, n) = (0, 0)$ in the given image. The convergence factor μ was determined by trial and error for different images as suggested by Hadhoud and Thomas [206], and set to 0.4×10^{-7} ; the ROS used was 3×3 .

The 2D LMS algorithm applies a gradually changing filter that tends to suppress noise in a relatively uniform manner over the image. Although this typically results in lower values of the MSE, it also tends to smooth or blur edges and other structured features in the image, and to leave excessive noise in uniform regions of the image. One explanation to this effect could be the fact that the adaptive weights of the filter, $w_l(p, q)$, depend on the model of the original image which is approximated by the decorrelated image $d(m, n)$. Because $d(m, n)$ is not an accurate approximation of the original image $f(m, n)$, the updated filter weights are not optimal, and thus the algorithm does not give the optimal MSE value. The 2D LMS algorithm did not perform particularly well on the test image in Figure 3.45 (b). The resultant image in Figure 3.45 (c) appears significantly blurred, with remnant noise.

3.7.4 The adaptive rectangular window LMS filter

In order to overcome the limitations due to the assumption of stationarity, a Wiener filter approach using an adaptive-sized rectangular window (ARW) to estimate the filter coefficients was proposed by Song and Pearlman [207, 208, 209], and refined by Mahesh et al. [210]. Using the same image degradation model as that in Equation 3.109 but with the additional assumption that the image processes also have zero mean, the estimate used by Mahesh et al. is of the form

$$\tilde{f}(m, n) = \alpha(m, n) g(m, n). \quad (3.150)$$

The problem reduces to that of finding the factor $\alpha(m, n)$ at each pixel location using the same minimum MSE criterion as that of the standard Wiener filter. The error is given by

$$e(m, n) = f(m, n) - \tilde{f}(m, n) = f(m, n) - \alpha(m, n) g(m, n). \quad (3.151)$$

Minimization of the MSE requires that the error signal e be orthogonal to the image g , that is,

$$E\{[f(m, n) - \alpha(m, n) g(m, n)] g(m, n)\} = 0. \quad (3.152)$$

Solving for α , we obtain

$$\alpha(m, n) = \frac{\sigma_f^2(m, n)}{\sigma_f^2(m, n) + \sigma_\eta^2(m, n)}. \quad (3.153)$$

If the original image is not a zero-mean process, Equation 3.150 can still be used by first subtracting the mean from the images f and g . Because the noise is of zero mean for all pixels, the *a posteriori* mean $\mu_g(m, n)$ of the image g at pixel position (m, n) is equal to the *a priori* mean $\mu_f(m, n)$ of the original image $f(m, n)$, and the estimate of Equation 3.150 thus becomes

$$\tilde{f}(m, n) = \mu_g(m, n) + \frac{\sigma_f^2(m, n)}{\sigma_f^2(m, n) + \sigma_\eta^2(m, n)} [g(m, n) - \mu_g(m, n)]. \quad (3.154)$$




FIGURE 3.45

(a) “Shapes3”: a 128×128 test image with various geometrical objects placed at random. (b) Image in (a) with Gaussian noise added, RMS error = 14.24. Result of filtering the image in (b) with: (c) the 2D LMS filter, RMS error = 15.40; (d) two passes of the ARW-LMS filter, RMS error = 7.07; (e) the ANNS filter, RMS error = 6.68; (f) two passes of the ANNS filter, RMS error = 5.10. Reproduced with permission from R.B. Paranjape, T.F. Rabie, and R.M. Rangayyan, “Image restoration by adaptive-neighborhood noise subtraction”, *Applied Optics*, 33(14):2861–2869, 1994.

© Optical Society of America.

Implementation of the ARW-LMS filter: In implementing Equation 3.154, it is necessary to make the assumption that the image pixel values in the immediate neighborhood of a pixel (m, n) are samples from the same ensemble as that of $f(m, n)$; that is, a globally nonstationary process can be considered to be locally stationary and ergodic over a small region. Thus, if we can accurately determine the size of a neighborhood in which the image values have the same statistical parameters, the sample statistics can approximate the *a posteriori* parameters needed for the estimate in Equation 3.154.

The view taken by Song and Pearlman [207, 208, 209] was to identify the size of a stationary square region for each pixel in the image, and to calculate the local statistics of the image within that region. The size of the window changes according to a measure of signal activity; an effective algorithm was proposed to determine the window size that improved the performance of various point estimators. The effect of the improved performance was greater smoothing in relatively flat (signal-free) regions in the image and less smoothing across edges.

In deriving the local sample statistics denoted as $\tilde{\mu}_g(m, n)$ and $\tilde{\sigma}_f^2(m, n)$, Mahesh et al. [210] made use of ARWs of length L_r in the row direction and L_c in the column direction. Because the window is required to be centered about the pixel being processed, the ARW lengths need to be odd. Except near the borders of the image, the ARW dimensions can be expressed as $L_r = 2N_r + 1$ and $L_c = 2N_c + 1$, where N_r and N_c are the dimensions of the one-sided neighborhood. Within this window, the local mean and variance are calculated as

$$\tilde{\mu}_g(m, n) = \frac{1}{L_r L_c} \sum_{p=-N_r}^{+N_r} \sum_{q=-N_c}^{+N_c} g(m+p, n+q), \quad (3.155)$$

and

$$\tilde{\sigma}_g^2(m, n) = \frac{1}{L_r L_c} \sum_{p=-N_r}^{+N_r} \sum_{q=-N_c}^{+N_c} [g(m+p, n+q) - \tilde{\mu}_g(m, n)]^2. \quad (3.156)$$

The local variance of the original image $\tilde{\sigma}_f^2$ is estimated as

$$\tilde{\sigma}_f^2 = \begin{cases} \tilde{\sigma}_g^2 - \sigma_\eta^2 & \text{if } \tilde{\sigma}_g^2 > \sigma_\eta^2 \\ 0 & \text{otherwise.} \end{cases} \quad (3.157)$$

Using the local sample statistics as above, Equation 3.154 becomes

$$\tilde{f}(m, n) = \tilde{\mu}_g(m, n) + \frac{\tilde{\sigma}_f^2(m, n)}{\tilde{\sigma}_f^2(m, n) + \sigma_\eta^2(m, n)} [g(m, n) - \tilde{\mu}_g(m, n)]. \quad (3.158)$$

It should be noted that the parameters $L_r, L_c, N_r, N_c, \tilde{\mu}_g, \tilde{\sigma}_g^2$, and $\tilde{\sigma}_f^2$ as well as the other parameters that follow are computed for each pixel (m, n) ,

and should be denoted as $L_r(m, n)$, etc.; this detail has been suppressed for convenience of notation. It should also be noted that, although the noise variance σ_η^2 is usually not known *a priori*, it can be easily estimated from a window in a flat (signal-free) area of the degraded image.

For accurate estimation of μ_g and σ_g^2 , the pixels in the ARW should belong to the same ensemble as that of the central pixel being filtered. If relatively large windows are used, the windows may cross over the boundaries of different regions and include pixels from other ensembles. In such a case, blurring could result across the edges present within the windows. On the other hand, if the windows are too small, the lack of samples would result in poor estimates of the mean and variance, and consequently, insufficient noise suppression would occur over uniform regions. Thus, it is desirable to use small windows where the image intensity changes rapidly, and large windows where the image contrast is relatively low.

In the method of Mahesh et al. [210], the ARW lengths L_r and L_c are varied depending upon a signal activity parameter S_r defined as

$$S_r(m, n) = \frac{1}{L_r L_c} \sum_{p=-N_r}^{+N_r} \sum_{q=-N_c}^{+N_c} [g(m+p, n+q) - \tilde{\mu}_r]^2 - \sigma_\eta^2, \quad (3.159)$$

where $\tilde{\mu}_r$ is the local mean evaluated in the row direction as

$$\tilde{\mu}_r = \frac{1}{L_r} \sum_{p=-N_r}^{+N_r} g(m+p, n). \quad (3.160)$$

S_r is a measure of local roughness of the image in the row direction, being equal to the variance of the original image in the same direction. A similar signal activity parameter is defined in the column direction. If the signal activity parameter S_r in the row direction is large, indicating the presence of an edge or other information, the window size in the row direction N_r is decremented so that points from other ensembles are not included. If S_r is small, indicating that the current pixel lies in a low-contrast region, N_r is incremented so that a better estimate of the mean and variance may be obtained. In order to make this decision, the signal activity parameter in the row direction is compared to a threshold T_r , and N_r is updated as follows:

$$N_r \leftarrow N_r - 1, \text{ if } S_r \geq T_r, \quad (3.161)$$

$$\text{or } N_r \leftarrow N_r + 1, \text{ if } S_r < T_r. \quad (3.162)$$

A similar procedure is applied in the column direction to update N_c .

Prespecified minimum and maximum values for N_r and N_c are used to limit the size of the ARW to reasonable dimensions, which could be related to the size of the details present in the image. The threshold T_r is defined as

$$T_r = \frac{\kappa \sigma_\eta^2}{L_r}, \quad (3.163)$$

where κ is a weighting factor that controls the rate at which the window size changes. The threshold varies in direct proportion to the noise variance and in inverse proportion to the ARW dimension. Thus, if the noise variance is high, the threshold will be high, and the window length is more likely to be incremented than decremented. This will lead to a large window and effective smoothing of the noise. If the window size is large, the threshold will be small. Thus, the window size is likely to be decremented for the next pixel. A similar threshold is defined in the column direction. This helps the ARW length to converge to a certain range.

Example: The result of application of the ARW-LMS filter to the noisy test image in Figure 3.45 (b) is shown in part (d) of the same figure. The ARW size was restricted to be a minimum of 1×1 and a maximum of 5×5 ; the value of the weighting factor κ in Equation 3.163 was fixed at 7. The ARW-LMS algorithm has resulted in much less smoothing at the edges of the objects in the image than in the uniform regions. As a result, a layer of noise remains in the filtered image surrounding each of the objects. Although this is clearly objectionable in the synthesized image, the effect was not as pronounced in the case of natural scenes, because ideal edges are not common in natural scenes.

The ARW-LMS output image appears to be clearer and sharper than the 2D LMS restored image [see Figure 3.45 (c)], because the ARW-LMS algorithm tends to concentrate the error around the edges of objects where the human visual system tends to ignore the artifact. On the other hand, the 2D LMS algorithm tends to reduce the error uniformly, which also leads to smoothing across the edges; the decreased sharpness of the edges makes the result less pleasing to the viewer.

3.7.5 The adaptive-neighborhood filter

An adaptive-neighborhood paradigm was proposed by Paranjape et al. [211, 212] to filter additive, signal-independent noise, and extended to multiplicative noise by Das and Rangayyan [213], to signal-dependent noise by Rangayyan et al. [201], and to color images by Ciuc et al. [214]. Unlike the other methods where statistics of the noise and signal are estimated locally within a fixed-size, fixed-shape (usually rectangular) neighborhood, the adaptive-neighborhood filtering approach consists of computing statistics within a variable-size, variable-shape neighborhood that is determined individually for *every* pixel in the image. It is desired that the adaptive neighborhood grown for the pixel being processed (which is called the “seed”) contains only those spatially connected pixels that are similar to the seed, that is, the neighborhood does not grow over edges but overlaps a stationary area. If this condition is fulfilled, the statistics computed using the pixels inside the region are likely to be closer to the true statistics of the local signal and noise components than the statistics computed within fixed neighborhoods. Hence, adaptive-neighborhood filtering should yield more accurate results than fixed-neighborhood methods. The

approach should also prevent the blurring or distortion of edges because adaptive neighborhoods, if grown with an appropriate threshold, should not mix the pixels of an object with those belonging to its background.

There are two major steps in adaptive-neighborhood filtering: adaptive region growing and estimation of the noise-free value for the seed pixel using statistics computed within the region.

Region growing for adaptive-neighborhood filtering: In adaptive-neighborhood filtering, a region needs to be grown for the pixel being processed (the seed) such that it contains only pixels belonging to the same object or image feature as the seed. Paranjape et al. [211, 212] used the following region-growing procedure for images corrupted by signal-independent Gaussian additive noise: the absolute difference between each of the 8-connected neighbors $g(p, q)$ and the seed $g(m, n)$ is computed as

$$d_{pq} = |g(p, q) - g(m, n)|. \quad (3.164)$$

Pixels $g(p, q)$ having $d_{pq} \leq T$, where T is a fixed, predefined threshold, are included in the region. The procedure continues by checking the neighbors of the newly included pixels in the same manner, and stops when the inclusion criterion is not fulfilled for any neighboring pixel. An adaptive neighborhood is grown for *each pixel* in the image.

In addition to the foreground region, an adaptive background region is also grown for each pixel. The background is obtained by expanding (dilating) the outermost boundary of the foreground region by a prespecified number of pixels. This provides a ribbon of a certain thickness that surrounds the foreground region. The foreground region may be further modified by combining isolated pixels into local regions. Observe that a foreground region could contain several disjoint regions that are not part of the foreground due to differences in gray level that are larger than that permitted by the threshold applied; such regions could be considered to be part of the background, although they are enclosed by the foreground region.

Example: Figure 3.46 illustrates the growth of an adaptive neighborhood. The foreground part of the adaptive neighborhood is shown in a light shade of gray. The black pixels within the foreground have the same gray level as that of the seed pixel from where the process was commenced; the use of a simple threshold for region growing as in Equation 3.164 will result in the same region being grown for all such pixels: for this reason, they could be called *redundant seed pixels*. The region-growing procedure need not be applied to such pixels, which results in computational savings. Furthermore, all redundant seed pixels within a region get the same output value as computed for the original seed pixel from where the process was commenced.

Adaptive-neighborhood mean and median filters: Paranjape et al. [211] proposed the use of mean and median values computed using adaptive neighborhoods to filter noise. The basic premise was that context-dependent adaptive neighborhoods provide a larger population of pixels to compute local statistics than 3×3 or 5×5 neighborhoods, and that edge distortion


**FIGURE 3.46**

Illustration of the growth of an adaptive neighborhood (left to right and top to bottom). The neighborhood being grown is shown in a light shade of gray. The black pixels within the neighborhood have the same gray level as that of the seed from where the process was commenced: they are called redundant seed pixels. The last figure shows a region in a darker shade of gray that surrounds the foreground region: this is known as the adaptive background region. Figure courtesy of W.M. Morrow [215].

is prevented because such neighborhoods are not expected to transgress the boundaries of the objects or features in the image. They also showed that the method could be iterated to improve the results.

Example: Figures 3.47 (a) and (b) show a test image and its noisy version with additive Gaussian-distributed noise. Parts (c) and (d) of the figure show the results of filtering the noisy image using the 3×3 mean and median, respectively. The blurring of edges caused by the mean, and the distortion of shape caused by the median, are clearly seen in the results. Parts (e) and (f) of the figure illustrate the results of the adaptive-neighborhood mean and median filters, respectively. Both methods have been equally effective in removing the noise without causing any blurring or distortion of edges. However, in other experiments with high levels of noise, and when the process was iterated, it was observed that the adaptive-neighborhood filters could lead to the loss of objects that have small gray-level differences with respect to their surroundings. The adaptive neighborhood, in such a case, could mix an object and its surroundings if the threshold is low compared to the noise level and the contrast of the object.

Adaptive-neighborhood noise subtraction (ANNS): Paranjape et al. [212] proposed the ANNS method to remove additive, signal-independent noise. The algorithm estimates the noise value at the seed pixel $g(m, n)$ by using an adaptive neighborhood, and then subtracts the noise value from the seed pixel to obtain an estimate of the original undegraded value $\tilde{f}(m, n)$.

The strategy used in deriving the ANNS filter is based upon the same principles as those of the ARW-LMS algorithm; the image process f is assumed to be a zero-mean process of variance σ_f^2 that is observed in the presence of additive white Gaussian noise, resulting in the image g . The noise process η is assumed to have zero mean and variance of σ_η^2 , and assumed to be uncorrelated to f . An estimate of the additive noise at the pixel (m, n) is obtained from the corresponding adaptive neighborhood grown in the corrupted image g as

$$\tilde{\eta}(m, n) = \alpha g(m, n), \quad (3.165)$$

where α is a scale factor which depends on the characteristics of the adaptive neighborhood grown. Then, the estimate of $f(m, n)$ is

$$\tilde{f}(m, n) = g(m, n) - \tilde{\eta}(m, n), \quad (3.166)$$


which reduces to

$$\tilde{f}(m, n) = \beta g(m, n), \quad (3.167)$$

where $\beta = 1 - \alpha$.

As described in Section 3.7.4 on the ARW-LMS algorithm, if the images used are of nonzero mean, the estimate of Equation 3.167 can be used by first subtracting the mean of each image from both sides of the equation. Then, the estimate may be expressed as

$$\tilde{f}(m, n) = \mu_g(m, n) + (1 - \alpha) [g(m, n) - \mu_g(m, n)], \quad (3.168)$$

**FIGURE 3.47**

(a) “Shapes2”: a 128×128 test image with various geometrical objects placed at random. (b) Image in (a) with Gaussian noise added, RMS error = 8.24. Result of filtering the image in (b) with: (c) the 3×3 mean, RMS error = 9.24; (d) the 3×3 median, RMS error = 6.02; (e) the adaptive-neighborhood mean, RMS error = 3.16; (f) the adaptive-neighborhood median, RMS error = 4.01. Images courtesy of R.B. Paranjape.

where μ_g is the *a posteriori* mean of the degraded image $g(m, n)$, which is also equal to the *a priori* mean μ_f of the original image $f(m, n)$ for zero-mean noise.

The problem now is to find the factor α , which is based upon the criterion that the estimated noise variance $\sigma_{\tilde{\eta}}^2$ be equal to the original noise variance σ_{η}^2 . The solution is obtained as follows:

$$\begin{aligned}\sigma_{\tilde{\eta}}^2 &= E[\tilde{\eta}^2] \\ &= E\left[\{\alpha [g(m, n) - \mu_g]\}^2\right] \\ &= \alpha^2 \sigma_g^2 \\ &= \alpha^2 (\sigma_f^2 + \sigma_{\eta}^2).\end{aligned}\tag{3.169}$$

The noise estimation factor α is then given by

$$\alpha = \sqrt{\frac{\sigma_{\eta}^2}{\sigma_f^2 + \sigma_{\eta}^2}}.\tag{3.170}$$

Thus, the estimate of Equation 3.168 becomes

$$\tilde{f}(m, n) = \mu_g(m, n) + \left(1 - \sqrt{\frac{\sigma_{\eta}^2(m, n)}{\sigma_f^2(m, n) + \sigma_{\eta}^2(m, n)}}\right) [g(m, n) - \mu_g(m, n)].\tag{3.171}$$

The indices (m, n) have been introduced in the expression above to emphasize the point that the statistical parameters are computed for every pixel (using the corresponding adaptive neighborhood). The estimate $\tilde{f}(m, n)$ given by Equation 3.171 may be considered to be an approximation of the original image $f(m, n)$ if we are able to obtain accurate values for the statistical parameters $\mu_g(m, n)$ and $\sigma_f^2(m, n)$.

Implementation of the ANNS filter: In implementing Equation 3.171, we need to derive the local (sample) statistics from the adaptive neighborhood grown at every seed pixel location (m, n) in the degraded image. This could be achieved in a manner similar to that described in Section 3.7.4 for the ARW-LMS filter.

Paranjape et al. [212] defined the tolerance used for growing adaptive neighborhoods in an adaptive manner, depending upon the signal activity in the region and the features surrounding it, as follows. A limit of Q pixels is set for the adaptive neighborhoods. An initial region is first grown with the tolerance set to the full dynamic range of the input image (that is, $T = 256$). This results in a square region of size Q pixels being formed. Using the foreground pixels in the adaptive neighborhood, a measure of the uncorrupted signal activity in the adaptive neighborhood of the seed pixel is given by the local signal variance $\tilde{\sigma}_f^2$. The signal variance in the adaptive neighborhood is then compared with the noise variance σ_{η}^2 , and if $\tilde{\sigma}_f^2 > 2\sigma_{\eta}^2$, it is assumed that

the adaptive neighborhood has identified a region with significant structural characteristics such as an edge, or other distinct objects. This is contrary to the desired characteristics of an adaptive neighborhood. An adaptive neighborhood is to be formed such that it includes relatively uniform structures (or background) in the original image, so that the primary source of variation in the adaptive neighborhood is the additive noise. Therefore, the gray-level tolerance used to define the adaptive neighborhood is modified to $T = 2\tilde{\sigma}_f$, with the notion that the signal standard deviation $\tilde{\sigma}_f$ be used to define the new adaptive neighborhood. The adaptive neighborhood is grown again using the new tolerance. Because the tolerance has been reduced, the new adaptive neighborhood developed, presumably, will not contain edges or structural features in the image, but will rather grow up to but not include such features. Using this approach to define the adaptive neighborhoods, the statistics of the adaptive neighborhoods are used in Equation 3.171 to estimate the uncorrupted image at each pixel location. In the situation that the foreground has only one pixel, the adaptive neighborhood is enlarged to include the background layer of pixels (of size 3×3); this approach is particularly useful in the presence of impulse noise or outliers in the corrupted image.

The advantage of the ANNS method lies in the fact that, in flat or slowly varying regions, the signal variance will be small compared to $2\sigma_\eta^2$, and the adaptive neighborhood foreground will grow to the maximum foreground bound of Q pixels (but of arbitrary shape depending upon the local image features). On the other hand, in busy regions where the signal variance is high compared to $2\sigma_\eta^2$, the tolerance will be reduced and the foreground will grow up to any edge present but not across the edge. This results in the removal of noise up to the edges present in the image, which the ARW-LMS filter fails to do [see Figure 3.45 (d)].

Example: Figure 3.45 (e) shows the result of application of the ANNS method to the noisy test image in part (b) of the same figure. The maximum adaptive neighborhood size Q was set to 25 pixels; this allows for direct comparison of the performance of the ANNS method against the ARW-LMS method. However, the ANNS algorithm uses a variable-shape window (adaptive neighborhood) in order to compute the filtered image, whereas the ARW-LMS method is restricted to rectangular windows. Unlike the ARW-LMS filter window, the size of the adaptive neighborhood is not compromised near the edges in the image; rather, its shape changes according to the contextual details present in the image. This aspect of the ANNS method allows for better estimation of the noise near edges, and thus permits greater noise suppression in such areas. Both the ANNS and ARW-LMS methods appear to have reduced the noise equally well in relatively uniform regions of the image; however, the ARW-LMS filter output contains residual noise around the edges of the objects in the image.

Repeated (iterative) application is a powerful and useful attribute of the ARW-LMS and ANNS methods. Figure 3.45 (f) shows the result of two-pass ANNS filtering of the test image in part (b) of the figure. The ARW-LMS

output in part (d) of the figure was obtained using two iterations. For both of the algorithms, updated values of the noise variance are required for the second and subsequent passes through the filter. The second-pass noise variance was estimated by calculating the variance of the output image after the first pass, and subtracting from it an estimate of the variance of the noise-free image $f(m, n)$. The resulting images [see Figures 3.45 (d) and (f)] show significant improvement over the results from a single pass through the algorithms. The major artifact after the first pass through the algorithms was the retention of noise around distinct edges in the image. With both algorithms, such layers of noise were greatly reduced after the second application; the ANNS method performed better than the ARW-LMS method after the second pass. A second artifact that was observed was the patchy appearance of the originally uniform background regions in the result after the first pass; this artifact too was reduced after the second pass with both algorithms.

Extension to filter signal-dependent noise: Rangayyan et al. [201] used the same basic region-growing procedure as above, because most types of noise (Poisson, film-grain, and speckle) can be modeled as additive noise, with modifications with respect to the fact that the noise is signal-dependent. The first modification is to let the threshold T vary across the image according to the local statistics of noise, which depend upon the average brightness of the corrupted area. It is obvious that the performance of the filter strongly depends upon the threshold value. A small T would lead to small regions over which the noise statistics cannot be reliably estimated, whereas a large T would lead to regions that may grow over edges present in the image, with the result that stationarity in such regions is no longer guaranteed. A reasonable value for T , large enough to allow inclusion in the region of representative samples of noise and small enough to prevent the region from growing over edges, is the *local standard deviation of noise*:

$$T = \sigma_{\eta}(m, n). \quad (3.172)$$

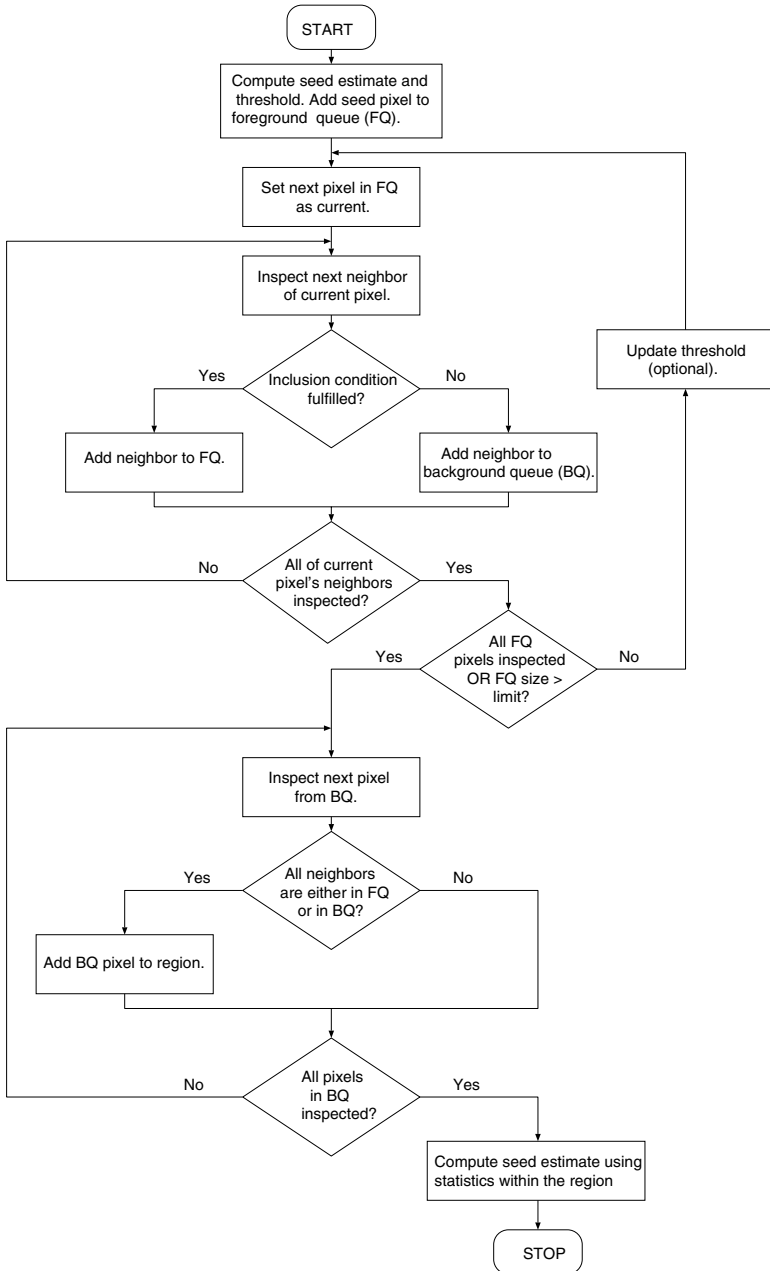
Before growing the region for the pixel located at the coordinates (m, n) , one has to estimate $\sigma_{\eta}(m, n)$. A coarse estimation would consist of taking the actual value of the seed $g(m, n)$ for the statistical expectation values in Equation 3.27 for Poisson noise, Equation 3.29 for film-grain noise, and Equation 3.31 for speckle noise. However, when the noise level is high, more accurate estimation is required, because the corrupted seed value may differ significantly from its average (expected value).

Estimation of seed and threshold values: Rangayyan et al. [201] defined an initial estimate for the seed pixel value $g(m, n)$, defined as the α -trimmed mean value $g_{\alpha TM}(m, n)$ or the median, computed within a 3×3 window. This step was found to be useful in order to prevent the use of a pixel significantly altered by noise (an outlier) for region growing. Then, the noise standard deviation was computed using the initial estimate in place of $E\{g(m, n)\}$ in Equation 3.27, 3.29, or 3.31, depending upon the type of noise.

The threshold T was defined as in Equation 3.172, and the seed's neighbors were inspected for inclusion in the region according to Equation 3.164, in which the actual seed value $g(m, n)$ was replaced by the initial estimate. Because the initial estimate provides only a basic estimate of the seed value, the estimated noise standard deviation and, consequently, the threshold value T might differ from their true or optimal values. To overcome this drawback, T may be continuously updated while the region is being grown, its value being computed as before, but by using the mean of the pixels already included in the region, for the expectation values in Equation 3.27, 3.29, or 3.31.

Revising the region to reduce bias: While the region is being grown, pixels that are inspected but do not meet the inclusion criterion are marked as background pixels. After the region-growing procedure stops, there are three types of pixels: region (or foreground) pixels, background pixels, and pixels that were not inspected as they are not connected to any pixel in the region. The area that holds the foreground pixels is 8-connected, although not necessarily compact, whereas the background is composed of several disconnected areas, many of which may be, at most, two pixels-wide, and could be lying within the foreground area as well. As a region is desired to be compact, that is, it does not contain holes with one or two pixels, some of the background pixels could be further checked for inclusion in the region. One can interpret such pixels as belonging to the same region as the seed, that, due to noise, did not meet the inclusion criterion in the first step of region growing. Ignoring such pixels would result in biased estimation of the signal or noise statistics. On the other hand, there would be background pixels that are adjacent to the external border of the foreground area that should not be included in the region because they, most likely, belong to other objects. For these reasons, Rangayyan et al. [201] derived a criterion for further inclusion of selected background pixels in the region: they included in the region all background pixels whose 8-connected neighbors are all either in the foreground or in the background. If only one of a background pixel's neighbor was not inspected in the initial region-growing step, then that pixel was exempted from inclusion in the foreground. This criterion was found to work efficiently in many trials with different types of noise. Figure 3.48 presents a flowchart of the adaptive region-growing procedure.

The criterion for the inclusion of background pixels in the region described above does not take into account the values of the background pixels, and is based only on their spatial relationships. It was implicitly assumed that all objects in the image are at least three pixels wide in any direction. However, for images that may contain small objects or regions (as in fine texture), the values of the inspected background pixels should also be taken into account. It was suggested that, in such a case, if a background pixel fulfills the spatial inclusion criterion, it should be added to the region only if its value does not differ significantly from the average value of the pixels in the region (for example, the absolute difference between the inspected background pixel value

**FIGURE 3.48**

Flowchart of the adaptive region-growing procedure. Reproduced with permission from R.M. Rangayyan, M. Ciuc, and F. Faghii, "Adaptive neighborhood filtering of images corrupted by signal-dependent noise", *Applied Optics*, 37(20):4477-4487, 1998. © Optical Society of America.

and the average value of the pixels in the region is smaller than twice the local variance of the noise).

Figure 3.49 illustrates a sample result of the adaptive region-growing procedure. The foreground region size was limited to a predetermined number of pixels (100) in order to reduce computing time.

Adaptive-region-based LLMMSE filter: Once an adaptive region is grown, statistics of the signal, and, according to them, statistics of the noise, are computed using the pixels in the foreground region obtained. The mean and variance of the noisy signal are computed using the pixels in the foreground instead of using a rectangular window as commonly done. The variance of the noise is computed using Equation 3.27, 3.29, or 3.31, depending upon the type of noise. Finally, the LLMMSE estimate is computed according to Equation 3.135 and assigned to the seed pixel location in the output image.

The second term in the LLMMSE estimate formula can be interpreted as a correction term whose contribution to the final result is important when the variance of the signal is much larger than the variance of the noise. Over flat regions, where the variations in the image are due mostly to the noise, the major contribution is given by the first term; that is, the mean of the noisy signal. This is always the case within a region that is grown to be relatively uniform, as in adaptive region growing. Hence, the mean of the foreground pixels by itself represents a good estimator of the noise-free seed pixel; this is advantageous when computational time specifications are restrictive.

Example: Figure 3.50 (a) shows a test image of a clock. The image contains a significant amount of noise, suspected to be due to poor shielding of the video-signal cable between the camera and the digitizing frame buffer. Parts (b)–(f) of the figure show the results of application of the 3×3 mean, 3×3 median, refined LLMMSE, NURW, and adaptive-neighborhood LLMMSE filters to the test image. The filters have suppressed noise to similar levels; however, the mean filter has caused significant blurring of the image, the median has resulted in some shape distortion in the numerals of the clock, and the refined LLMMSE has led to a patchy appearance. The NURW and adaptive-neighborhood LLMMSE filters have provided good noise suppression without causing edge degradation.

3.8 Comparative Analysis of Filters for Noise Removal

In a comparative study [201] of several of the filters described in the preceding sections, the Shapes and Peppers images were corrupted by different types of noise and filtered. The performance of the filters was assessed using the MSE and by visual inspection of the results. The following paragraphs provide the details of the study.



**FIGURE 3.49**

Illustration of the steps in adaptive region growing: (a) 25×25 pixel-wide portion of the original Peppers image. (b) Image corrupted by Poisson noise with $\lambda=0.1$. (c) The seed pixel, shown in white and located at the center of the image. (d) First step of region growing on the corrupted image: foreground pixels are in white, background pixels are in light gray. The foreground size has been limited to 100 pixels. (e) Region after inclusion of interior background pixels. (f) Filtered image. In (c), (d), and (e), the region has been superimposed over the uncorrupted image for convenience of display. Figure courtesy of M. Ciuc, Laboratorul de Analiza și Prelucrarea Imaginilor, Universitatea Politehnica București, Bucharest, Romania.

**FIGURE 3.50**

(a) Clock test image. Result of filtering the image in (a) using: (b) 3×3 mean; (c) 3×3 median; (d) the refined LLMMSE filter; (e) NURW filter; and (f) adaptive-neighborhood LLMMSE filter. Figure courtesy of M. Ciuc, Laboratorul de Analiza și Prelucrarea Imaginilor, Universitatea Politehnica București, Bucharest, Romania.

Additive, Gaussian noise: Random noise with Gaussian distribution having zero mean and standard deviation of 20 was added to the test images. Figures 3.51 and 3.52 show the original images, their noisy versions, and the results of a few selected filters. The MSE values of the noisy and filtered images are listed in Tables 3.1 and 3.2. The application of the LLMMSE filter using the adaptive-neighborhood paradigm led to the best results with both images. The NURW and adaptive-neighborhood mean filters also provided good results with the Peppers image.


Additive, uniformly distributed noise: The Shapes and Peppers test images were corrupted by adding uniformly distributed noise with $\mu = 0$, $\sigma = 20$. The results of filtering the noisy images are shown in Figures 3.53 and 3.54. The MSE values of the images are listed in Tables 3.1 and 3.2. The application of the LLMMSE filter using the adaptive-neighborhood paradigm led to the best result with the Shapes image. With the Peppers image, the NURW filter provided better results than the adaptive-neighborhood LLMMSE filter.

Poisson noise: The test images were corrupted by Poisson noise with $\lambda = 0.1$. The results of filtering the noisy images are shown in Figures 3.55 and 3.56. The MSE values of the images are listed in Tables 3.1 and 3.2. The application of the LLMMSE filter using the adaptive-neighborhood paradigm led to the best result with the Shapes image. With the Peppers image, the NURW filter provided results comparable to those of the adaptive-neighborhood mean and LLMMSE filters.

Film-grain noise: The test images were corrupted by purely signal-dependent film-grain noise, with $\eta_2(m, n) = 0$, $\kappa = 3.3$, and $\sigma_{\eta_1} = 1$ in the model given in Equation 3.28. The results of filtering the noisy images are shown in Figures 3.57 and 3.58. The MSE values of the images are listed in Tables 3.1 and 3.2. The LLMMSE filter using the adaptive-neighborhood paradigm provided the best result with the Shapes image. With the Peppers image, the NURW filter provided results comparable to those of the adaptive-neighborhood mean and LLMMSE filters.

Speckle noise: Speckle noise was simulated by multiplicative noise, the noise having exponential distribution given by

$$p_{\eta(m,n)}(x) = \exp[-\eta(m, n)]. \quad (3.173)$$

**FIGURE 3.51**

(a) Shapes: a 128×128 test image. (b) Image in (a) with Gaussian noise added, with $\mu = 0$, $\sigma = 20$, $\text{MSE} = 228.75$. Result of filtering the noisy image in (b) using: (c) 3×3 LLMMSE, $\text{MSE} = 108.39$; (d) NURW, $\text{MSE} = 132.13$; (e) adaptive-neighborhood mean, $\text{MSE} = 205.04$; and (f) adaptive-neighborhood LLMMSE, $\text{MSE} = 78.58$. Figure courtesy of M. Ciuc, Laboratorul de Analiza și Prelucrarea Imaginilor, Universitatea Politehnica București, Bucharest, Romania.




FIGURE 3.52

(a) 512×512 Peppers test image. (b) Image in (a) with Gaussian noise added, with $\mu = 0$, $\sigma = 20$, $MSE = 389.87$. Result of filtering the noisy image in (b) using: (c) refined LLMMSE, $MSE = 69.49$; (d) NURW, $MSE = 54.70$; (e) adaptive-neighborhood mean, $MSE = 55.21$; and (f) adaptive-neighborhood LLMMSE, $MSE = 52.32$. Figure courtesy of M. Ciuc, Laboratorul de Analiza și Prelucrarea Imaginilor, Universitatea Politehnica București, Bucharest, Romania.

TABLE 3.1MSE values of the Noisy and Filtered Versions of the 128×128 Shapes Image.

Noise type	Noisy	3 Mean	3 Med.	5 Mean	5 Med.	3 LL	5 LL	R-LL	NURW	AN Mean	AN Med.	AN LL
Gaussian	228.75	469.26	213.71	772.76	518.17	108.39	122.29	124.88	132.13	205.04	197.93	78.58
Uniform	226.20	479.68	236.55	785.62	530.75	113.83	130.51	133.63	144.71	216.52	204.90	93.08
Poisson	241.07	441.04	266.85	743.73	657.35	108.70	130.14	131.47	147.87	215.18	249.41	62.57
Film-grain	275.11	450.92	283.74	746.90	665.78	119.81	147.27	141.64	166.42	236.27	296.81	69.98
Speckle	255.43	445.61	278.76	749.00	665.02	119.15	147.43	138.49	166.75	236.09	286.90	68.01
Salt & pepper	1740.86	642.20	206.63	835.46	557.75	1739.37	1405.09	1739.10	1740.84	213.02	205.72	1686.13

Note: 3 = 3×3 . 5 = 5×5 . Med. = Median. LL = LLMMSE. R = Refined. AN = Adaptive neighborhood.

TABLE 3.2
MSE Values of the Noisy and Filtered Versions of the 512×512 Peppers Image.

Noise type	Noisy	3 Mean	3 Med.	5 Mean	5 Med.	3 LL	5 LL	R-LL	NURW	AN Mean	AN Med.	AN LL
Gaussian	389.87	74.89	93.55	84.88	71.80	86.53	68.19	69.49	54.70	55.21	57.50	52.32
Uniform	391.43	75.09	129.08	85.30	89.10	76.17	63.02	65.25	54.39	62.53	70.98	58.62
Poisson	1132.56	159.29	239.26	116.29	139.50	197.87	121.71	133.82	85.32	88.83	110.10	87.48
Film-grain	1233.43	168.25	245.77	119.59	135.32	212.03	125.25	117.54	88.83	90.54	101.19	89.07
Speckle	988.84	142.62	204.39	110.91	119.98	172.35	105.67	100.76	77.59	81.53	91.45	79.26
Salt & pepper	947.69	144.01	22.93	117.02	38.70	886.03	832.25	821.88	872.37	34.49	29.27	861.01


Note: 3 = 3×3 . 5 = 5×5 . Med. = Median. LL = LLMMSE. R = Refined. AN = Adaptive neighborhood.

Both the mean and the variance of the noise PDF above are equal to unity, which lead to a highly noisy image with $SNR = 1$. When applying filtering methods directly on such images, the results are typically of poorer quality than with other types of noise. A common preprocessing step in speckle noise reduction is to obtain several realizations of the same image and to average them, which reduces the power of the noise [186, 187]. In the present study, the filters were applied after averaging four frames of corrupted versions of the same image, which reduced the noise variance to one half of its initial value. The results of filtering the noisy images are shown in Figures 3.59 and 3.60. The MSE values of the images are listed in Tables 3.1 and 3.2. The application of the LLMMSE filter using the adaptive-neighborhood method provided the best result with the Shapes image. With the Peppers image, the NURW filter gave results comparable to those of the adaptive-neighborhood mean and LLMMSE filters.


Salt-and-pepper noise: The test images were corrupted with salt-and-pepper noise such that 5% of the pixels would become outliers. Although only the median filter would be appropriate for filtering salt-and-pepper noise, all of the filters in the comparative study were applied to the noisy image. The results of filtering the noisy images are shown in Figures 3.61 and 3.62. The MSE values of the images are listed in in Tables 3.1 and 3.2.

Discussion: The results presented above indicate that the LLMMSE estimate, especially when computed using the adaptive-neighborhood paradigm or when applied repeatedly, can successfully remove several types of signal-independent and signal-dependent noise. The use of local statistics in an adaptive and nonlinear filter is a powerful approach to remove noise while retaining the edges in the images with minimal distortion. In the case of salt-and-pepper noise, the local median is the most appropriate method, due to the high incidence of outliers that render the use of the local variance inappropriate. Although the use of the 3×3 median as the initial estimate for region growing in the adaptive-neighborhood procedure led to clipping of the corners in some cases, the application of the LLMMSE filter within the same context corrected this distortion to some extent.


The large number of examples provided also illustrate the difficulty in comparing the results provided by several filters. Although the MSE or the RMS error is commonly used for this purpose and has been provided with several illustrations in this chapter, it is seen that, in some cases, an image with a larger MSE may be preferred to one with a lower MSE but with some distortion present. In practical applications, it is important to obtain an assessment of the results by the end-user or by a specialist in the relevant area of application.

**FIGURE 3.53**


(a) Shapes test image. (b) Image in (a) with uniformly distributed noise added, with $\mu = 0$, $\sigma = 20$; $\text{MSE} = 226.20$. Result of filtering the noisy image in (b) using: (c) 3×3 LLMSE; $\text{MSE} = 113.83$. (d) NURW; $\text{MSE} = 144.71$. (e) adaptive-neighborhood mean; $\text{MSE} = 216.52$. (f) adaptive-neighborhood LLMSE; $\text{MSE} = 93.08$. Figure courtesy of M. Ciuc, Laboratorul de Analiza și Prelucrarea Imaginilor, Universitatea Politehnică București, Bucharest, Romania.

**FIGURE 3.54**


(a) Peppers test image. (b) Image in (a) with uniformly distributed noise added, with $\mu = 0$, $\sigma = 20$; $\text{MSE} = 391.43$. Result of filtering the noisy image in (b) using: (c) refined LLMMSE; $\text{MSE} = 65.25$. (d) NURW; $\text{MSE} = 54.39$. (e) adaptive-neighborhood mean; $\text{MSE} = 62.53$. (f) adaptive-neighborhood LLMMSE; $\text{MSE} = 58.62$. Figure courtesy of M. Ciuc, Laboratorul de Analiza și Prelucrarea Imaginilor, Universitatea Politehnica București, Bucharest, Romania.

**FIGURE 3.55**


(a) Shapes test image. (b) Image in (a) with Poisson noise, with $\lambda = 0.1$. MSE = 241.07. Result of filtering the noisy image in (b) using: (c) 3×3 LLMMSE; MSE = 108.70. (d) NURW; MSE = 147.87. (e) adaptive-neighborhood mean; MSE = 215.18. (f) adaptive-neighborhood LLMMSE; MSE = 62.57. Figure courtesy of M. Ciuc, Laboratorul de Analiza și Prelucrarea Imaginilor, Universitatea Politehnica București, Bucharest, Romania.

**FIGURE 3.56**


(a) Peppers test image. (b) Image in (a) with Poisson noise, with $\lambda = 0.1$. MSE = 1132.56. Result of filtering the noisy image in (b) using: (c) refined LLMMSE; MSE = 133.82. (d) NURW; MSE = 85.32. (e) adaptive-neighborhood mean; MSE = 88.83. (f) adaptive-neighborhood LLMMSE; MSE = 87.48. Figure courtesy of M. Ciuc, Laboratorul de Analiza și Prelucrarea Imaginilor, Universitatea Politehnica București, Bucharest, Romania.

**FIGURE 3.57**


(a) Shapes test image. (b) Image in (a) with film-grain noise; $\text{MSE} = 275.11$. Result of filtering the noisy image in (b) using: (c) 3×3 LLMMSE; $\text{MSE} = 119.81$. (d) NURW; $\text{MSE} = 166.42$. (e) adaptive-neighborhood mean; $\text{MSE} = 236.27$. (f) adaptive-neighborhood LLMMSE; $\text{MSE} = 69.98$. Figure courtesy of M. Ciuc, Laboratorul de Analiza și Prelucrarea Imaginilor, Universitatea Politehnica București, Bucharest, Romania.

**FIGURE 3.58**


(a) Peppers test image. (b) Image in (a) with film-grain noise; $MSE = 1233.43$. Result of filtering the noisy image in (b) using: (c) refined LLMMSE; $MSE = 117.54$. (d) NURW; $MSE = 88.83$. (e) adaptive-neighborhood mean; $MSE = 90.54$. (f) adaptive-neighborhood LLMMSE; $MSE = 89.07$. Figure courtesy of M. Ciuc, Laboratorul de Analiza și Prelucrarea Imaginilor, Universitatea Politehnica București, Bucharest, Romania.

**FIGURE 3.59**


(a) Shapes test image. (b) Image in (a) with speckle noise, $MSE = 255.43$. Result of filtering the noisy image in (b) using: (c) 3×3 LLMMSE, $MSE = 119.15$; (d) NURW, $MSE = 116.75$; (e) adaptive-neighborhood mean, $MSE = 236.09$; (f) adaptive-neighborhood LLMMSE, $MSE = 68.01$. Figure courtesy of M. Ciuc, Laboratorul de Analiza și Prelucrarea Imaginilor, Universitatea Politehnica București, Bucharest, Romania.

**FIGURE 3.60**

(a) Peppers test image. (b) Image in (a) with speckle noise, $MSE = 988.84$. Result of filtering the noisy image in (b) using: (c) refined LLMMSE, $MSE = 100.76$; (d) NURW, $MSE = 77.59$; (e) adaptive-neighborhood mean, $MSE = 81.54$; (f) adaptive-neighborhood LLMMSE, $MSE = 79.26$. Figure courtesy of M. Ciuc, Laboratorul de Analiza și Prelucrarea Imaginilor, Universitatea Politehnica București, Bucharest, Romania.

**FIGURE 3.61**

(a) Shapes test image. (b) Image in (a) with salt-and-pepper noise, $MSE = 1740.86$. Result of filtering the noisy image in (b) using: (c) 3×3 mean, $MSE = 642.20$; (d) 3×3 median, $MSE = 206.63$; (e) adaptive-neighborhood mean, $MSE = 213.02$; (f) adaptive-neighborhood median, $MSE = 205.72$. Figure courtesy of M. Ciuc, Laboratorul de Analiza și Prelucrarea Imaginilor, Universitatea Politehnica București, Bucharest, Romania.

**FIGURE 3.62**

(a) Peppers test image. (b) Image in (a) with salt-and-pepper noise, $MSE = 947.69$. Result of filtering the noisy image in (b) using: (c) 5×5 mean, $MSE = 117.02$; (d) 3×3 median, $MSE = 22.93$; (e) adaptive-neighborhood mean, $MSE = 34.49$; (f) adaptive-neighborhood median, $MSE = 29.27$. Figure courtesy of M. Ciuc, Laboratorul de Analiza și Prelucrarea Imaginilor, Universitatea Politehnica București, Bucharest, Romania.

3.9 Application: Multiframe Averaging in Confocal Microscopy

The confocal microscope uses a laser beam to scan and image finely focused planes within fluorescent-dye-tagged specimens that could be a few mm in thickness [216, 217, 218]. The use of a coherent light source obviates the blur caused in imaging with ordinary white light, where the different frequency components of the incident light are reflected and refracted at different angles by the specimen. Laser excitation causes the dyes to emit light (that is, fluoresce) at particular wavelengths. The use of multiple dyes to stain different tissues and structures within the specimen permits their separate and distinct imaging.

The confocal microscope uses a pinhole to permit the passage of only the light from the plane of focus; light from the other planes of the specimen is blocked. Whereas the use of the pinhole permits fine focusing, it also reduces significantly the amount of light that is passed for further detection and viewing. For this reason, a PMT is used to amplify the light received. A scanning mechanism is used to raster-scan the sample in steps that could be as small as $0.1 \mu m$. The confocal microscope facilitates the imaging of multiple focal planes separated by distances of the order of $1 \mu m$; several such slices may be acquired and combined to build 3D images of the specimen.

The use of a laser beam for scanning and imaging carries limitations. The use of high-powered laser beams to obtain strong emitted light could damage the specimen by heating. On the other hand, low laser power levels result in weak emitted light, which, during amplification by the PMT, could suffer from high levels of noise. Scanning with a low-power laser beam over long periods of time to reduce noise could lead to damage of the specimen by photo-bleaching (the affected molecules permanently lose their capability of fluorescence). Images could, in addition, be contaminated by noise due to autofluorescence of the specimen. A technique commonly used to improve image quality in confocal microscopy is to average multiple acquisitions of each scan line or of the full image frame (see Section 3.2).

Figure 3.63 shows images of cells from the nucleus pulposus (the central portion of the intervertebral discs, which are cartilaginous tissues lying between the bony vertebral bodies) [216]. The specimen was scanned using a laser beam of wavelength $488 nm$. The red-dye (long-pass cutoff at $585 nm$) and green-dye (pass band of $505 - 530 nm$) components show distinctly and separately the cell nuclei and the actin filament structure, respectively, of the specimen, in a single focal plane representing a thickness of about $1 \mu m$. The component and composite images would be viewed in the colors mentioned on the microscope, but are illustrated in gray scale in the figure. Images of this nature have been found to be useful in studies of injuries and diseases that affect the intervertebral discs and the spinal column [216].

Figures 3.64 (a) and (b) show two single-frame acquisitions of the composite image as in Figure 3.63 (c). Figures 3.64 (c) and (d) show the results of averaging four and eight single-frame acquisitions of the specimen, respectively. Multiframe averaging has clearly reduced the noise and improved the quality of the image.

Al-Kofahi et al. [219] used confocal microscopy to study the structure of soma and dendrites via 3D tracing of neuronal topology.

Comparison of filtering with space-domain and frequency-domain lowpass filters: Figures 3.65 (a) and (b) show the results of 3×3 mean and median filtering of the single-frame acquisition of the composite image of the nucleus pulposus in Figure 3.64 (a). Figures 3.65 (c) and (d) show the results of 5×5 mean and median filtering of the same image. It is evident that neighborhood filtering, while suppressing noise to some extent, has caused blurring of the sharp details in the images.

Figure 3.66 shows the results of application of the 5×5 LLMMSE, refined LLMMSE, NURW, and adaptive-neighborhood LLMMSE filters to the noisy image in Figure 3.64 (a). The noise was assumed to be multiplicative, with the normalized mean $\mu = 1$ and normalized standard deviation $\sigma = 0.2$. The optimal, adaptive, and nonlinear filters have performed well in suppressing noise without causing blurring.

The results of Fourier-domain ideal and Butterworth lowpass filtering of the noisy image are shown in Figures 3.67 (c) and (d); parts (a) and (b) of the same figure show the original image and its Fourier log-magnitude spectrum, respectively. The spectrum indicates that most of the energy of the image is concentrated within a small area around the center; that is, around $(u, v) = (0, 0)$. The filters have caused some loss of sharpness while suppressing noise. Observe that the ideal lowpass filter has given the noisy areas a mottled appearance. Comparing the results in Figures 3.65, 3.66, and 3.67 with the results of multiframe averaging in Figures 3.64 and 3.63, it is seen that multiframe averaging has provided the best results.

3.10 Application: Noise Reduction in Nuclear Medicine Imaging

Nuclear medicine images are typically acquired under low-photon conditions, which leads to a significant presence of Poisson noise in the images. Counting the photons emitted over long periods of time reduces the effect of noise and improves the quality of the image. However, imaging over long periods of time may not be feasible due to motion artifacts and various practical limitations.

Figure 3.68 shows the SPECT images of one section of a resolution phantom acquired over 2, 15, and 40 s. Each image has been scaled such that its





FIGURE 3.63

(a) The red-dye (cell nuclei) component of the confocal microscope image of the nucleus pulposus of a dog. (b) The green-dye (actin filament structure) component. (c) Combination of the images in (a) and (b) into a composite image. The images would be viewed in the colors mentioned on the microscope. The width of each image corresponds to $145 \mu\text{m}$. Each image was acquired by averaging eight frames. Images courtesy of C.J. Hunter, J.R. Matyas, and N.A. Duncan, McCaig Centre for Joint Injury and Arthritis Research, University of Calgary.

**FIGURE 3.64**

(a) A single-frame acquisition of the composite image of the nucleus pulposus; see also Figure 3.63. (b) A second example of a single-frame acquisition as in (a). (c) The result of averaging four frames including the two in (a) and (b). (d) The result of averaging eight frames including the two in (a) and (b). The width of each image corresponds to $145 \mu\text{m}$. Images courtesy of C.J. Hunter, J.R. Matyas, and N.A. Duncan, McCaig Centre for Joint Injury and Arthritis Research, University of Calgary.





FIGURE 3.65

Results of filtering the single-frame acquisition of the composite image of the nucleus pulposus in Figure 3.64 (a) with: (a) the 3×3 mean filter; (b) the 3×3 median filter; (c) the 5×5 mean filter; and (d) the 5×5 median filter.

**FIGURE 3.66**

Results of filtering the single-frame acquisition of the composite image of the nucleus pulposus in Figure 3.64 (a) with: (a) the 5×5 LLMMSE filter; (b) the refined LLMMSE filter; (c) the NURW filter; and (d) the adaptive-neighborhood LLMMSE filter. Figure courtesy of M. Ciuc, Laboratorul de Analiza și Prelucrarea Imaginilor, Universitatea Politehnica București, Bucharest, Romania.




FIGURE 3.67

(a) The single-frame acquisition of the composite image of the nucleus pulposus of a dog, as in Figure 3.64 (a). (b) Fourier log-magnitude spectrum of the image in (a). Results of filtering the image in (a) with: (c) the ideal lowpass filter with cutoff $D_0 = 0.4$, as in Figure 3.28 (a); and (d) the Butterworth lowpass filter with cutoff $D_0 = 0.4$ and order $n = 2$, as in Figure 3.28 (b).

minimum and maximum values are mapped to the display range of $[0, 255]$. Also shown is a schematic representation of the section: the circles represent cross-sections of cylindrical holes in a plexiglass block; the diameter of the entire phantom is 200 mm ; the two large circles at the extremes of the two sides are of diameter 39 mm ; the two inner arrays have circles of diameter $22, 17, 14, 12, 9, 8, 6,$ and 5 mm . The phantom was filled with a radiopharmaceutical such that the cylindrical holes would be filled with the radioactive material. It is evident that the image quality improves as the photon counting time is increased. The circles of diameter 9 and 8 mm are distinctly visible only in image (c). However, the circles of diameter 5 mm are not visible in any image.

Figure 3.69 shows six nuclear medicine (planar) images of the chest of a patient in the gated blood-pool mode of imaging the heart using ^{99m}Tc . Frame (a) displays the left ventricle in its fully relaxed state (at the end of diastole). The subsequent frames show the left ventricle at various stages of contraction through one cardiac cycle. Frame (d) shows the left ventricle at its smallest, being fully contracted in systole. Frame (f) completes the cycle at a few milliseconds before the end of diastole. Each frame represents the sum of 16 gated frames acquired over 16 cardiac cycles. The ECG signal was used to time photon counting such that each frame gets the photon counts acquired during an interval equal to $\frac{1}{16}$ th of the duration of each heart beat, at exactly the same phase of the cardiac cycle; see Figure 3.70. This procedure is akin to synchronized averaging, with the difference that the procedure may be considered to be integration instead of averaging; the net result is the same except for a scale factor. (Clinical imaging systems do not provide the data corresponding to the intervals of an individual cardiac cycle, but provide only the images integrated over 16 or 32 cardiac cycles.)

3.11 Remarks

In this chapter, we have studied several types of artifacts that could arise in biomedical images, and developed a number of techniques to characterize, model, and remove them. Starting with simple averaging over multiple image frames or over small neighborhoods within an image, we have seen how statistical parameters may be used to filter noise of different types. We have also examined frequency-domain derivation and application of filters. The class of filters based upon mathematical morphology [8, 192, 220, 221, 222] has not been dealt with in this book.

The analysis of the results of several filters has demonstrated the truth behind the adage *prevention is better than cure*: attempts to remove one type of artifact could lead to the introduction of others! Regardless, adaptive and




FIGURE 3.68

128×128 SPECT images of a resolution phantom obtained by counting photons over: (a) 2 s; (b) 15 s; and (c) 40 s. (d) Schematic representation of the section. Images courtesy of L.J. Hahn, Foothills Hospital, Calgary.




FIGURE 3.69

(a) 64×64 gated blood-pool images at six phases of the cardiac cycle, obtained by averaging over 16 cardiac cycles. Images courtesy of L.J. Hahn, Foothills Hospital, Calgary.




FIGURE 3.70

Use of the ECG signal in synchronized averaging or accumulation of photon counts in gated blood-pool imaging. Two cycles of cardiac activity are shown by the ECG signal. The ECG waves have the following connotation: P — atrial contraction; QRS — ventricular contraction (systole); T — ventricular relaxation (diastole). Eight frames representing the gated images are shown over each cardiac cycle. Counts over the same phase of the cardiac cycle are added to the same frame over several cardiac cycles.

nonlinear filters have proven themselves to be useful in removing noise without creating significant artifacts. Preprocessing of images to remove artifact is an important step before other methods may be applied for further enhancement or analysis of the features in the images.

Notwithstanding the models that were used in deriving some of the filters presented in this chapter, most of the methods applied in practice for noise removal are considered to be *ad hoc* approaches: methods that have been shown to work successfully in similar situations encountered by other researchers are tried to solve the problem on hand. Difficulties arise due to the fact that some of the implicit models and assumptions may not apply well to the image or noise processes of the current problem. For this reason, it is common to try several previously established techniques. However, the assessment of the results of filtering operations and the selection of the most appropriate filter for a given application remain a challenge. Although several measures of image quality are available (see Chapter 2), visual assessment of the results by a specialist in the area of application may remain the most viable approach for comparative analysis and selection.

3.12 Study Questions and Problems

(*Note:* Some of the questions may require background preparation with other sources on the basics of signals and systems as well as digital signal and image processing, such as Lathi [1], Oppenheim et al. [2], Oppenheim and Schaffer [7], Gonzalez and Woods [8], Pratt [10], Jain [12], Hall [9], and Rosenfeld and Kak [11].)

Selected data files related to some of the problems and exercises are available at the site

www.enel.ualgary.ca/People/Ranga/enel697

1. Explain the differences between linear and circular (or periodic) convolution. Given two 1D signals $h(n) = [3, 2, 1]$, and $f(n) = [5, 4, 1]$, for $n = 0, 1, 2$, compute by hand the results of linear and circular convolution. Demonstrate how the result of circular convolution may be made equivalent to that of linear convolution in the above example.
2. The two 1D signals $\mathbf{f} = [1, 4, 2, 4]^T$ and $\mathbf{h} = [1, 2, -1]^T$ (given as vectors) are to be convolved. Prepare a circulant matrix \mathbf{h} such that the matrix product $\mathbf{g} = \mathbf{h} \mathbf{f}$ will provide results equivalent to the *linear convolution* $g = h * f$ of the two signals.
3. What are the impulse response and frequency response (transfer function or MTF) of the 3×3 mean filter?
4. An image is processed by applying the 3×3 mean filter mask (a) once, (b) twice, and (c) thrice in series. What are the impulse responses of the three operations?

5. Perform linear 2D convolution of the following two images:

$$\begin{bmatrix} 1 & 2 & 3 \\ 2 & 4 & 1 \\ 2 & 3 & 1 \end{bmatrix} \quad (3.174)$$

and

$$\begin{bmatrix} 2 & 5 & 1 & 2 \\ 2 & 4 & 6 & 3 \\ 5 & 5 & 5 & 1 \\ 4 & 2 & 1 & 2 \end{bmatrix}. \quad (3.175)$$

6. The image

$$\begin{bmatrix} 3 & 2 & 1 & 3 \\ 2 & 1 & 3 & 4 \\ 6 & 5 & 4 & 5 \\ 4 & 3 & 2 & 1 \\ 2 & 1 & 1 & 2 \end{bmatrix} \quad (3.176)$$

is passed through a linear shift-invariant system having the impulse response

$$\begin{bmatrix} 1 & 3 & 2 \\ 0 & 5 & 1 \\ 2 & 4 & 3 \end{bmatrix}. \quad (3.177)$$

Compute the output of the system.

7. The output of a filter at (m, n) is defined as the average of the four immediate neighbors of (m, n) ; the pixel at (m, n) is itself not used.

Derive the MTF of the filter and describe its characteristics.

8. The Fourier transform of a 1D periodic train of impulses, represented as $\sum_{n=-\infty}^{+\infty} \delta(t - nT)$, where T is the period or interval between the impulses, is given by another periodic train of impulses in the frequency domain as $\omega_0 \sum_{n=-\infty}^{+\infty} \delta(\omega - n\omega_0)$, where $\omega_0 = \frac{2\pi}{T}$.

Using the information provided above, derive the 2D Fourier transform of an image made up of a periodic array of strips parallel to the x axis. The thickness of each strip is W , the spacing between the strips is S , and the image is of size $A \times B$, with $\{A, B\} \gg \{W, S\}$.

Draw a schematic sketch of the spectrum of the image.

9. A 3×3 window of a noisy image contains the following pixel values:

$$\begin{bmatrix} 52 & 59 & 41 \\ 62 & 74 & 66 \\ 56 & 57 & 59 \end{bmatrix}. \quad (3.178)$$

Compute the outputs of the 3×3 mean and median filters for the pixel at the center of the window. Show all steps in your computation.

10. A digital image contains periodic grid lines in the horizontal and vertical directions with a spacing of 1 cm. The sampling interval is 1 mm, and the size of the image is 20 cm × 20 cm.

The spectrum of the image is computed using the FFT with an array size of 256 × 256, including zero-padding in the area not covered by the original image.

Sketch a schematic diagram of the spectrum of the image, indicating the nature and exact locations of the frequency components of the grid lines.

Propose a method to remove the grid lines from the image.

11. In deriving the Wiener filter, it is assumed that the processes generating the image \mathbf{f} and noise $\boldsymbol{\eta}$ are statistically independent of each other, that the mean of the noise process is zero, and that both the processes are second-order stationary. A degraded image is observed as $\mathbf{g} = \mathbf{f} + \boldsymbol{\eta}$. The following expression is encountered for the MSE between the Wiener estimate $\tilde{\mathbf{f}} = \mathbf{L}\mathbf{g}$ and the original image \mathbf{f} :

$$\varepsilon^2 = E \left[\text{Tr} \left\{ (\mathbf{f} - \tilde{\mathbf{f}})(\mathbf{f} - \tilde{\mathbf{f}})^T \right\} \right]. \quad (3.179)$$

Reduce the expression above to one containing \mathbf{L} and autocorrelation matrices only. Give reasons for each step of your derivation.

3.13 Laboratory Exercises and Projects

1. Add salt-and-pepper noise to the image `shapest.tif`. Apply the median filter using the neighborhood shapes given in Figure 3.14, and compare the results in terms of MSE values and edge distortion.
2. From your collection of test images, select two images: one with strong edges of the objects or features present in the image, and the other with smooth edges and features.

Prepare several noisy versions of the images by adding

- (i) Gaussian noise, and
- (ii) salt-and-pepper noise

at various levels.

Filter the noisy images using

- (i) the median filter with the neighborhoods given in Figure 3.14 (a), (b), and (k); and
- (ii) the 3 × 3 mean filter with the condition that the filter is applied only if the difference between the pixel being processed and the average of its 8-connected neighbors is less than a threshold. Try different thresholds and study the effect.

Compare the results in terms of noise removal, MSE, and the effect of the filters on the edges present in the images.

3. Select two of the noisy images from the preceding exercise. Apply the ideal lowpass filter and the Butterworth lowpass filter using two different cutoff frequencies for each filter.

Study the results in terms of noise removal and the effect of the filters on the sharpness of the edges present in the images.

Image Enhancement

In spite of the significant advances made in biomedical imaging techniques over the past few decades, several practical factors often lead to the acquisition of images with less than the desired levels of contrast, visibility of detail, or overall quality. In the preceding chapters, we reviewed several practical limitations, considerations, and trade-offs that could lead to poor images. When the nature of the artifact that led to the poor quality of the image is known, such as noise as explained in Chapter 3, we may design specific methods to remove or reduce the artifact. When the degradation is due to a blur function, deblurring and restoration techniques, described in Chapter 10, may be applied to reverse the phenomenon. In some applications of biomedical imaging, it becomes possible to include additional steps or modifications in the imaging procedure to improve image quality, although at additional radiation dose to the subject in the case of some X-ray imaging procedures, as we shall see in the sections to follow.

In several situations, the understanding of the exact cause of the loss of quality is limited or nonexistent, and the investigator is forced to attempt to improve or enhance the quality of the image on hand using several techniques applied in an *ad hoc* manner. In some applications, a nonspecific improvement in the general appearance of the given image may suffice. Researchers in the field of image processing have developed a large repertoire of image enhancement techniques that have been demonstrated to work well under certain conditions with certain types of images. Some of the enhancement techniques, indeed, have an underlying philosophy or hypothesis, as we shall see in the following sections; however, the practical application of the techniques may encounter difficulties due to a mismatch between the applicable conditions or assumptions and those that relate to the problem on hand.

A few biomedical imaging situations and applications where enhancement of the feature of interest would be desirable are:

- Microcalcifications in mammograms.
- Lung nodules in chest X-ray images.
- Vascular structure of the brain.
- Hair-line fractures in the ribs.

Some of the features listed above could be difficult to see in the given image due to their small size, subtlety, small differences in characteristics with respect to their surrounding structures, or low contrast; others could be rendered not readily visible due to superimposed structures in planar images. Enhancement of the contrast, edges, and general detail visibility in the images, without causing any distortion or artifacts, would be desirable in the applications mentioned above.

In this chapter, we shall explore a wide range of image enhancement techniques that can lead to improved contrast or visibility of certain image features such as edges or objects of specific characteristics. In extending the techniques to other applications, it should be borne in mind that ad hoc procedures borrowed from other areas may not lead to the best possible or optimal results. Regardless, if the improvement so gained is substantial and consistent as judged by the users and experts in the domain of application, one may have on hand a practically useful technique. (See the July 1972 and May 1979 issues of the *Proceedings of the IEEE* for reviews and articles on digital image processing, including historically significant images.)

4.1 Digital Subtraction Angiography

In digital subtraction angiography (DSA), an X-ray contrast agent (such as an iodine compound) is injected so as to increase the density (attenuation coefficient) of the blood within a certain organ or system of interest. A number of X-ray images are taken as the contrast agent spreads through the arterial network and before the agent is dispersed via circulation throughout the body. An image taken before the injection of the agent is used as the “mask” or reference image, and subtracted from the “live” images obtained with the agent in the system to obtain enhanced images of the arterial system of interest.

Imaging systems that perform contrast-enhanced X-ray imaging (without subtraction) in a motion or cine mode are known as cine-angiography systems. Such systems are useful in studying circulation through the coronary system to detect sclerosis (narrowing or blockage of arteries due to the deposition of cholesterol, calcium, and other substances).

Figures 4.1 (a), (b), and (c) show the mask, live, and the result of DSA, respectively, illustrating the arterial structure in the brain of a subject [223, 224, 225]. The arteries are barely visible in the live image [Figure 4.1 (b)], in spite of the contrast agent. Subtraction of the skull and the other parts that have remained unchanged between the mask and the live images has resulted in greatly improved visualization of the arteries in the DSA image [Figure 4.1 (c)]. The mathematical procedure involved may be expressed simply as

$$\mathbf{f} = \alpha \mathbf{f}_1 - \beta \mathbf{f}_2, \text{ or}$$

$$f(m, n) = \alpha f_1(m, n) - \beta f_2(m, n), \quad (4.1)$$

where f_1 is the live image, f_2 is the mask image, α and β are weighting factors (if required), and f is the result of DSA.

The simple mathematical operation of subtraction (on a pixel-by-pixel basis) has, indeed, a significant application in medical imaging. The technique, however, is sensitive to motion, which causes misalignment of the components to be subtracted. The DSA result in Figure 4.1 (c) demonstrates motion artifacts in the lowest quarter and around the periphery of the image. Methods to minimize motion artifact in DSA have been proposed by Meijering et al. [223, 224, 225]. Figure 4.1 (d) shows the DSA result after correction of motion artifacts. Regardless of its simplicity, DSA carries a certain risk of allergic reaction, infection, and occasionally death, due to the injection of the contrast agent.

4.2 Dual-energy and Energy-subtraction X-ray Imaging

Different materials have varying energy-dependent X-ray attenuation coefficients. X-ray measurements or images obtained at multiple energy levels (also known as energy-selective imaging) could be combined to derive information about the distribution of specific materials in the object or body imaged. Weighted combinations of multiple-energy images may be obtained to display soft-tissue and hard-tissue details separately [5]. The disadvantages of dual-energy imaging exist in the need to subject the patient to two or more X-ray exposures (at different energy or kV). Furthermore, due to the time lapse between the exposures, motion artifacts could arise in the resulting image.

In a variation of the dual-energy method, MacMahon [226, 227] describes energy-subtraction imaging using a dual-plate CR system. The Fuji FCR 9501ES (Fujifilm Medical Systems USA, Stamford, CT) digital chest unit uses two receptor plates instead of one. The plates are separated by a copper filter. The first plate acquires the full-spectrum X-ray image in the usual manner. The copper filter passes only the high-energy components of the X rays on to the second plate. Because bones and calcium-containing structures would have preferentially absorbed the low-energy components of the X rays, and because the high-energy components would have passed through low-density tissues with little attenuation, the transmitted high-energy components could be expected to contain more information related to denser tissues than to lighter tissues. The two plates capture two different views derived from the same X-ray beam; the patient is not subjected to two different imaging exposures, but only one. Weighted subtraction of the two images as in Equation 4.1 provides various results that can demonstrate soft tissues or bones and calcified tissues in enhanced detail; see Figures 4.2 and 4.3.




FIGURE 4.1

(a) Mask image of the head of a patient for DSA. (b) Live image. (c) DSA image of the cerebral artery network. (d) DSA image after correction of motion artifacts. Image data courtesy of E.H.W. Meijering and M.A. Viergever, Image Sciences Institute, University Medical Center Utrecht, Utrecht, The Netherlands. Reproduced with permission from E.H.W. Meijering, K.J. Zuiderveld, and M.A. Viergever, "Image registration for digital subtraction angiography", *International Journal of Computer Vision*, 31(2/3): 227 – 246, 1999. © Kluwer Academic Publishers.

Energy-subtraction imaging as above has been found to be useful in detecting fracture of the ribs, in assessing the presence of calcification in lung nodules (which would indicate that they are benign, and hence, need not be examined further or treated), and in detecting calcified pleural plaques due to prolonged exposure to asbestos [226, 227]. The bone-detail image in Figure 4.3 (a) shows, in enhanced detail, a small calcified granuloma near the lower-right corner of the image.



FIGURE 4.2

Full-spectrum PA chest image (CR) of a patient. See also Figure 4.3. Image courtesy of H. MacMahon, University of Chicago, Chicago, IL. Reproduced with permission from H. MacMahon, “Improvement in detection of pulmonary nodules: Digital image processing and computer-aided diagnosis”, *RadioGraphics*, 20(4): 1169–1171, 2000. © RSNA.




FIGURE 4.3

(a) Bone-detail image, and (b) soft-tissue detail image obtained by energy subtraction. See also Figure 4.2. Images courtesy of H. MacMahon, University of Chicago, Chicago, IL. Reproduced with permission from H. MacMahon, “Improvement in detection of pulmonary nodules: Digital image processing and computer-aided diagnosis”, *RadioGraphics*, 20(4): 1169–1171, 2000. © RSNA.

4.3 Temporal Subtraction

Temporal or time-lapse subtraction of images could be useful in detecting normal or pathological changes that have occurred over a period of time. MacMahon [226] describes and illustrates the use of temporal subtraction in the detection of lung nodules that could be difficult to see in planar chest images due to superimposed structures. DR and CR imaging facilitate temporal subtraction.

In temporal subtraction, it is desired that normal anatomic structures are suppressed and pathological changes are enhanced. Registration of the images is crucial in temporal subtraction; misregistration could lead to artifacts similar to those due to motion in DSA. Geometric transformation and warping techniques are useful in matching landmark features that are not expected to have changed in the interval between the two imaging sessions [223, 224, 225]. Mazur et al. [228] describe image correlation and geometric transformation techniques for the registration of radiographs for temporal subtraction.

4.4 Gray-scale Transforms

The gray-level histogram of an image gives a global impression of the presence of different levels of density or intensity in the image over the dynamic range available (see Section 2.7 for details and illustrations). When the pixels in a given image do not make full use of the available dynamic range, the histogram will indicate low levels of occurrences of certain gray-level values or ranges. The given image may also contain large areas representing objects with certain specific ranges of gray level; the histogram will then indicate large populations of pixels occupying the corresponding gray-level ranges. Based upon a study of the histogram of an image, we could design gray-scale transforms or look-up tables (LUTs) that alter the overall appearance of the image, and could improve the visibility of selected details.

4.4.1 Gray-scale thresholding

When the gray levels of the objects of interest in an image are known, or can be determined from the histogram of the given image, the image may be thresholded to obtain a variety of images that can display selected features of interest. For example, if it is known that the objects of interest in the image have gray-level values greater than L_1 , we could create an image for display

as

$$g(m, n) = \begin{cases} 0 & \text{if } f(m, n) \leq L_1 \\ 255 & \text{if } f(m, n) \geq L_1 \end{cases}, \quad (4.2)$$

where $f(m, n)$ is the original image; $g(m, n)$ is the thresholded image to be displayed; and the display range is $[0, 255]$. The result is a bilevel or binary image. Thresholding may be considered to be a form of image enhancement in the sense that the objects of interest are perceived better in the resulting image. The same operation may also be considered to be a detection operation; see Section 5.1.

If the values less than L_1 were to be considered as noise (or features of no interest), and the gray levels within the objects of interest that are greater than L_1 are of interest in the displayed image, we could also define the output image as

$$g(m, n) = \begin{cases} 0 & \text{if } f(m, n) \leq L_1 \\ f(m, n) & \text{if } f(m, n) \geq L_1 \end{cases}. \quad (4.3)$$

The resulting image will display the features of interest including their gray-level variations.

Methods for the derivation of optimal thresholds are described in Sections 5.4.1, 8.3.2, and 8.7.2.


Example: A CT slice image of a patient with neuroblastoma is shown in Figure 4.4 (a). A binarized version of the image, with thresholding as in Equation 4.2 using $L_1 = 200 \text{ HU}$, is shown in part (b) of the figure. As expected, the bony parts of the image appear in the result; however, the calcified parts of the tumor, which also have high density comparable to that of bone, appear in the result. The result of thresholding the image as in Equation 4.3 with $L_1 = 200 \text{ HU}$ is shown in part (c) of the figure. The relative intensities of the hard bone and the calcified parts of the tumor are evident in the result.

4.4.2 Gray-scale windowing

If a given image $f(m, n)$ has all of its pixel values in a narrow range of gray levels, or if certain details of particular interest within the image occupy a narrow range of gray levels, it would be desirable to stretch the range of interest to the full range of display available. In the absence of reason to employ a nonlinear transformation, a linear transformation as follows could be used for this purpose:

$$g(m, n) = \begin{cases} 0 & \text{if } f(m, n) \leq f_1 \\ \frac{f(m, n) - f_1}{f_2 - f_1} & \text{if } f_1 < f(m, n) < f_2 \\ 1 & \text{if } f(m, n) \geq f_2 \end{cases}, \quad (4.4)$$

where $f(m, n)$ is the original image; $g(m, n)$ is the windowed image to be displayed, with its gray-scale normalized to the range $[0, 1]$; and $[f_1, f_2]$ is the range of the original gray-level values to be displayed in the output after

**FIGURE 4.4**

(a) CT image of a patient with neuroblastoma. The tumor, which appears as a large circular region on the left-hand side of the image, includes calcified tissues that appear as bright regions. The HU range of $[-200, 400]$ has been linearly mapped to the display range of $[0, 255]$; see also Figures 2.15 and 2.16. Image courtesy of Alberta Children's Hospital, Calgary. (b) The image in (a) thresholded at the level of $200 HU$ as in Equation 4.2. Values above $200 HU$ appear as white, and values below this threshold appear as black. (c) The image in (a) thresholded at the level of $200 HU$ as in Equation 4.3. Values above $200 HU$ appear at their original level, and values below this threshold appear as black. (d) The HU range of $[0, 400]$ has been linearly mapped to the display range of $[0, 255]$ as in Equation 4.4. Pixels corresponding to tissues lighter than water appear as black. Pixels greater than $400 HU$ are saturated at the maximum gray level of 255.

stretching to the full range. Note that the range $[0, 1]$ in the result needs to be mapped to the display range available, such as $[0, 255]$, which is achieved by simply multiplying the normalized values by 255. Details (pixels) below the lower limit f_1 will be eliminated (rendered black) and those above the upper limit f_2 will be saturated (rendered white) in the resulting image. The details within the range $[f_1, f_2]$ will be displayed with increased contrast and latitude, utilizing the full range of display available.

Example: A CT slice image of a patient with neuroblastoma is shown in Figure 4.4 (a). This image displays the range of $[-200, 400]$ HU linearly mapped to the display range of $[0, 255]$ as given by Equation 4.4. The full range of HU values in the image is $[-1000, 1042]$ HU . Part (d) of the figure shows another display of the same original data, but with mapping of the range $[0, 400]$ HU to $[0, 255]$ as given by Equation 4.4. In this result, pixels corresponding to tissues lighter than water appear as black; pixels greater than 400 HU are saturated at the maximum gray level of 255. Gray-level thresholding and mapping are commonly used for detailed interpretation of CT images.


Example: Figure 4.5 (a) shows a part of the chest X-ray image in Figure 1.11 (b), downsampled to 512×512 pixels. The histogram of the image is shown in Figure 4.6 (a); observe the large number of pixels with the gray level zero. Figure 4.6 (b) shows two linear gray-scale transformations (LUTs) that map the range $[0, 0.6]$ (dash-dot line) and $[0.2, 0.7]$ (solid line) to the range $[0, 1]$; the results of application of the two LUTs to the image in Figure 4.5 (a) are shown in Figures 4.5 (b) and (c), respectively. The image in Figure 4.5 (b) shows the details in and around the heart with enhanced visibility; however, large portions of the original image have been saturated. The image in Figure 4.5 (c) provides an improved visualization of a larger range of tissues than the image in (b); regardless, the details with normalized gray levels less than 0.2 and greater than 0.7 have been lost.

Example: Figure 4.7 (a) shows an image of a myocyte. Figure 4.8 (a) shows the normalized histogram of the image. Most of the pixels in the image have gray levels within the limited range of $[50, 150]$; the remainder of the available range $[0, 255]$ is not used effectively.

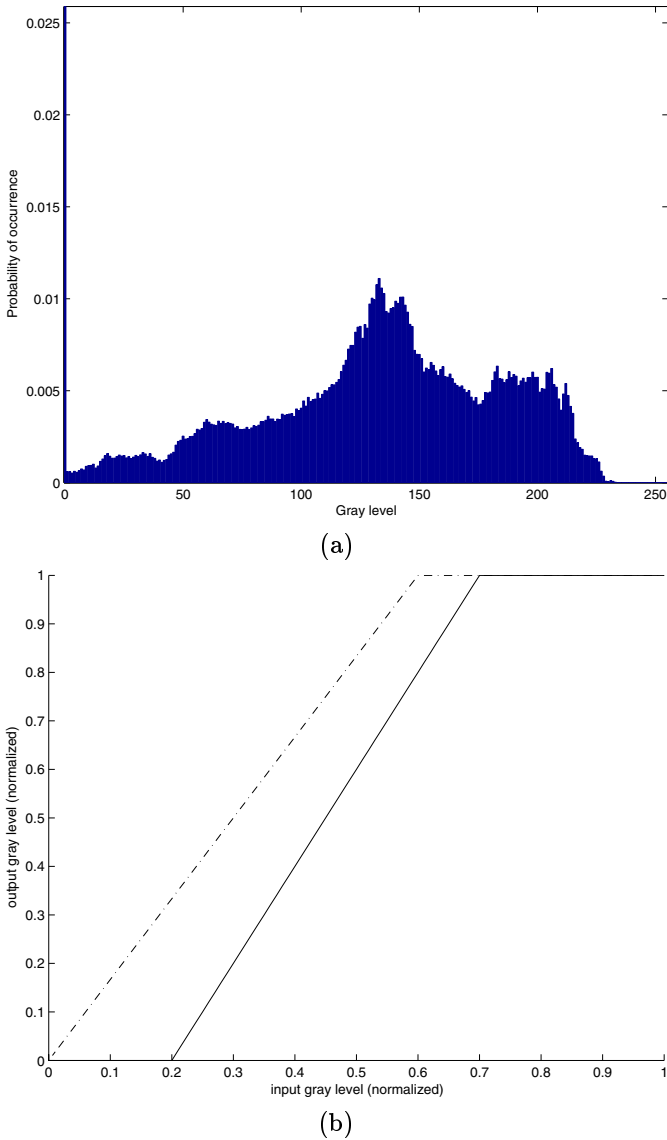
Figure 4.7 (b) shows the image in (a) after the normalized gray-level range of $[0.2, 0.6]$ was stretched to the full range of $[0, 1]$ by the linear transformation in Equation 4.4. The details within the myocyte are visible with enhanced clarity in the transformed image. The corresponding histogram in Figure 4.8 (b) shows that the image now occupies the full range of gray scale available; however, several gray levels within the range are unoccupied, as indicated by the white stripes in the histogram.

4.4.3 Gamma correction


Figure 2.6 shows the H-D curves of two devices. The slope of the curve is known as γ . An imaging system with a large γ could lead to an image with

**FIGURE 4.5**

(a) Part of a chest X-ray image. The histogram of the image is shown in Figure 4.6 (a). (b) Image in (a) enhanced by linear mapping of the range $[0, 0.6]$ to $[0, 1]$. (c) Image in (a) enhanced by linear mapping of the range $[0.2, 0.7]$ to $[0, 1]$. See Figure 4.6 (b) for plots of the LUTs.

**FIGURE 4.6**

(a) Normalized histogram of the chest X-ray image in Figure 4.5 (a); entropy = 7.55 *bits*. (b) Linear density-windowing transformations that map the ranges $[0, 0.6]$ to $[0, 1]$ (dash-dot line) and $[0.2, 0.7]$ to $[0, 1]$ (solid line).

**FIGURE 4.7**

(a) Image of a myocyte as acquired originally. (b) Image in (a) enhanced by linear mapping of the normalized range $[0.2, 0.6]$ to $[0, 1]$. See Figure 4.8 for the histograms of the images.

high contrast; however, the image may not utilize the full range of the available gray scale. On the other hand, a system with a small γ could result in an image with wide latitude but poor contrast. Gamma correction is a nonlinear transformation process by which we may alter the transition from one gray level to the next, and change the contrast and latitude of gray scale in the image. The transformation may be expressed as [203]



$$g(m, n) = [f(m, n)]^\gamma, \quad (4.5)$$

where $f(m, n)$ is the given image with its gray scale normalized to the range $[0, 1]$, and $g(m, n)$ is the transformed image. (*Note:* Lindley [229] provides a different definition as


$$g(m, n) = \exp \left[\frac{\ln \{f(m, n)\}}{\gamma} \right], \quad (4.6)$$

which would be equivalent to the operation given by Equation 4.5 if the gray levels were not normalized, that is, the gray levels were to remain in a range such as $0 - 255$.) Gray-scale windowing as in Equation 4.4 could also be incorporated into Equation 4.5.


Example: Figure 4.9 (a) shows a part of a chest X-ray image. Figure 4.10 illustrates three transforms with $\gamma = 0.3, 1.0$, and 2.0 . Parts (b) and (c) of Figure 4.9 show the results of gamma correction with $\gamma = 0.3$ and $\gamma = 2.0$, respectively. The two results demonstrate enhanced visibility of details in the darker and lighter gray-scale regions (with reference to the original image).

**FIGURE 4.8**

Normalized histograms of (a) the image in Figure 4.7 (a), entropy = 4.96 bits; and (b) the image in Figure 4.7 (b), entropy = 4.49 bits.

**FIGURE 4.9**

(a) Part of a chest X-ray image. (b) Image in (a) enhanced with $\gamma = 0.3$. (c) Image in (a) enhanced with $\gamma = 2.0$. See Figure 4.10 for plots of the gamma-correction transforms (LUTs).

**FIGURE 4.10**

Gamma-correction transforms with $\gamma = 0.3$ (solid line), $\gamma = 1.0$ (dotted line), and $\gamma = 2.0$ (dash-dot line).

4.5 Histogram Transformation

As we saw in Section 2.7, the histogram of an image may be normalized and interpreted as a PDF. Then, based upon certain principles of information theory, we reach the property that maximal information is conveyed when the PDF of a process is uniform, that is, the corresponding image has all possible gray levels with equal probability of occurrence (see Section 2.8). Based upon this property, the technique of histogram equalization has been proposed as a method to enhance the appearance of an image [9, 8, 11]. Other techniques have also been proposed to map the histogram of the given image into a different “desired” type of histogram, with the expectation that the transformed image so obtained will bear an enhanced appearance. Although the methods often do not yield useful results in biomedical applications, and although the underlying assumptions may not be applicable in many practical situations, histogram-based methods for image enhancement are popular. The following sections provide the details and results of a few such methods.

4.5.1 Histogram equalization

Consider an image $f(m, n)$ of size $M \times N$ pixels, with gray levels $l = 0, 1, 2, \dots, L-1$. Let the histogram of the image be represented by $P_f(l)$ as defined in Equation 2.12. Let us normalize the gray levels by dividing by the maximum level available or permitted, as $r = \frac{l}{L-1}$, such that $0 \leq r \leq 1$. Let $p_f(r)$ be the normalized histogram or PDF as given by Equation 2.15.

If we were to apply a transformation $s = T(r)$ to the random variable r , the PDF of the new variable s is given by [8]

$$p_g(s) = p_f(r) \left. \frac{dr}{ds} \right|_{r=T^{-1}(s)}, \quad (4.7)$$

where g refers to the resulting image $g(m, n)$ with the normalized gray levels $0 \leq s \leq 1$. Consider the transformation

$$s = T(r) = \int_0^r p_f(w) dw; \quad 0 \leq r \leq 1. \quad (4.8)$$

This is the cumulative (probability) distribution function of r . $T(r)$ has the following important and desired properties:

- $T(r)$ is single-valued and monotonically increasing over the interval $0 \leq r \leq 1$. This is necessary to maintain the black-to-white transition order between the original and processed images.
- $0 \leq T(r) \leq 1$ for $0 \leq r \leq 1$. This is required in order to maintain the same range of values in the input and output images.

It follows that $\frac{ds}{dr} = p_f(r)$. Then, we have

$$p_g(s) = \left[p_f(r) \frac{1}{p_f(r)} \right]_{r=T^{-1}(s)} = 1; \quad 0 \leq s \leq 1. \quad (4.9)$$

Thus, $T(r)$ equalizes the histogram of the given image; that is, the histogram or PDF of the resulting image $g(m, n)$ is uniform. As we saw in Section 2.8, a uniform PDF has maximal entropy.

Discrete version of histogram equalization: For a digital image $f(m, n)$ with a total of $P = MN$ pixels and L gray levels $r_k, k = 0, 1, \dots, L - 1, 0 \leq r_k \leq 1$, occurring n_k times, respectively, the PDF may be approximated by the histogram

$$p_f(r_k) = \frac{n_k}{P}; \quad k = 0, 1, \dots, L - 1. \quad (4.10)$$

The histogram-equalizing transformation is approximated by

$$s_k = T(r_k) = \sum_{i=0}^k p_f(r_i) = \sum_{i=0}^k \frac{n_i}{P}; \quad k = 0, 1, \dots, L - 1. \quad (4.11)$$

Note that this transformation may yield values of s_k that may not equal the available quantized gray levels. The values will have to be quantized, and hence the output image may only have an approximately uniform histogram.

In practical applications, the resulting values in the range $[0, 1]$ have to be scaled to the display range, such as $[0, 255]$. Histogram equalization is usually implemented via an LUT that lists the related (s_k, r_k) pairs as given by Equation 4.11. It should be noted that a quantized histogram-equalizing transformation is likely to contain several segments of many-to-one gray-level transformation: this renders the transformation nonunique and irreversible.

Example: Figure 4.11 (a) shows a 240×288 image of a girl in a snow cave: the high reflectivity of the snow has caused the details inside the cave to have poor visibility. Part (b) of the same figure shows the result after histogram equalization; the histograms of the original and equalized images are shown in Figure 4.12. Although the result of equalization shows some of the features of the girl within the cave better than the original, several details remain dark and unclear.

The histogram of the equalized image in Figure 4.12 (b) indicates that, while a large number of gray levels have higher probabilities of occurrence than their corresponding levels in the original [see Figure 4.12 (a)], several gray levels are unoccupied in the enhanced image (observe the white stripes in the histogram, which indicate zero probability of occurrence of the corresponding gray levels). The equalizing transform (LUT), shown in Figure 4.13, indicates that there are several many-to-one gray-level mappings: note the presence of several horizontal segments in the LUT. It should also be observed that the original image has a well-spread histogram, with an entropy of 6.93 bits; due to the absence of several gray levels in the equalized image, its entropy of 5.8 bits turns out to be lower than that of the original.

Figure 4.11 (c) shows the result of linear stretching or windowing of the range $[0, 23]$ in the original image in Figure 4.11 (a) to the full range of $[0, 255]$. The result shows the details of the girl and the inside of the cave more clearly than the original or the equalized version; however, the high-intensity details outside the cave have been washed out.

Figure 4.11 (d) shows the result of enhancing the original image in Figure 4.11 (a) with $\gamma = 0.3$. Although the details inside the cave are not as clearly seen as in Figure 4.11 (c), the result has maintained the details at all gray levels.






FIGURE 4.11

(a) Image of a girl in a snow cave (240×288 pixels). (b) Result of histogram equalization. (c) Result of linear mapping (windowing) of the range $[0, 23]$ to $[0, 255]$. (d) Result of gamma correction with $\gamma = 0.3$. Image courtesy of W.M. Morrow [215, 230].

Example: Figure 4.14 (a) shows a part of a chest X-ray image; part (b) of the same figure shows the corresponding histogram-equalized image. Al-



(a)



(b)

FIGURE 4.12

Normalized histograms of (a) the image in Figure 4.11 (a), entropy = 6.93 *bits*; and (b) the image in Figure 4.11 (b), entropy = 5.8 *bits*. See also Figure 4.13.




FIGURE 4.13

Histogram-equalizing transform (LUT) for the image in Figure 4.11 (a); see Figure 4.12 for the histograms of the original and equalized images.

though some parts of the image demonstrate improved visibility of features, it should be observed that the low-density tissues in the lower right-hand portion of the image have been reduced to poor levels of visibility. The histogram of the equalized image is shown in Figure 4.15 (a); the equalizing transform is shown in part (b) of the same figure. It is seen that several gray levels are unoccupied in the equalized image; for this reason, the entropy of the enhanced image was reduced to 5.95 *bits* from the value of 7.55 *bits* for the original image.

Example: Figure 4.16 (b) shows the histogram-equalized version of the myocyte image in Figure 4.16 (a). The corresponding equalizing transform, shown in Figure 4.17 (b), indicates a sharp transition from the darker gray levels to the brighter gray levels. The rapid transition has caused the output to have high contrast over a small effective dynamic range, and has rendered the result useless. The entropies of the original and enhanced images are 4.96 *bits* and 4.49 *bits*, respectively.

4.5.2 Histogram specification

A major limitation of histogram equalization is that it can provide only one output image, which may not be satisfactory in many cases. The user has




FIGURE 4.14

(a) Part of a chest X-ray image. The histogram of the image is shown in Figure 4.6 (a). (b) Image in (a) enhanced by histogram equalization. The histogram of the image is shown in Figure 4.15 (a). See Figure 4.15 (b) for a plot of the LUT.

no control over the procedure or the result. In a related procedure known as histogram specification, a series of histogram-equalization steps is used to obtain an image with a histogram that is expected to be close to a *prespecified* histogram. Then, by specifying several histograms, it is possible to obtain a range of enhanced images, from which one or more may be selected for further analysis or use.

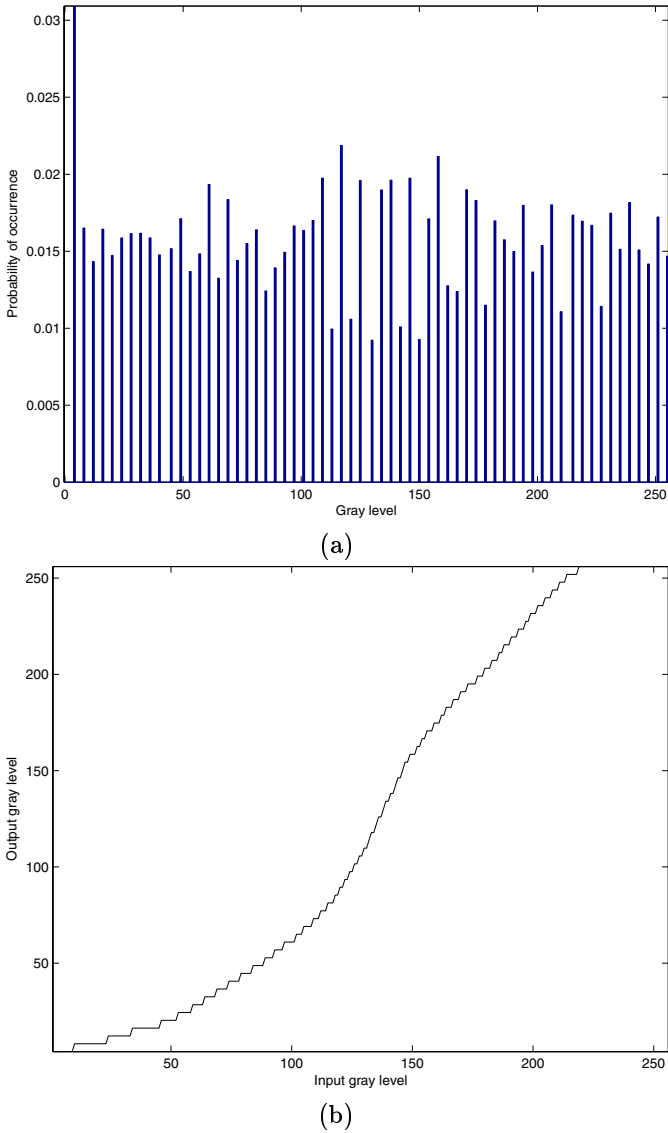
Suppose that the desired or specified normalized histogram is $p_d(t)$, with the desired image being represented as d , having the normalized gray levels $t = 0, 1, 2, \dots, L - 1$. Now, the given image f with the PDF $p_f(r)$ may be histogram-equalized by the transformation

$$s = T_1(r) = \int_0^r p_f(w) dw; \quad 0 \leq r \leq 1, \quad (4.12)$$


as we saw in Section 4.5.1, to obtain the image g with the normalized gray levels s . We may also derive a histogram-equalizing transform for the desired (but as yet unavailable) image as

$$q = T_2(t) = \int_0^t p_d(w) dw; \quad 0 \leq t \leq 1. \quad (4.13)$$

Observe that, in order to derive a histogram-equalizing transform, we need only the PDF of the image; the image itself is not needed. Let us call the (hypothetical) image so obtained as e , having the gray levels q . The inverse

**FIGURE 4.15**

(a) Normalized histogram of the histogram-equalized chest X-ray image in Figure 4.14 (b); entropy = 5.95 *bits*. (b) The histogram-equalizing transformation (LUT). See Figure 4.6 (a) for the histogram of the original image.

**FIGURE 4.16**


(a) Image of a myocyte. The histogram of the image is shown in Figure 4.8 (a). (b) Image in (a) enhanced by histogram equalization. The histogram of the image is shown in Figure 4.17 (a). See Figure 4.17 (b) for a plot of the LUT.

of the transform above, which we may express as $t = T_2^{-1}(q)$, will map the gray levels q back to t .


Now, $p_g(s)$ and $p_e(q)$ are both uniform PDFs, and hence are identical functions. The desired PDF may, therefore, be obtained by applying the transform T_2^{-1} to s ; that is, $t = T_2^{-1}(s)$. It is assumed here that $T_2^{-1}(s)$ exists, and is a single-valued (unique) transform. Based on the above, the procedure for histogram specification is as follows:

1. Specify the desired histogram and derive the equivalent PDF $p_d(t)$.
2. Derive the histogram-equalizing transform $q = T_2(t)$.
3. Derive the histogram-equalizing transform $s = T_1(r)$ from the PDF $p_f(r)$ of the given image f .
4. Apply the inverse of the transform T_2 to the PDF obtained in the previous step and obtain $t = T_2^{-1}(s)$. This step may be directly implemented as $t = T_2^{-1}[T_1(r)]$.
5. Apply the transform as above to the given image f ; the result provides the desired image d with the specified PDF $p_d(t)$.

Although the procedure given above can theoretically lead us to an image having the specified histogram, the method faces limitations in practice. Difficulty arises in the very first step of specifying a meaningful histogram or



(a)



(b)

FIGURE 4.17

(a) Normalized histogram of the histogram-equalized myocyte image in Figure 4.16 (b). (b) The histogram-equalizing transformation (LUT). See Figure 4.8 (a) for the histogram of the original image.

PDF; several trials may be required before a usable image is obtained. More importantly, in a practical implementation with discrete gray levels, it will be difficult, if not impossible, to derive the inverse transform T_2^{-1} . The possible existence of many-to-one mapping segments in the histogram-equalizing transform T_2 , as we saw in the examples in Section 4.5.1, may render inversion impossible. Appropriate specification of the desired PDF could facilitate the design of an LUT to approximately represent T_2^{-1} . The LUTs corresponding to T_1 and T_2^{-1} may be combined into one LUT that may be applied to the given image f to obtain the desired image d in a single step. Note that the image obtained as above may have a histogram that only approximates the one specified.

4.5.3 Limitations of global operations

Global operators such as gray-scale and histogram transforms provide simple mechanisms to manipulate the appearance of images. Some knowledge about the range of gray levels of the features of interest can assist in the design of linear or nonlinear LUTs for the enhancement of selected features in a given image. Although histogram equalization can lead to useful results in some situations, it is quite common to result in poor images. Even if we keep aside the limitations related to nonunique transforms, a global approach to image enhancement ignores the nonstationary nature of images, and hence could lead to poor results. The results of histogram equalization of the chest X-ray and myocyte images in Figures 4.14 and 4.16 demonstrate the limitations of global transforms. Given the wide range of details of interest in medical images, such as the hard tissues (bone) and soft tissues (lung) in a chest X-ray image, it is desirable to design local and adaptive transforms for effective image enhancement.

4.5.4 Local-area histogram equalization

Global histogram equalization tends to result in images where features having gray levels with low probabilities of occurrence in the original image are merged upon quantization of the equalizing transform, and hence are lost in the enhanced image. Ketchum [231] attempted to address this problem by suggesting the application of histogram equalization on a local basis. In local-area histogram equalization (LAHE), the histogram of the pixels within a 2D sliding rectangular window, centered at the current pixel being processed, is equalized, and the resulting transform is applied only to the central pixel; the process is repeated for every pixel in the image. The window provides the local context for the pixel being processed. The method is computationally expensive because a new transform needs to be computed for every pixel.

Pizer et al. [232], Leszczynski and Shalev [233], and Rehm and Dallas [234] proposed variations of LAHE, and extended the method to the enhancement of medical images. In one of the variations of LAHE, the histogram-equalizing

transforms are computed not for every pixel, but only for a number of nonoverlapping rectangular blocks spanning the image. The pixels at the center of each block are processed using the corresponding transform. Pixels that are not at the centers of the blocks are processed using interpolated versions of the transforms corresponding to the four neighboring center pixels. The success of LAHE depends upon the appropriate choice of the size of the sliding window in relation to the sizes of the objects present in the image, and of the corresponding background areas.

Example: The images in Figures 4.18 (c) and (d) show the results of application of the LAHE method to the image in part (a) of the figure, using windows of size 11×11 and 101×101 pixels, respectively. The result of global histogram equalization is shown in part (b) of the figure for comparison. Although the results of LAHE provide improved visualization of some of the details within the snow cave, the method has led to gray-level inversion in a few regions (black patches in white snow areas); this effect is due to the spreading of the gray levels in a small region over the full range of $[0, 255]$, which is not applicable to all local areas in a given image. The overall quality of the results of LAHE has been downgraded by this effect.

4.5.5 Adaptive-neighborhood histogram equalization

A limitation of LAHE lies in the use of rectangular windows: although such a window provides the local context of the pixel being processed, there is no apparent justification to the choice of the rectangular shape for the moving window. Furthermore, the success of the method depends significantly upon proper choice of the size of the window; the use of a fixed window of a prespecified size over an entire image has no particular reasoning.

Paranjape et al. [230] proposed an adaptive-neighborhood approach to histogram equalization. As we saw in Section 3.7.5, the adaptive-neighborhood image processing paradigm is based upon the identification of variable-shape, variable-size neighborhoods for each pixel by region growing. Because the region-growing procedure used for adaptive-neighborhood image processing leads to a relatively uniform region, with gray-level variations limited to that permitted by the specified threshold, the local histogram of such a region will tend to span a limited range of gray levels. Equalizing such a histogram and permitting the occurrence of the entire range of gray levels in any and every local context is inappropriate. In order to provide an increased context to histogram equalization, Paranjape et al. included in the local area not only the foreground region grown, but also a background composed of a ribbon of pixels molded to the foreground; see Figure 3.46. The extent of the local context provided depends upon the tolerance specified for region growing, the width of the background ribbon of pixels, and the nature of gray-level variability present in the given image. The method adapts to local details present in the given image; regions of different size and shape are grown for each pixel.




FIGURE 4.18

(a) Image of a girl in a snow cave (240×288 pixels). (b) Result of global histogram equalization. Results of LAHE with (c) a 11×11 window and (d) a 101×101 window. Results of adaptive-neighborhood histogram equalization with (e) growth tolerance 16 and background width 5 pixels, and (f) growth tolerance 64 and background width 8 pixels. Reproduced with permission from R.B. Paranjape, W.M. Morrow, and R.M. Rangayyan, "Adaptive-neighborhood histogram equalization for image enhancement", *CVGIP: Graphical Models and Image Processing*, 54(3):259–267, 1992. © Academic Press.

After obtaining the histogram of the local region, the equalizing transform is derived, and applied only to the seed pixel from where the process was started. The same value is applied to all redundant seed pixels in the region; that is, to the pixels that have the same gray-level value as the seed (for which the same region would have been grown using a simple tolerance).

In an extension of adaptive-neighborhood histogram equalization to color images proposed by Ciuc et al. [235], instead of equalizing the local histogram, an adaptive histogram stretching operation is applied to the local histograms. The enhancement operation is applied only to the intensity of the image; undesired changes to the color balance (hue) are prevented by this method.

Example: Figure 4.19 shows a simple test image with square objects of different gray levels, as well as its enhanced versions using global, local-area, and adaptive-neighborhood histogram equalization. The limitations of global histogram equalization are apparent in the fact that the brighter, inner square on the right-hand side of the image remains almost invisible. The result of LAHE permits improved visualization of the inner squares; however, the artifacts due to block-wise processing are obvious and disturbing. Adaptive-neighborhood histogram equalization has provided the best result, with enhanced visibility of the inner squares and without any artifacts.




FIGURE 4.19

(a) A test image and its enhanced versions by: (b) global or full-frame histogram equalization, (c) LAHE, and (d) adaptive-neighborhood histogram equalization. Image courtesy of R.B. Paranjape.

Example: The images in Figures 4.18 (e) and (f) show the results of application of the adaptive-neighborhood histogram equalization method to the image in part (a) of the figure. The two images were obtained using growth tolerance values of 16 and 64, and background width of 5 and 8 pixels. The larger tolerance and larger background width provide for larger areas of the local context to be included in the local histogram. The result of global histogram equalization is shown in part (b) of the figure for comparison. The results of adaptive-neighborhood histogram equalization provide improved visualization of details and image features both inside and outside the snow cave. Furthermore, the result with the larger growth tolerance and background ribbon width is relatively free of the gray-level inversion (black patches in otherwise white areas) present in the results of LAHE, shown in parts (c) and (d) of the same figure.

4.6 Convolution Mask Operators

Filtering images using 3×3 convolution masks is a popular approach. Several such masks have been proposed and are in practical use for image enhancement. Equation 3.39 demonstrates the use of a simple 3×3 mask to represent the local mean filter. We shall explore a few other 3×3 convolution masks for image enhancement in the following sections.

4.6.1 Unsharp masking

When an image is blurred by some unknown phenomenon, we could assume that each pixel in the original image contributes, in an additive manner, a certain fraction of its value to the neighboring pixels. Then, each pixel is composed of its own true value, plus fractional components of its neighbors. The spreading of the value of a pixel into its neighborhood may be viewed as the development of a local fog or blurred background.

In an established photographic technique known as unsharp masking, the given degraded image, in its negative form, is first blurred, and a positive transparency is created from the result. The original negative and the positive are held together, and a (positive) print is made of the combination. The procedure leads to the subtraction of the local blur or fog component, and hence to an improved and sharper image.

A popular 3×3 convolution mask that mimics unsharp masking is given by

$$\begin{bmatrix} -\frac{1}{8} & -\frac{1}{8} & -\frac{1}{8} \\ -\frac{1}{8} & 2 & -\frac{1}{8} \\ -\frac{1}{8} & -\frac{1}{8} & -\frac{1}{8} \end{bmatrix}. \quad (4.14)$$

Observe that the net sum of the values in the mask equals unity; therefore, there is no net change in the local average intensity.

The operation above may be generalized to permit the use of other local window sizes and shapes as

$$f_e(m, n) = [g(m, n) - \mu_g(m, n)] + \alpha g(m, n). \quad (4.15)$$

This expression indicates that the pixel at the location (m, n) in the enhanced image $f_e(m, n)$ is given as a weighted combination of the corresponding pixel $g(m, n)$ in the given degraded image, and the difference between the pixel and the local mean $\mu_g(m, n)$. The expression is equivalent to the mask in Equation 4.14, with $\alpha = 1$ and the local mean being computed as the average of the eight neighbors of the pixel being processed. Note that because the mask possesses symmetry about both the x and y axes, reversal has no effect, and hence is not required, in performing convolution.

The relative weighting between the pixel being processed and the local difference could be modified depending upon the nature of the image and the desired effect, leading to various values at the central location in the mask given in Equation 4.14. Equivalently, different values of α could be used in Equation 4.15. Because the local difference in Equation 4.15 is a measure of the local gradient, and because gradients are associated with edges, combining the given image with its local gradient could be expected to lead to edge enhancement or high-frequency emphasis.

Example: Figure 4.20 (a) shows a test image of a clock; part (b) of the same figure shows the result of unsharp masking using the 3×3 mask in Equation 4.14. It is evident that the details in the image, such as the numerals, have been sharpened by the operation. However, it is also seen that the high-frequency emphasis property of the filter has led to increased noise in the image.

Figures 4.21 (a), 4.22 (a), 4.23 (a), and 4.24 (a) show the image of a myocyte, a part of a chest X-ray image, an MR image of a knee, and the Shapes test image; the results of enhancement obtained by the unsharp masking operator are shown in parts (b) of the same figures. The chest image, in particular, has been enhanced well by the operation: details of the lungs in the dark region in the lower-right quadrant of the image are seen better in the enhanced image than in the original.

An important point to observe from the result of enhancement of the Shapes test image is that the unsharp masking filter performs *edge enhancement*. Fur-

thermore, strong edges will have a clearly perceptible overshoot and undershoot; this could be considered to be a form of ringing artifact. The images in Figure 4.25 illustrate the artifact in an enlarged format. Although the artifact is not as strongly evident in the other test images, the effect is, indeed, present. Radiologists often do not prefer edge enhancement, possibly for this reason.

Note that the unsharp masking operation could lead to negative pixel values in the enhanced image; the user has to decide how to handle this aspect when displaying the result. The illustrations in this section were prepared by linearly mapping selected ranges of the results to the display range of $[0, 255]$, as stated in the figure captions; compression of the larger dynamic range in the enhanced image to a smaller display range could mute the effect of enhancement to some extent.

4.6.2 Subtracting Laplacian

Under certain conditions, a degraded image g may be modeled as being the result of a diffusion process that spreads intensity values over space as a function of time, according to the partial differential equation [11]

$$\frac{\partial g}{\partial t} = \kappa \nabla^2 g, \quad (4.16)$$

where t represents time, $\kappa > 0$ is a constant, and

$$\nabla^2 g = \frac{\partial^2 g}{\partial x^2} + \frac{\partial^2 g}{\partial y^2}. \quad (4.17)$$

In the initial state at $t = 0$, we have $g(x, y, 0) = f(x, y)$, the original image. At some time instant $t = \tau > 0$, the degraded image $g(x, y, \tau)$ is observed. The degraded image may be expressed in a Taylor series as

$$g(x, y, \tau) = g(x, y, 0) + \tau \frac{\partial g}{\partial t}(x, y, \tau) - \frac{\tau^2}{2} \frac{\partial^2 g}{\partial t^2}(x, y, \tau) + \dots \quad (4.18)$$


Ignoring the quadratic and higher-order terms, letting $g(x, y, 0) = f(x, y)$, and using the diffusion model in Equation 4.16, we get

$$f_e = g - \kappa \tau \nabla^2 g, \quad (4.19)$$


where f_e represents an approximation to f . Thus, we have an enhanced image obtained as a weighted subtraction of the given image and its Laplacian (gradient).

A discrete implementation of the Laplacian is given by the 3×3 convolution mask


$$\begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix}; \quad (4.20)$$

**FIGURE 4.20**

(a) Clock test image. (b) Result of unsharp masking; display range $[-50, 250]$ out of $[-68, 287]$. (c) Laplacian (gradient) of the image; display range $[-50, 50]$ out of $[-354, 184]$. (d) Result of the subtracting Laplacian; display range $[-50, 250]$ out of $[-184, 250]$.

**FIGURE 4.21**

(a) Image of a myocyte; the range from the minimum to the maximum of the image has been linearly mapped to the display range $[0, 255]$. (b) Result of unsharp masking; display range $[-20, 180]$ out of $[-47, 201]$. (c) Laplacian (gradient) of the image; display range $[-20, 20]$ out of $[-152, 130]$. (d) Result of the subtracting Laplacian; display range $[-50, 200]$ out of $[-130, 282]$.

**FIGURE 4.22**

(a) Part of a chest X-ray image. (b) Result of unsharp masking; display range $[-30, 230]$ out of $[-59, 264]$. (c) Laplacian (gradient) of the image; display range $[-5, 5]$ out of $[-134, 156]$. (d) Result of the subtracting Laplacian; display range $[-50, 250]$ out of $[-156, 328]$.





FIGURE 4.23

(a) MR image of a knee. (b) Result of unsharp masking; display range $[-40, 250]$ out of $[-72, 353]$. (c) Laplacian (gradient) of the image; display range $[-50, 50]$ out of $[-302, 365]$. (d) Result of the subtracting Laplacian; display range $[-50, 250]$ out of $[-261, 549]$.

**FIGURE 4.24**

(a) Shapes test image. (b) Result of unsharp masking; display range $[-100, 250]$ out of $[-130, 414]$. See also Figure 4.25. (c) Laplacian (gradient) of the image; display range $[-50, 50]$ out of $[-624, 532]$. (d) Result of the subtracting Laplacian; display range $[-300, 300]$ out of $[-532, 832]$.




FIGURE 4.25

Enlarged views of a part of (a) the Shapes test image and (b) the result of unsharp masking; see also Figure 4.24 (a) and (b). Observe the edge-enhancement artifact.

see also Equation 2.82 and the associated discussion. Observe that the net weight of the coefficients in the Laplacian mask is zero; therefore, the mask performs a differentiation operation that will lead to the loss of intensity information (that is, the result in an area of any uniform brightness value will be zero).

Letting the weighting factor $\kappa \tau = 1$ in Equation 4.19, we get the following 3×3 mask known as the subtracting Laplacian:

$$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix}. \quad (4.21)$$

Because the net weight of the mask is equal to unity, the mask retains the local average intensity in the image.

Comparing Equations 4.21 and 4.14, we see that they have a similar structure, the main difference being in the number of the neighboring pixels used in computing the local gradient or difference. For this reason, the unsharp masking filter is referred to as the generalized (subtracting) Laplacian by some authors. On the same note, the subtracting Laplacian is also an unsharp masking filter. For the same reasons as in the case of the unsharp masking filter, the subtracting Laplacian also leads to edge enhancement or high-frequency emphasis; see also Equation 2.82 and the associated discussion.

Example: Part (c) of Figure 4.20 shows the Laplacian of the test image in part (a) of the same figure. The Laplacian shows large values (positive or

negative) at the strong edges that are present in the image. Part (d) of the figure shows the result of the subtracting Laplacian, which demonstrates the edge-enhancing property of the filter.


Figures 4.21 (c), 4.22 (c), 4.23 (c), and 4.24 (c) show the Laplacian of the corresponding images in parts (a) of the same figures. Parts (d) of the figures show the results of the subtracting Laplacian operator. The subtracting Laplacian has provided higher levels of sharpening than the unsharp masking filter in most cases; the result is also noisier in the case of the Clock test image.

Observe that the Laplacian does not maintain the intensity information present in the image, whereas the subtracting Laplacian does maintain this information; the former results in a depiction of the edges (gradient) present in the image, whereas the latter provides a sharper image. As in the case of unsharp masking, the subtracting Laplacian could lead to negative pixel values in the enhanced image; the user has to decide how to handle this aspect when displaying the result. The illustrations in this section were prepared by linearly mapping selected ranges of the results to the display range of $[0, 255]$, as stated in the figure captions; compression of the larger dynamic range in the enhanced image to a smaller display range could mute the effect of enhancement to some extent, and also alter the intensity values of parts of the image.

Similar to the artifact introduced by the unsharp-masking operator as illustrated in Figure 4.25, the subtracting Laplacian could also introduce disturbing overshoot and undershoot artifacts around edges; see Figure 4.24 (d). This characteristic of the operator is illustrated using a 1D signal in Figure 4.26. Such artifacts could affect the quality and acceptance of images enhanced using the subtracting Laplacian.

4.6.3 Limitations of fixed operators

Fixed operators, such as the unsharp-masking and subtracting-Laplacian filters, apply the same mathematical operation at every location over the entire space of the given image. The coefficients and the size of such filters do not vary, and hence the filters cannot adapt to changes in the nature of the image from one location to another. For these reasons, fixed operators may encounter limited success in enhancing large images with complex and space-variant features. In medical images, we encounter a wide latitude of details; for example, in a chest X-ray image, we see soft-tissue patterns in the lungs and hard-tissue structures such as ribs. Similar changes in density may be of concern in one anatomical region or structure, but not in another. The spatial scale of the details of diagnostic interest could also vary significantly from one part of an image to another; for example, from fine blood vessels or bronchial tubes to large bones such as the ribs in chest X-ray images. Operators with fixed coefficients and fixed spatial scope of effect cannot take these factors into

**FIGURE 4.26**

Top to bottom: a rectangular pulse signal smoothed with a Gaussian blur function; the first derivative of the signal; the second derivative of the signal; and the result of a filter equivalent to the subtracting Laplacian. The derivatives are shown with enlarged amplitude scales as compared to the original and filtered signals.

consideration. Adaptive filters and operators are often desirable to address these concerns.

4.7 High-frequency Emphasis

Highpass filters are useful in detecting edges, under the assumption that high-frequency Fourier spectral components are associated with edges and large changes in the image. This property follows from the effect of differentiation of an image on its Fourier transform, as expressed by Equation 2.75.

The ideal highpass filter: The *ideal* highpass filter is defined in the 2D Fourier space as

$$H(u, v) = \begin{cases} 1 & \text{if } D(u, v) \geq D_0 \\ 0 & \text{otherwise.} \end{cases} \quad (4.22)$$

where $D(u, v) = \sqrt{u^2 + v^2}$ is the distance of the frequency component at (u, v) from the DC point $(u, v) = (0, 0)$, with the spectrum being centered such that the DC component is at its center (see Figures 2.26, 2.27, and 2.28). D_0 is the cutoff frequency, below which all components of the Fourier transform of the given image are set to zero. Figure 4.27 (a) shows the ideal highpass filter function; Figure 4.28 shows the profile of the filter.

The Butterworth highpass filter: As we saw in the case of lowpass filters (in Section 3.4.1), prevention of the ringing artifacts encountered with the ideal filter requires that the transition from the stopband to the passband be smooth. The Butterworth filter response is monotonic in the passband as well as in the stopband. (See Rangayyan [31] for details and illustrations of the 1D Butterworth filter.)

In 2D, the Butterworth highpass filter is defined as [8]

$$H(u, v) = \frac{1}{1 + (\sqrt{2} - 1) \left[\frac{D_0}{D(u, v)} \right]^{2n}}, \quad (4.23)$$

where n is the order of the filter, $D(u, v) = \sqrt{u^2 + v^2}$, and D_0 is the half-power 2D radial cutoff frequency [the scale factor in the denominator leads to the gain of the filter being $\frac{1}{\sqrt{2}}$ at $D(u, v) = D_0$]. The filter's transition from the stopband to the passband becomes steeper as the order n is increased. Figure 4.27 (b) illustrates the magnitude (gain) of the Butterworth highpass filter with the normalized cutoff $D_0 = 0.2$ and order $n = 2$. Figure 4.28 shows the profile of the filter.

Because the gain of a highpass filter is zero at DC, the intensity information is removed by the filter. This leads to a result that depicts only the edges present in the image. Furthermore, the result will have positive and negative




FIGURE 4.27

(a) The magnitude transfer function of an ideal highpass filter. The cutoff frequency D_0 is 0.2 times the maximum frequency. (b) The magnitude transfer function of a Butterworth highpass filter, with normalized cutoff $D_0 = 0.2$ and order $n = 2$. The $(u, v) = (0, 0)$ point is at the center. Black represents a gain of zero, and white represents a gain of unity. See also Figure 4.28.


values. If the enhancement rather than the extraction of edges is desired, it is necessary to maintain the intensity information. This effect could be achieved by using a high-emphasis filter, defined simply as a highpass filter plus a constant in the (u, v) space. The Butterworth high-emphasis filter may be specified as

$$H(u, v) = \kappa_1 + \frac{\kappa_2}{1 + (\sqrt{2} - 1) \left[\frac{D_0}{D(u, v)} \right]^{2n}}, \quad (4.24)$$

which is similar to the Butterworth highpass filter in Equation 4.23 except for the addition of the factors κ_1 and κ_2 .

The high-emphasis filter has a nonzero gain at DC. High-frequency components are emphasized with respect to the low-frequency components in the image; however, the low-frequency components are not removed entirely. Figure 4.28 shows the profile of the Butterworth high-emphasis filter, with $\kappa_1 = 0.5$, $\kappa_2 = 1.0$, $D_0 = 0.2$, and $n = 2$.

Examples: Figure 4.29 (a) shows a test image of a clock; part (b) of the same figure shows the result of the ideal highpass filter. Although the edges in the image have been extracted by the filter, the strong presence of ringing artifacts diminishes the value of the result. Part (c) of the figure shows the result of the Butterworth highpass filter, where the edges are seen without the ringing artifact. The result of the Butterworth high-emphasis filter, shown in

**FIGURE 4.28**

Profiles of the magnitude transfer functions of an ideal highpass filter (solid line), a Butterworth highpass filter (dash-dot line, normalized cutoff $D_0 = 0.2$ and order $n = 2$), and a Butterworth high-emphasis filter (dashed line). See also Figure 4.27.

part (d) of the figure, demonstrates edge enhancement; however, the relative intensities of the objects have been altered.

Figures 4.30 (a), 4.31 (a), 4.32 (a), and 4.33 (a) show the image of a myocyte, a part of a chest X-ray image, an MR image of a knee, and the Shapes test image, respectively. The results of the ideal highpass filter, Butterworth highpass filter, and Butterworth high-emphasis filter are shown in parts (b), (c), and (d), respectively, of the same figures. The distinction between edge enhancement and edge extraction is demonstrated by the examples.

4.8 Homomorphic Filtering for Enhancement


We have studied several linear filters designed to separate images that were added together. The question asked has been, given $g(x, y) = f(x, y) + \eta(x, y)$, how could one extract $f(x, y)$ only? Given that the Fourier transform is linear, we know that the Fourier transforms of the images as above are also combined in an additive manner: $G(u, v) = F(u, v) + \eta(u, v)$. Therefore, a linear filter will facilitate the separation of $F(u, v)$ and $\eta(u, v)$, with the assumption that they have significant portions of their energies in different frequency bands.

Suppose now that we are presented with an image that contains the product of two images, such as $g(x, y) = f(x, y) s(x, y)$. From the multiplication or convolution property of the Fourier transform we have $G(u, v) = F(u, v) * S(u, v)$, where $*$ represents 2D convolution in the frequency domain. How would we be able to separate $f(x, y)$ from $s(x, y)$?


Furthermore, suppose we have $g(x, y) = h(x, y) * f(x, y)$, where $*$ stands for 2D convolution, as in the case of the passage of the original image $f(x, y)$ through an LSI system or filter with the impulse response $h(x, y)$. The Fourier transforms of the signals are related as $G(u, v) = H(u, v) F(u, v)$. How could we attempt to separate $f(x, y)$ and $h(x, y)$?

4.8.1 Generalized linear filtering


Given that linear filters are well established and understood, it is attractive to consider extending their application to images that have been combined by operations other than addition, especially by multiplication and convolution as indicated in the preceding paragraphs. An interesting possibility to achieve this is via conversion of the operation combining the images into addition by one or more transforms. Under the assumption that the transformed images occupy different portions of the transform space, linear filters may be applied to separate them. The inverses of the transforms used initially would then take us back to the original space of the images. This approach was proposed in a series of papers by Bogert et al. [236] and Oppenheim et al. [237, 238]; see

**FIGURE 4.29**

(a) Clock test image. Result of (b) the ideal highpass filter, display range $[-50, 50]$ out of $[-79, 113]$; (c) the Butterworth highpass filter, display range $[-40, 60]$ out of $[-76, 115]$; and (d) the Butterworth high-emphasis filter, display range $[-40, 160]$ out of $[-76, 204]$.

**FIGURE 4.30**

(a) Image of a myocyte; the range from the minimum to the maximum of the image has been linearly mapped to the display range $[0, 255]$. Result of (b) the ideal highpass filter, display range $[-20, 20]$ out of $[-60, 65]$; (c) the Butterworth highpass filter, display range $[-20, 20]$ out of $[-61, 61]$; and (d) the Butterworth high-emphasis filter, display range $[-20, 100]$ out of $[-52, 138]$.


**FIGURE 4.31**

(a) Part of a chest X-ray image. Result of (b) the ideal highpass filter, display range $[-5, 5]$ out of $[-74, 91]$; (c) the Butterworth highpass filter, display range $[-5, 5]$ out of $[-78, 95]$; and (d) the Butterworth high-emphasis filter, display range $[-50, 130]$ out of $[-78, 192]$.



FIGURE 4.32

(a) MR image of a knee. Result of (b) the ideal highpass filter, display range $[-50, 50]$ out of $[-117, 127]$; (c) the Butterworth highpass filter, display range $[-50, 50]$ out of $[-126, 139]$; and (d) the Butterworth high-emphasis filter, display range $[-30, 150]$ out of $[-78, 267]$.

**FIGURE 4.33**

(a) Shapes test image. Result of (b) the ideal highpass filter, display range $[-100, 100]$ out of $[-154, 183]$; (c) the Butterworth highpass filter, display range $[-100, 100]$ out of $[-147, 176]$; and (d) the Butterworth high-emphasis filter, display range $[-100, 200]$ out of $[-147, 296]$.

also Childers et al. [239] and Rangayyan [31] for details and illustrations of application to biomedical signals. Because the procedure extends the application of linear filters to multiplied and convolved images, it has been referred to as *generalized linear filtering*. Furthermore, as the operations can be represented by algebraically linear transformations between the input and output vector spaces, they have been called *homomorphic systems*.

As a simple illustration of a homomorphic system for multiplied images, consider again the image

$$g(x, y) = f(x, y) s(x, y). \quad (4.25)$$

Given the goal of converting the multiplication operation to addition, it is evident that a simple logarithmic transformation is appropriate:

$$\log[g(x, y)] = \log[f(x, y) s(x, y)] = \log[f(x, y)] + \log[s(x, y)], \quad (4.26)$$

$f(x, y) \neq 0$, $s(x, y) \neq 0 \forall (x, y)$. The logarithms of the two images are now combined in an additive manner. Taking the Fourier transform, we get

$$G_l(u, v) = F_l(u, v) + S_l(u, v), \quad (4.27)$$

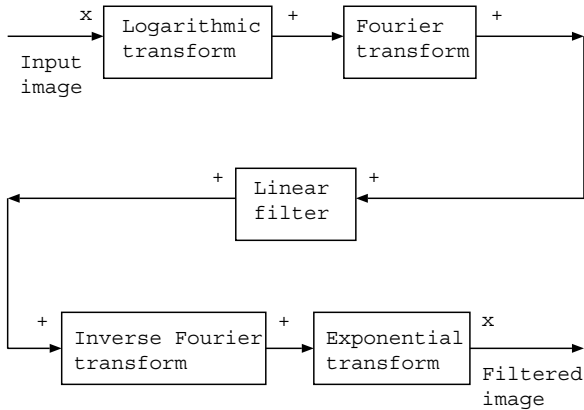
where the subscript l indicates that the Fourier transform has been applied to a log-transformed version of the image.

Assuming that the logarithmic transformation has not affected the separability of the Fourier components of the two images $f(x, y)$ and $s(x, y)$, a linear filter (lowpass, highpass, etc.) may now be applied to $G_l(u, v)$ to separate them. An inverse Fourier transform will yield the filtered image in the space domain. An exponential operation will complete the reversal procedure. This procedure was proposed by Stockham [240] for image processing in the context of a visual model.

Figure 4.34 illustrates the operations involved in a multiplicative homomorphic system (or filter). The symbol at the input or output of each block indicates the operation that combines the image components at the corresponding step. A system of this nature is useful in image enhancement, where an image may be treated as the product of an illumination function and a transmittance or reflectance function. The homomorphic filter facilitates the separation of the illumination function and correction for nonuniform lighting. The method has been used to achieve simultaneous dynamic range compression and contrast enhancement [7, 8, 237, 240].

The extension of homomorphic filtering to separate convolved signals is described in Section 10.3.

Example: The test image in Figure 4.35 (a) shows a girl inside a snowcave. The intensity of illumination of the scene differs significantly between the outside and the inside of the snowcave. Although there is high contrast between the outside and the inside of the snowcave, there is poor contrast of the details within the snowcave. Because the image possesses a large dynamic

**FIGURE 4.34**

Homomorphic filtering for enhancement of images combined by multiplication.

range, linear stretching of the gray-level range of the full image is not viable. (However, a part of the range may be stretched to the full range, as illustrated in Figure 4.11.)

Figure 4.35 (b) shows the result of logarithmic transformation of the image in part (a) of the figure. Although the girl is now visible, the image is not sharp. The image was filtered using a Butterworth high-emphasis filter, as illustrated in Figure 4.36, within the context of the homomorphic system shown in Figure 4.34. The filter was specified as in Equation 4.24, with $\kappa_1 = 0.1$, $\kappa_2 = 0.5$, $D_0 = 0.6$, and $n = 1$. The result, shown in Figure 4.35 (c), demonstrates reduced dynamic range in terms of the difference in illumination between the inside and the outside of the snowcave, but increased contrast and sharpness of the details within the snowcave. Application of the Butterworth high-emphasis filter without the homomorphic system resulted in the image in Figure 4.35 (d), which does not present the same level of enhancement as seen in Figure 4.35 (c).

Example: A part of a mammogram containing calcifications is shown in Figure 4.37 (a). The multiplicative model of an illuminated scene does not apply to X-ray imaging; however, the image has nonuniform brightness (density) that affects the visibility of details in the darker regions, and could benefit from homomorphic enhancement. Figure 4.37 (b) shows the result of logarithmic transformation of the image in part (a) of the figure; the result of filtering using a Butterworth high-emphasis filter is shown in part (c). The log operation has improved the visibility of the calcifications in the dark region in the upper-central part of the image (arranged along an almost-vertical linear pattern); application of the Butterworth high-emphasis filter (illustrated in Figure 4.36) has further sharpened these features. The result [Figure 4.37 (c)],




FIGURE 4.35

(a) Test image of a girl in a snowcave. Result of (b) log transformation; (c) homomorphic filtering including a Butterworth high-emphasis filter; and (d) the Butterworth high-emphasis filter only. The test image in this illustration is of size 256×256 pixels, and is slightly different from that in Figures 4.11 and 4.18; regardless, comparison of the results indicates the advantages of homomorphic filtering. The Butterworth high-emphasis filter used is shown in Figure 4.36. Image courtesy of W.M. Morrow [215, 230].




FIGURE 4.36

Profile of the high-emphasis Butterworth filter used to enhance high-frequency components along with homomorphic filtering as illustrated in Figures 4.34, 4.35, and 4.37.

however, does not depict the distinction between high-density tissues (bright areas) and low-density tissues (dark areas).

The result of application of the Butterworth high-emphasis filter without the homomorphic system is shown in Figure 4.37 (d). This operation has also resulted in improved depiction of the calcifications in the dark regions, albeit not to the same extent as within the context of the homomorphic procedure.

Yoon et al. [241] extended the application of homomorphic high-emphasis filtering to the wavelet domain for contrast enhancement of mammograms.

4.9 Adaptive Contrast Enhancement


Diagnostic features in medical images, such as mammograms, vary widely in size and shape. Classical image enhancement techniques cannot adapt to the varying characteristics of such features. The application of a global transform or a fixed operator to an entire image often yields poor results in at least some parts of the given image. It is, therefore, necessary to design methods that can adapt the operation performed or the pixel collection used to derive measures to the local details present in the image. The following section provides the details of an adaptive-neighborhood approach to contrast enhancement of images.

4.9.1 Adaptive-neighborhood contrast enhancement

Morrow et al. [123, 215] proposed an adaptive-neighborhood contrast enhancement technique for application to mammograms. As we saw in Section 3.7.5, in adaptive-neighborhood or region-based image processing, an adaptive neighborhood is defined about each pixel in the image, the extent of which is dependent on the characteristics of the image feature in which the pixel being processed is situated. This neighborhood of similar pixels is called an adaptive neighborhood or region.

Note that in image segmentation, groups of pixels are found that have some property in common (such as similar gray level), and are used to define disjoint image regions called segments. Region-based processing may be performed by initially segmenting the given image and then processing each segment in turn. Alternatively, for region-based processing, we may define possibly overlapping regions for each pixel, and process each of the regions independently.

Regions, if properly defined, should correspond to image features. Then, features in the image are processed as whole units, rather than pixels being processed using arbitrary groups of neighboring pixels (for example, 3×3 masks). Region-based processing could also be designated as pixel-independent

**FIGURE 4.37**

(a) Original image of a part of mammogram with malignant calcifications. Result of (b) log transformation; (c) homomorphic filtering including a Butterworth high-emphasis filter; and (d) the Butterworth high-emphasis filter only. See also Figures 4.40 and 4.41.

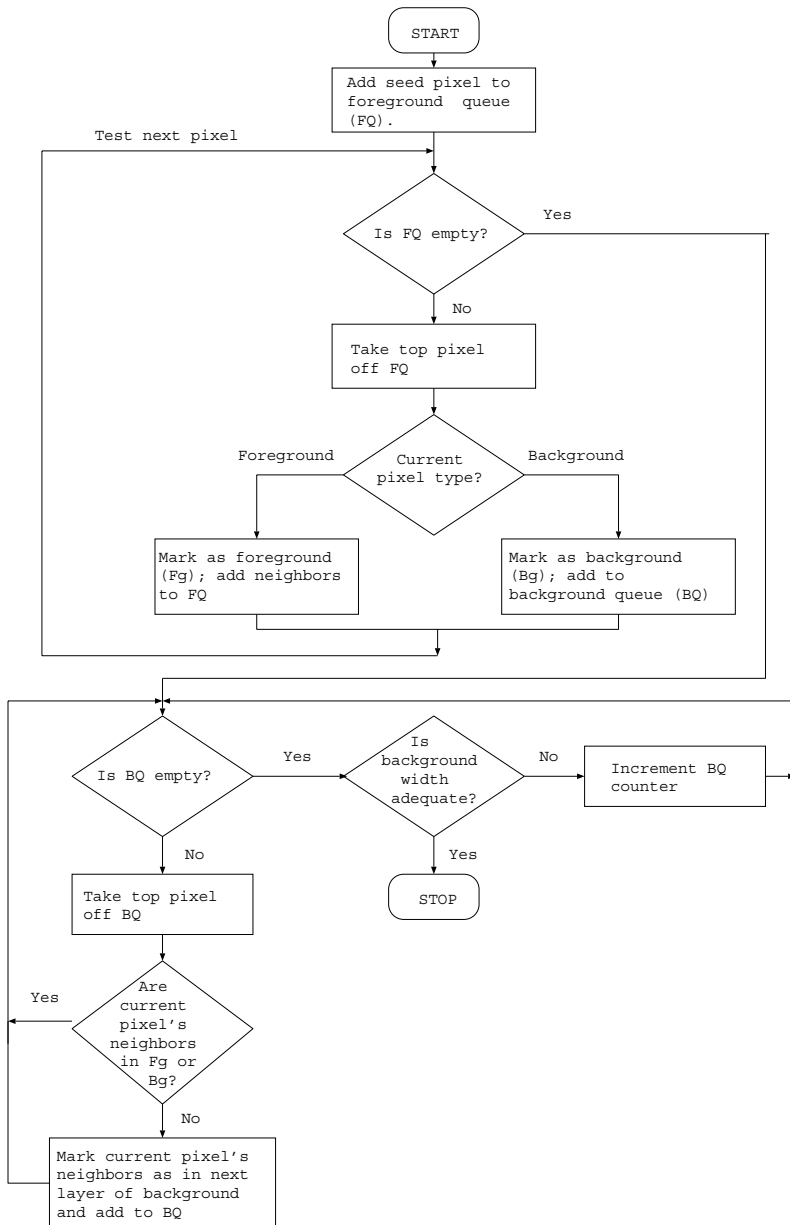
processing [242, 243, 244], feature-based processing, adaptive-neighborhood processing, or object-oriented processing.

The fundamental step in adaptive-neighborhood image processing is defining the extent of regions in the image. Overlapping regions are used in this application because disjoint segmentation of an image, with subsequent enhancement of the segments, would result in noticeable edge artifacts and an inferior enhanced image.

Seed-fill region growing: Morrow et al. [123, 215] used a region-growing technique based on a simple graphical seed-fill algorithm, also known as pixel aggregation [8]. In this method, regions consist of spatially connected pixels that fall within a specified gray-level deviation from the starting or seed pixel. For high-resolution digitized mammograms, 4-connectivity was found, by visual comparison, to be adequate to allow accurate region growing, although small features were better matched with 8-connected regions. The use of 8-connectivity for region growing requires longer computing time than 4-connectivity.

The flowchart in Figure 4.38 illustrates the region-growing algorithm. The algorithm starts with the pixel being processed, called the seed pixel, or simply the seed. The seed is placed in an initially empty queue that holds pixels to be evaluated for inclusion in, or exclusion from, the region being grown. The main loop is then entered. If the queue is empty, the program exits the loop; otherwise, the first pixel is taken from the queue. This pixel is called the current pixel; if its gray level value is within the specified deviation from the seed, it is labeled as a foreground pixel. The immediate neighbors (either 4-connected or 8-connected, as specified) of the current pixel could possibly qualify to be foreground pixels, and are added to the queue, if they are not already in the queue from being connected to previously checked pixels. If the current pixel is outside the permitted gray-level range, it is marked as a background pixel, and a border pixel of the region has been reached. A region may have a number of internal borders, in addition to the encompassing external border. Thus, the background may consist of more than one set of pixels, with each such set being disconnected from the others. After all of the current pixel's neighbors have been checked, control is directed back to the start of the loop, to check the next pixel in the queue.

The final step in growing a region around the seed is completing the background. This is done by starting with the existing background points, as found during foreground region growing. The neighbors of this set of pixels are examined to see if they belong to either the foreground or background. If not, they are set to be the next layer of the background. The new layer is then used to grow another layer, and so on, until the specified background width is achieved. The region-growing procedure as described above does have inefficiencies, in that a given pixel may be checked more than once for placement in the queue. More complicated algorithms may be used to grow regions along line segments, and thereby partially eliminate this inefficiency [245]. Preliminary testing of a scan-line based algorithm showed minimal improvement with

**FIGURE 4.38**

Procedure for region growing for adaptive-neighborhood contrast enhancement of mammograms. Reproduced with permission from W.M. Morrow, R.B. Paranjape, R.M. Rangayyan, and J.E.L. Desautels, "Region-based contrast enhancement of mammograms" *IEEE Transactions on Medical Imaging*, 11(3):392-406, 1992. © IEEE.

mammogram images, because the type of regions grown in mammograms are usually complex.

The adaptive-neighborhood contrast enhancement procedure may be stated in algorithmic form as follows [246]:

1. The first pixel (or the next unprocessed pixel) in the image is taken as the seed pixel.
2. The immediate neighbors (8-connected pixels) of the seed are checked for inclusion in the region. Each neighbor pixel is checked to see if its gray-level value is within the specified deviation from the seed pixel's gray-level value. The growth tolerance or deviation is specified as

$$\frac{f(m, n) - \text{seed}}{\text{seed}} \leq \tau, \quad (4.28)$$

where $f(m, n)$ is the gray-level value of the neighbor pixel being checked for inclusion, and the threshold $\tau = 0.05$.

3. If a neighbor pixel's gray-level value is within the specified deviation, it is added to a queue of foreground pixels that will make up the region being grown. A pixel is added to the queue only if it has not already been included while processing another connected pixel.
4. A pixel $f(m, n)$ is taken from the start of the foreground queue. This becomes the current pixel whose 8-connected neighbors are checked against the seed's gray-level according to the tolerance specified, as in Steps 2 and 3 above.
5. If a neighbor pixel's gray-level value is outside the specified gray-level tolerance range, it is marked as a background pixel. A background pixel indicates that the border of the region has been reached at that position. However, if a neighbor pixel's gray-level value is within the specified deviation, it is added to the foreground.
6. Once all the current pixel's neighbors have been checked, the program goes back to Step 4 to check the connected neighbor pixels of the next pixel in the foreground queue.
7. Steps 4 – 6 are repeated until region growing stops (that is, no more pixels can be added to the foreground region).
8. The borders of the foreground region (marked as background pixels) are expanded in all directions by a prespecified number of pixels (three pixels in the work of Morrow et al.) to obtain a background region that is molded to the shape of the foreground region. The foreground and background regions together form the adaptive neighborhood of the seed pixel that was used to start the region-growing procedure. See Figure 3.46 for an example of region growing with an image.

9. The contrast of the region is computed as per Equation 2.7, and enhanced as desired; see Figure 4.39. The gray-level value of the seed pixel is modified as per Equation 4.30. All pixels in the foreground region having the same gray-level value as the seed, referred to as the redundant seed pixels, are also modified to the same value as for the seed pixel.
10. Steps 1 – 9 are executed until all the pixels in the image have been processed.

It should be noted that each pixel in the connected foreground that has the same gray level as the seed will lead to the same foreground and background. These pixels are called the region's redundant seed pixels. Considerable computation may be saved by using this redundancy and obviating the repeated growing of the same regions. Furthermore, the same final transformation that is applied to the region's seed pixel is also applicable to the region's redundant seed pixels. In high-resolution mammogram images, redundant seed pixels were seen to account for over 75% of the pixels in a given image; this large percentage is partially due to the dark background in the image off the projection of the breast, and due to the relatively smooth variations in gray levels in mammograms. The number of redundant seeds is also dependent upon the growth tolerance used for region growing.

Parameters for region growing: The crucial parameter in controlling seed-fill region growing is the criterion used to decide whether a pixel is to be included or excluded in the region. This criterion is defined by the growth tolerance, τ . The growth tolerance indicates the deviation (positive or negative) about the seed pixel's gray level that is allowed within the foreground region. For example, with a growth tolerance of 0.05, any pixel with a gray value between 0.95 and 1.05 times the seed pixel's value, which also satisfies the spatial-connectivity criterion, is included in the region. The reason for using this type of growth tolerance is found from a closer examination of the definition of contrast. Seed-fill region growing results in regions having contrast greater (in magnitude) than a certain minimum contrast, C_{\min} . It is desired that this minimum contrast be independent of a region's gray level, so that the results of enhancement will be independent of a multiplicative transformation of the image. A region with the minimum positive contrast C_{\min} will have a mean foreground value of f and a mean background value of $(1 - \tau)f$. Using Equation 2.7, the minimum contrast C_{\min} is

$$C_{\min} = \frac{f - (1 - \tau)f}{f + (1 - \tau)f} = \frac{\tau}{2 - \tau} \approx \frac{\tau}{2}. \quad (4.29)$$

The contrast C_{\min} is thus independent of the foreground gray level or the background gray level, and depends only upon the region-growing tolerance parameter τ . Weber's ratio of 2% for a just-noticeable feature suggests that the growth tolerance should be about 4%, in order to grow regions that are

barely noticeable prior to enhancement (and are subsequently enhanced to a contrast above the Weber ratio). A lower bound on τ may be established empirically, or, depending upon the class of images being enhanced, through an analysis of the noise present in the images.

Contrast enhancement: Equation 2.7 defines a region's contrast as a function of the mean gray levels of the foreground f and background b . The contrast of a region may be increased by changing f or b . Rearranging Equation 2.7, and replacing C with an increased contrast C_e gives

$$f_e = b \frac{1 + C_e}{1 - C_e}, \quad (4.30)$$

where f_e is the new foreground value. Only the seed pixel and the redundant seed pixels in the foreground are modified to the value f_e . The remaining pixels in the foreground obtain new values when they, in turn, act as seed pixels and are used to grow different regions. (If all the pixels in the foreground were replaced by f_e , the output image would depend on the order in which regions are grown; furthermore, the gray-level variations and details within each region would be lost, and the resulting image would be a collection of uniform regions.) The new contrast C_e for the region may be calculated using an analytic function of C [242, 243, 244, 247], or an empirically determined relationship between C_e and C . Morrow et al. [123] proposed an empirical relationship between C_e and C as shown in Figure 4.39, which was designed to boost the perceptibility of regions with low-to-moderate contrast (in the range 0.02 – 0.5), while not affecting high-contrast regions.

Example — Contrast enhancement of a cluster of calcifications: Figure 4.40 (a) shows a part of a mammogram with a cluster of calcifications. Some of the calcifications are linearly distributed, suggesting that they are intraductal. Cancer was suspected because of the irregular shape and size of the individual constituents of the calcification cluster, although hyperdense tissue could not be clearly seen in this area of the image. A biopsy was subsequently performed on the patient, which confirmed the presence of an invasive intraductal carcinoma.

Figure 4.40 (b) shows the same part of the image as in (a), after adaptive-neighborhood contrast enhancement was applied to the entire mammogram. The curve shown in Figure 4.39 was used as the contrast transformation curve, the growth tolerance was 3%, and a background width of three pixels was used. Increased contrast is apparent in the enhanced image, and subtle details are visible at higher contrast. Observe the presence of sharper edges between features; the contrast of the calcifications has been greatly increased in the processed image. The closed-loop feature immediately below the cluster of calcifications is possibly the cross-sectional projection of a mammary duct. If this interpretation is correct, the distorted geometry (different from the normally circular cross-section) could be indicative of intraductal malignancy. This feature is not readily apparent in the original image.




FIGURE 4.39

An empirical relationship between the contrast C of an adaptive neighborhood and the increased contrast C_e for enhancement of mammograms [123]. $C_e = C$ for $C \geq 0.5$.

In order to compare the results of the adaptive-neighborhood contrast enhancement method with those of other techniques, a simple nonlinear rescaling (or gamma-correction) procedure was applied, with the output being defined as $g(m, n) = f^{1.5}(m, n)$ without normalization of the gray scale. The result was linearly scaled to the display range of $[0, 255]$, and is shown in Figure 4.40 (c). Contrast in the area of the calcification cluster was increased, at the cost of decreased contrast in the darker areas of the image. Although the enhancement is not as good as with adaptive-neighborhood contrast enhancement, the advantage of this method is its simplicity.

The 3×3 unsharp masking filter was applied to the complete mammogram from which the image in Figure 4.40 (a) was obtained. The corresponding portion of the resulting image is shown in Figure 4.40 (d). The contrast and sharpness of the calcification cluster was increased, although not to the same degree as in the image generated using adaptive-neighborhood contrast enhancement. The overall appearance of the image was altered significantly from that of the original image.

Global histogram equalization of the full mammogram led to complete washout of the region with the calcifications. The result, shown in Figure 4.41, indicates the unsuitability of global techniques for the enhancement of mammograms.

The enhancement shown in the above case has limited practical value, because the characteristics of the calcification cluster in the original image are

sufficient to lead the radiologist to recommend biopsy. However, if mammary ducts and other anatomical features become more clearly visible in the enhanced image, as suggested above, the extent and degree of disease could be judged more accurately, and the biopsy method and location determined accordingly.


Example — Contrast enhancement of dense masses: Figure 4.42 (a) shows a portion of a mammogram, in the lower-right quadrant of which a dense mass with diffuse edges and a spiculated appearance is present. The probable presence of calcifications was suggested after examination of the film through a hand lens. Figure 4.42 (b) shows the corresponding part of the mammogram after adaptive-neighborhood contrast enhancement. The internal details of the mass are more readily seen in the enhanced image; the bright, irregular details were suspected to be calcifications. Also of interest is the appearance of the dense mass to the left of the spiculated mass. The mass has smooth margins, and a generally benign appearance. After enhancement, bright, irregularly shaped features are apparent in this mass, and may possibly be calcifications associated with malignancy as well.

Example — Contrast enhancement of a benign mass: Figure 4.43 (a) shows a part of a mammogram with a histologically verified benign cyst. The brighter regions at the center of the cyst do not demonstrate any irregular outline; they were interpreted to be the result of superimposition of crossing, linear, supporting tissues. The corresponding portion from the enhanced image is shown in Figure 4.43 (b). Few changes are apparent as compared with the original image, although contrast enhancement was perceived over the entire image. Enhancement did not affect the appearance or the assessment of the benign cyst.

4.10 Objective Assessment of Contrast Enhancement

The improvement in images after enhancement is often difficult to measure or assess. A processed image can be said to be an enhanced version of the original image if it allows the observer to perceive better the desired information in the image. With mammograms, the improvement in perception is difficult to quantify. The use of statistical measures of gray-level distribution as measures of local contrast enhancement (for example, variance or entropy) is not particularly meaningful for mammographic images.

Morrow et al. [123] proposed a new approach to assess image enhancement through the contrast histogram. The contrast histogram represents the distribution of contrast of all possible regions present in the image. If we measure the contrast of all regions in the image (as obtained by the region-growing procedure described in Section 4.9.1) prior to enhancement and subsequent

**FIGURE 4.40**

(a) Part of a mammogram with a cluster of calcifications, true size 43×43 mm. Results of enhancement by (b) adaptive-neighborhood contrast enhancement; (c) gamma correction; and (d) unsharp masking. See also Figures 4.37 and 4.41. Reproduced with permission from W.M. Morrow, R.B. Paranjape, R.M. Rangayyan, and J.E.L. Desautels, "Region-based contrast enhancement of mammograms" *IEEE Transactions on Medical Imaging*, 11(3):392–406, 1992.

© IEEE.




FIGURE 4.41

Result of enhancement of the image in Figure 4.40 (a) by global histogram equalization applied to the entire image. See also Figures 4.37 and 4.40. Reproduced with permission from W.M. Morrow, R.B. Paranjape, R.M. Rangayyan, and J.E.L. Desautels, "Region-based contrast enhancement of mammograms," *IEEE Transactions on Medical Imaging*, 11(3):392–406, 1992. © IEEE.





FIGURE 4.42

(a) Part of a mammogram with dense masses, true size 43×43 mm. (b) Result of enhancement by adaptive-neighborhood contrast enhancement. Reproduced with permission from W.M. Morrow, R.B. Paranjape, R.M. Rangayyan, and J.E.L. Desautels, "Region-based contrast enhancement of mammograms," *IEEE Transactions on Medical Imaging*, 11(3):392–406, 1992. © IEEE.

**FIGURE 4.43**

(a) Part of a mammogram with a benign cyst, true size $43 \times 43 \text{ mm}$. (b) Result of enhancement by adaptive-neighborhood contrast enhancement. Reproduced with permission from W.M. Morrow, R.B. Paranjape, R.M. Rangayyan, and J.E.L. Desautels, "Region-based contrast enhancement of mammograms," *IEEE Transactions on Medical Imaging*, 11(3):392–406, 1992. © IEEE.

to enhancement, the contrast histogram of the enhanced image should contain more counts of regions at higher contrast levels than the contrast histogram of the original image. Various enhancement methods can be quantitatively compared by measuring the properties of their respective contrast histograms.

The spread of a contrast histogram may be quantified by taking the second moment about the zero-contrast level. For a distribution of contrast values c_i , quantized so that there are N bins over the range $[-1, 1]$, the second moment M_2 is

$$M_2 = \sum_{i=1}^N c_i^2 p(c_i), \quad (4.31)$$

where $p(c_i)$ is the normalized number of occurrences of seed pixels (including redundant seed pixels) that lead to the growth of a region with contrast c_i . A low-contrast image, that is, an image with a narrow contrast histogram, will have a low value for M_2 ; an image with high contrast will have a broader contrast histogram, and hence a greater value of M_2 .

For the purpose described above, image contrast needs to be recomputed after the entire image has been enhanced, because the relative contrast between adjacent regions is dependent upon the changes made to each of the regions. In order to measure the contrast in an image after enhancement, region growing (using the same parameters as in the enhancement procedure)

is performed on the output enhanced image, and a contrast histogram is generated.


In general, the final contrast values in the output image of adaptive-neighborhood contrast enhancement will not match the contrast values specified by the contrast transformation in Equation 4.30. This is because Equation 4.30 is applied pixel-by-pixel to the input image, and the adaptive neighborhood for each pixel will vary. Only if all the pixels in an object have exactly the same gray-level value will they all have exactly the same adaptive neighborhood and be transformed in exactly the same way. Thus, the contrast enhancement curve is useful for identifying the ranges in which contrast enhancement is desired, but cannot specify the final contrast of the regions. The contrast of each region grown in the image is dependent on the value specified by the initial region contrast and the transformation curve, as well as the transformation applied to adjacent regions.

Figure 4.44 shows the contrast histograms of the complete mammograms corresponding to the images in Figure 4.40. The contrast distribution is plotted on a logarithmic scale in order to emphasize the small numbers of occurrence of features at high contrast values. The wider distribution and greater occurrence of regions at high contrast values in the histogram of the adaptive-neighborhood enhanced image show that it has higher contrast. The histograms of the results of gamma correction and unsharp masking also show some increase in the counts for larger contrast values than that of the original, but not to the same extent as the result of adaptive-neighborhood contrast enhancement. The values of M_2 for the four histograms in Figure 4.44 are 3.71×10^{-4} , 6.17×10^{-4} , 3.2×10^{-4} , and 4.4×10^{-4} . The contrast histogram and its statistics provide objective means for the analysis of image enhancement.

4.11 Application: Contrast Enhancement of Mammograms

The accurate diagnosis of breast cancer depends upon the quality of the mammograms obtained; in particular, the accuracy of diagnosis depends upon the visibility of small, low-contrast objects within the breast image. Unfortunately, the contrast between malignant tissue and normal tissue is often so low that the detection of malignant tissue becomes difficult. Hence, the fundamental enhancement needed in mammography is an increase in contrast, especially for dense breasts.

Drongers and Zwaag [248] suggested the use of reversal film rather than negative film for the implementation of a form of photographic contrast enhancement for mammograms. They found that the image quality produced



(a)






Figure 4.44 (b)



(c)



(d)

FIGURE 4.44

Contrast histograms of the full mammograms corresponding to the images in Figure 4.40. (a) Original, $M_2 = 3.71 \times 10^{-4}$; (b) adaptive-neighborhood contrast enhancement, $M_2 = 6.17 \times 10^{-4}$; (c) gamma correction, $M_2 = 3.2 \times 10^{-4}$; and (d) unsharp masking, $M_2 = 4.4 \times 10^{-4}$. Reproduced with permission from W.M. Morrow, R.B. Paranjape, R.M. Rangayyan, and J.E.L. Desautels, "Region-based contrast enhancement of mammograms," *IEEE Transactions on Medical Imaging*, 11(3):392–406, 1992. © IEEE.

was equal to that of conventional techniques without the need for special mammographic equipment. A photographic unsharp-masking technique for mammographic images was proposed by McSweeney et al. [249]. This procedure includes two steps: first, a blurred image is produced by copying the original mammogram through a sheet of glass or clear plastic that diffuses the light; then, by using subtraction print film, the final image is formed by subtracting the blurred image from the original mammogram. Although the photographic technique improved the visualization of mammograms, it was not widely adopted, possibly due to the variability in the image reproduction procedure.

Askins et al. [250] investigated autoradiographic enhancement of mammograms by using thiourea labeled with ^{35}S . Mammograms underexposed as much as tenfold could be autoradiographically intensified so that the enhanced image was comparable to a normally exposed film. The limitations to routine use of autoradiographic techniques include cost, processing time, and the disposal of radioactive solutions.

Digital image enhancement techniques have been used in radiography for more than three decades. (See Bankman [251] for a section including discussions on several enhancement techniques.) Ram [252] stated that images considered unsatisfactory for medical analysis may be rendered usable through various enhancement techniques, and further indicated that the application of such techniques in a clinical situation may reduce the radiation dose by about 50%. Rogowska et al. [253] applied digital unsharp masking and local contrast stretching to chest radiographs, and reported that the quality of images was improved. Chan et al. [254] investigated unsharp-mask filtering for digital mammography: according to their receiver operating characteristics (ROC) studies, unsharp masking could improve the detectability of calcifications on digital mammograms. However, this method also increased noise and caused some artifacts.

Algorithms based on adaptive-neighborhood image processing to enhance mammographic contrast were first reported on by Gordon and Rangayyan [242]. Rangayyan and Nguyen [243] defined a tolerance-based method for growing foreground regions that could have arbitrary shapes rather than square shapes. Morrow et al. [215, 123] further developed this approach with a new definition of background regions. Dhawan et al. [247] investigated the benefits of various contrast transfer functions, including \sqrt{C} , $\ln(1 + 3C)$, $1 - e^{-3C}$, and $\tanh(3C)$, where C is the original contrast, but used square adaptive neighborhoods. They found that while a suitable contrast function was important to bring out the features of interest in mammograms, it was difficult to select such a function. Later, Dhawan and Le Royer [255] proposed a tunable contrast enhancement function for improved enhancement of mammographic features.

Emphasis has recently been directed toward image enhancement based upon the characteristics of the human visual system [256], leading to innovative methods using nonlinear filters, scale-space filters, multiresolution filters, and

wavelet transforms. Attention has been paid to designing algorithms to enhance the contrast and visibility of diagnostic features while maintaining control on noise enhancement. Laine et al. [257] presented a method for nonlinear contrast enhancement based on multiresolution representation and the use of dyadic wavelets. A software package named MUSICA [258] (MUlti-Scale Image Contrast Amplification) has been produced by Agfa-Gevaert. Belikova et al. [259] discussed various optimal filters for the enhancement of mammograms. Qu et al. [260] used wavelet techniques for enhancement and evaluated the results using breast phantom images. Tahoces et al. [261] presented a multistage spatial filtering procedure for nonlinear contrast enhancement of chest and breast images. Qian et al. [262] reported on tree-structured nonlinear filters based on median filters and an edge detector. Chen et al. [263] proposed a regional contrast enhancement technique based on unsharp masking and adaptive density shifting.

The various mammogram enhancement algorithms that have been reported in the literature may be sorted into three categories: algorithms based on conventional image processing methods [253, 254, 259, 261, 264, 265]; adaptive algorithms based on the principles of human visual perception [123, 242, 247, 255, 256, 263, 266]; and multiresolution enhancement algorithms [257, 260, 262, 267, 268, 269, 270]. In order to evaluate the diagnostic utility of an enhancement algorithm, an ROC study has to be conducted; however, few of the above-mentioned methods [254, 264, 266, 267, 271] have been tested with ROC procedures; see Sections 12.8.1 and 12.10 for details on ROC analysis.

4.11.1 Clinical evaluation of contrast enhancement

In order to examine the differences in radiological diagnoses that could result from adaptive-neighborhood enhancement of mammograms, eight test cases from the teaching library of the Foothills Hospital (Calgary, Alberta, Canada) were studied in the work of Morrow et al. [123]. For each of the cases, the pathology was known due to biopsy or other follow-up procedures. For each case, a single mammographic film that presented the abnormality was digitized using an Eikonix 1412 scanner (Eikonix Inc., Bedford, MA) to 4,096 by about 2,048 pixels with 12-bit gray-scale resolution. (The size of the digitized image differed from film to film depending upon the size of the actual image in the mammogram.) The effective pixel size was about $0.054 \text{ mm} \times 0.054 \text{ mm}$. Films were illuminated by a Plannar 1417 light box (Gordon Instruments, Orchard Park, NY). Although the light box was designed to have a uniform light intensity distribution, it was necessary to correct for nonuniformities in illumination. After correction, pixel gray levels were determined to be accurate to 10 bits, with a dynamic range of approximately $0.02 - 2.52 \text{ OD}$ [174].

The images were enhanced using the adaptive-neighborhood contrast enhancement method. For all images, the tolerance τ for region growing was set at 0.05, the width of the background was set to three pixels, and the enhance-

ment curve used was that presented in Figure 4.39. The original and processed images were down-sampled by a factor of two for processing and display for interpretation on a MegaScan 2111 monitor (Advanced Video Products Inc., Littleton, MA). Although the memory buffer of the MegaScan system was of size $4,096 \times 4,096 \times 12$ bits, the display buffer was limited to $2,560 \times 2,048 \times 8$ bits, with panning and zooming facilities. The monitor displayed images at 72 noninterlaced frames per second.

In each case, the original, digitized mammogram was first presented on the MegaScan 2111 monitor. The image occupied about 20×15 cm on the screen. An experienced radiologist, while viewing the digitized original, described the architectural abnormalities that were observed. Subsequently, the enhanced image was added to the display. While observing both the enhanced mammogram and the original mammogram together, the radiologist described any new details or features that became apparent.

Case (1) was that of a 62-year-old patient with a history of diffuse nodularity in both breasts. The MLO view of the left breast was digitized for assessment. The unenhanced mammogram revealed two separate nodular lesions: one with well-defined boundaries, with some indication of lobular calcium; the other smaller, with poorly defined borders, some spiculation, but no microcalcifications. The unenhanced mammogram suggested that the smaller lesion was most likely associated with carcinoma; however, there was some doubt about the origins of the larger lesion. An examination of the enhanced mammogram revealed definite calcium deposits in the larger lesion and some indication of microcalcifications in the smaller lesion. The enhanced image suggested carcinoma as the origin of both lesions more strongly than the unenhanced mammogram. The biopsy report for both areas indicated intraductal infiltrating carcinoma, confirming the diagnosis from the enhanced mammogram.

Case (2) was that of a 64-year-old patient. The digitized original mammogram was the CC view of the left breast. The unenhanced mammogram contained two lesions. The lesion in the lower-outer part of the breast had irregular edges and coarse calcifications, whereas the other lesion appeared to be a cyst. Examination of the unenhanced mammogram suggested that both lesions were benign. Examination of the enhanced mammogram revealed no additional details that would suggest a change in the original diagnosis. The appearance of the lesions was not much different from that seen in the unenhanced mammogram; however, the details in the internal architecture of the breast appeared clearer, adding further weight to the diagnosis of benign lesions. Excision biopsies carried out at both sites confirmed this diagnosis.

Case (3) was that of a 44-year-old patient, for whom the MLO view of the left breast was digitized. The original digitized mammogram revealed multiple benign cysts as well as a spiculated mass in the upper-outer quadrant of the breast. There was some evidence of calcium, but it was difficult to confirm the same by visual inspection. A dense nodule was present adjacent to the spiculated mass. Examination of the enhanced mammogram revealed that the

spiculated mass did contain microcalcifications. The dense nodule appeared to be connected to the spiculated mass, suggesting a further advanced carcinoma than that suspected from the unenhanced mammogram. Biopsy reports were available only for the spiculated region, and indicated lobular carcinoma. No further information was available to verify the modified diagnosis from the enhanced mammogram.

Case (4) was that of a 40-year-old patient, whose mammograms indicated dense breasts. The image of the right breast indicated an area of uniform density. The CC view of the right breast was digitized and enhanced. The digitized original mammogram indicated a cluster of microcalcifications, all of approximately uniform density, centrally located above the nipple. The enhanced mammogram indicated a similar finding with a larger number of microcalcifications visible, and some irregularity in the density of the calcifications. Both the original and the enhanced mammograms suggested a similar diagnosis of intraductal carcinoma. Biopsy of the suspected area confirmed this diagnosis.

Case (5) was that of a 64-year-old patient with a history of a benign mass in the right breast. A digitized mammogram of the CC view of the right breast was examined. The unenhanced mammogram clearly showed numerous microcalcifications that were roughly linear in distribution, with some variation in density. The original mammogram clearly suggested intraductal carcinoma. The enhanced mammogram showed a greater number of calcifications, indicating a lesion of larger extent. The variation in the density of the calcifications was more evident. Biopsy indicated an infiltrating ductal carcinoma.

Case (6) was that of a 59-year-old patient whose right CC view was digitized. The original mammogram indicated a poorly defined mass with some spiculations. The lesion was irregular in shape, and contained some calcium. The unenhanced mammogram suggested intraductal carcinoma. The enhanced mammogram provided stronger evidence of carcinoma with poor margins of the lesion, a greater number of microcalcifications, and inhomogeneity in the density of the calcifications. Biopsy confirmed the presence of the carcinoma.

Case (7) involved the same patient as in Case (6); however, the mammogram was taken one year after that described in Case (6). The digitized mammogram was the CC view of the right breast. The unenhanced view showed significant architectural distortion due to segmental mastectomy. The unenhanced mammogram showed an area extending past the scarred region of fairly uniform density with irregular boundaries. The unenhanced mammogram along with the patient's history suggested the possibility of cancer, and biopsy was recommended. The enhanced mammogram suggested a similar finding, with added evidence of some small microcalcifications in the uniform area. Biopsy of the region showed that the mass was, in fact, a benign hematoma.

Case (8) was that of an 86-year-old patient; the MLO view of the left breast was digitized. In the unenhanced mammogram, a dense region was observed

with some speculations. The mammogram suggested the possibility of carcinoma and biopsy was recommended. The enhanced mammogram showed the same detail as the unenhanced mammogram, with the additional finding of some microcalcifications; this added to the suspicion of cancer. The biopsy of the region indicated intraductal invasive carcinoma with lymph-node metastasis present.

In each of the eight cases described above, the overall contrast in the enhanced mammogram was significantly improved. This allowed the radiologist to comment that “much better overall anatomical detail” was apparent in the enhanced mammograms, and that “overall detail (internal architecture) is improved” in the enhanced mammograms. In all cases, the radiological diagnosis was confirmed by biopsy. In seven of the eight cases, the enhanced mammogram added further weight to the diagnosis made from the original mammogram, and the diagnosis was confirmed by biopsy. In one case, the enhanced mammogram as well as the unenhanced mammogram suggested the possibility of carcinoma; however, the biopsy report indicated a benign condition. This case was, however, complicated by the fact that the patient’s history influenced the radiologist significantly. While it is not possible to make a quantitative assessment of the differences in diagnoses from the qualitative comparison as above, it appeared that a clearer indication of the patient’s condition was obtained by examination of the enhanced mammogram.

The adaptive-neighborhood contrast enhancement method was used in a preference study comparing the performance of enhancement algorithms by Sivaramakrishna et al. [125]. The other methods used in the study were adaptive unsharp masking, contrast-limited adaptive histogram equalization, and wavelet-based enhancement. The methods were applied to mammograms of 40 cases, including 10 each of benign and malignant masses, and 10 each of benign and malignant microcalcifications. The four enhanced images and the original image of each case were displayed randomly across three high-resolution monitors. Four expert mammographers ranked the images from 1 (best) to 5 (worst). In a majority of the cases with microcalcifications, the adaptive-neighborhood contrast enhancement algorithm provided the most-preferred images. In the set of images with masses, the unenhanced images were preferred in most of the cases.

See Sections 12.8.1, 12.8.2, and 12.10 for discussions on statistical analysis of the clinical outcome with enhanced mammograms.

4.12 Remarks

Quite often, an image acquired in a real-life application does not have the desired level of quality in terms of contrast, sharpness of detail, or the visibility

of the features of interest. We explored several techniques in this chapter that could assist in improving the quality of a given image. The class of filters based upon mathematical morphology [8, 192, 220, 221, 222] has not been dealt with in this chapter.

An understanding of the exact phenomenon that caused the poor quality of the image at the outset could assist in the design of an appropriate technique to address the problem. However, in the absence of such information, one could investigate the suitability of existing and established models of degradation, as well as the associated enhancement techniques to improve the quality of the image on hand. It may be desirable to obtain several enhanced versions using a variety of approaches; the most suitable image may then be selected from the collection of the processed images for further analysis. In situations as above, there is no single or optimal solution to the problem. Several enhanced versions of the given image may also be analyzed simultaneously; however, this approach could demand excessive time and resources, and may not be feasible in a large-scale screening application.

Given the subjective nature of image quality, and in spite of the several methods we studied in Chapter 2 to characterize image quality and information content, the issue of image enhancement is nonspecific and elusive. Regardless, if a poor-quality image can be enhanced to the satisfaction of the user, and if the enhanced image leads to improved analysis — and more accurate or confident diagnosis in the biomedical context — an important achievement could result.

The topic of image restoration — image quality improvement when the exact cause of degradation is known and can be represented mathematically — is investigated in Chapter 10.

4.13 Study Questions and Problems

(*Note:* Some of the questions may require background preparation with other sources on the basics of signals and systems as well as digital signal and image processing, such as Lathi [1], Oppenheim et al. [2], Oppenheim and Schaffer [7], Gonzalez and Woods [8], Pratt [10], Jain [12], Hall [9], and Rosenfeld and Kak [11].)

Selected data files related to some of the problems and exercises are available at the site

www.enel.ualgary.ca/People/Ranga/enel697

1. A poorly exposed image was found to have gray levels limited to the range 25 – 90. Derive a linear transform to stretch this range to the display range of 0 – 255.
Give the display values for the original gray levels of 45 and 60.

2. Explain the differences between the Laplacian and subtracting Laplacian operators in the spatial and frequency domains.
3. Compute by hand the result of linear convolution of the following two images:

$$\begin{bmatrix} 3 & 5 & 5 & 5 \\ 0 & 0 & 1 & 3 \\ 4 & 4 & 4 & 3 \\ 2 & 2 & 2 & 2 \\ 2 & 2 & 2 & 2 \end{bmatrix} \quad (4.32)$$

and

$$\begin{bmatrix} 3 & 5 & 1 \\ 4 & 2 & 3 \\ 1 & 3 & 2 \end{bmatrix}. \quad (4.33)$$

4. Explain the differences between the 3×3 mean and median filters. Would you be able to compare the filters in the Fourier domain? Why (not)?
5. Derive the frequency response of the 3×3 unsharp masking filter and explain its characteristics.
6. An image has a uniform PDF (normalized gray-level histogram) over the range $[0, 255]$. A novice researcher derives the transform to perform histogram equalization. Derive an analytical representation of the transform. Explain its effects on the image in terms of the modification of gray levels and the histogram.
7. An image has a uniform PDF (normalized gray-level histogram) over the range $[25, 90]$, with the probability being zero outside this interval within the available range of $[0, 255]$. Derive an analytical representation of the transform to perform histogram equalization. Explain its effects on the image in terms of the modification of gray levels and the histogram.
8. Give an algorithmic representation of the method to linearly map a selected range of gray-level values $[x_1, x_2]$ to the range $[y_1, y_2]$ in an image of size $M \times N$. Values below x_1 are to be mapped to y_1 , and values above x_2 mapped to y_2 . Use pseudocode format and show all the necessary programming steps and details.
9. An 8×8 image with an available gray-level range of $0 - 7$ at 3 bits/pixel has the following pixel values:

$$\begin{bmatrix} 3 & 5 & 5 & 5 & 4 & 4 & 4 & 3 \\ 3 & 3 & 4 & 5 & 5 & 3 & 3 & 2 \\ 0 & 0 & 1 & 3 & 4 & 4 & 4 & 4 \\ 4 & 5 & 5 & 5 & 3 & 2 & 2 & 4 \\ 4 & 4 & 4 & 3 & 3 & 3 & 3 & 2 \\ 2 & 2 & 2 & 2 & 2 & 1 & 1 & 1 \\ 2 & 2 & 2 & 2 & 1 & 1 & 1 & 1 \\ 2 & 2 & 2 & 2 & 1 & 1 & 1 & 1 \end{bmatrix}. \quad (4.34)$$

Derive the transformation and look-up table for enhancement of the image by histogram equalization. Clearly show all of the steps involved, and give the

pixel values in the enhanced image using the available gray-level range of 3 bits/pixel.

Draw the histograms of the original image and the enhanced image. Explain the differences between them as caused by histogram equalization.

10. Write the expression for the convolution of an $N \times N$ digital image with an $M \times M$ digital image (or filter function) with $M \ll N$.

Using pseudocode format, show all of the necessary programming steps and details related to the implementation of convolution as above.

Explain how you handle the size and data at the edges of the resulting image.

11. Prepare a 5×5 image with zero pixel values. Add a square of size 3×3 pixels with the value 100 at the center of the image. Apply

(a) the subtracting Laplacian operator,

and

(b) the Laplacian operator

to the image. Examine the pixel values inside and around the edges of the square in the resulting images. Give reasons for the effects you find.

12. Apply

(a) the subtracting Laplacian operator,

and

(b) the Laplacian operator

to the image in Equation 4.34. Give reasons for the effects you find.

13. Derive the MTF of the 3×3 unsharp masking operator.

Explain its characteristics.

14. An image is processed by applying the subtracting Laplacian mask and then by applying the 3×3 mean filter mask.

What is the impulse response of the complete system?

What is the MTF of the complete system?

Explain the effect of each operator.

15. Derive the MTF of the 3×3 subtracting Laplacian operator and explain its characteristics.

16. What causes ringing artifact in frequency-domain filtering?

How do you prevent the artifact?

17. Discuss the differences between highpass filtering and high-frequency emphasis filtering in the frequency domain in terms of their

(a) transfer functions, and

(b) effects on image features.

18. List the steps of computation required in order to perform lowpass filtering of an image in the frequency domain by using the Fourier transform.

4.14 Laboratory Exercises and Projects

1. Select two underexposed images, or images with bright and dark regions such that the details in some parts are not clearly visible, from your collection. Apply histogram equalization, gamma adjustment, and linear gray-level mapping transforms to the images.

Compare the results in terms of the enhancement of the visibility of details, saturation or loss of details at the high or low ends of the gray scale, and overall visual quality.

Plot the histograms of the resulting images and compare them with the histograms of the original images. Comment upon the differences.

2. Select two images from your collection, with one containing relatively sharp and well-defined edges, and the other containing smooth features.

Apply the unsharp masking filter, the Laplacian operator, and the subtracting Laplacian filter to the images. Study the results in terms of edge enhancement.

Create noisy versions of the images by adding Gaussian noise. Apply the enhancement methods as above to the noisy images. Study the results in terms of edge enhancement and the effect of noise.

3. Select two images from your collection, with one containing relatively sharp and well-defined edges, and the other containing smooth features.

Apply the ideal highpass filter, the Butterworth highpass filter, and the Butterworth high-emphasis filter to the images. Use at least two different cutoff frequencies. Study the results in terms of edge enhancement or edge extraction.

Create noisy versions of the images by adding Gaussian noise. Apply the filters as above to the noisy images. Study the results in terms of edge enhancement or extraction and the effect of noise.

Detection of Regions of Interest

Although a physician or a radiologist, of necessity, will carefully examine an image on hand in its entirety, more often than not, diagnostic features of interest manifest themselves in local regions. It is uncommon that a condition or disease will alter an image over its entire spatial extent. In a screening situation, the radiologist scans the entire image and searches for features that could be associated with disease. In a diagnostic situation, the medical expert concentrates on the region of suspected abnormality, and examines its characteristics to decide if the region exhibits signs related to a particular disease.

In the CAD environment, one of the roles of image processing would be to detect the region of interest (ROI) for a given, specific, screening or diagnostic application. Once the ROIs have been detected, the subsequent tasks would relate to the characterization of the regions and their classification into one of several categories. A few examples of ROIs in different biomedical imaging and image analysis applications are listed below.

- Cells in cervical-smear test images (Papanicolaou or Pap-smear test) [272, 273].
- Calcifications in mammograms [274].
- Tumors and masses in mammograms [275, 276, 277].
- The pectoral muscle in mammograms [278].
- The breast outline or skin-air boundary in mammograms [279].
- The fibroglandular disc in mammograms [280].
- The air-way tree in lungs.
- The arterial tree in lungs.
- The arterial tree of the left ventricle, and constricted parts of the same due to plaque development.

Segmentation is the process that divides an image into its constituent parts, objects, or ROIs. Segmentation is an essential step before the description, recognition, or classification of an image or its constituents. Two major approaches to image segmentation are based on the detection of the following characteristics:

- **Discontinuity** — Abrupt changes in gray level (corresponding to edges) are detected.
- **Similarity** — Homogeneous parts are detected, based on gray-level thresholding, region growing, and region splitting/merging.

Depending upon the nature of the images and the ROIs, we may attempt to detect the edges of the ROIs (if distinct edges are present), or we may attempt to grow regions to approximate the ROIs. It should be borne in mind that, in some cases, an ROI may be composed of several disjoint component areas (for example, a tumor that has metastasized into neighboring regions and calcifications in a cluster). Edges that are detected may include disconnected parts that may have to be matched and joined. We shall explore several techniques of this nature in the present chapter.

Notwithstanding the stated interest in local regions as above, applications do exist where entire images need to be analyzed for global changes in patterns: for example, changes in the orientational structure of collagen fibers in ligaments (see Figure 1.8), and bilateral asymmetry in mammograms (see Section 8.9). Furthermore, in the case of clustered calcifications in mammograms, cells in cervical smears, and other examples of images with multicomponent ROIs, analysis may commence with the detection of single units of the pattern of interest, but several such units present in a given image may need to be analyzed, separately and together, in order to reach a decision regarding the case.

5.1 Thresholding and Binarization

If the gray levels of the objects of interest in an image are known from prior knowledge, or can be determined from the histogram of the given image, the image may be thresholded to detect the features of interest and reject other details. For example, if it is known that the objects of interest in the image have gray-level values greater than L_1 , we could create a binary image for display as

$$g(m, n) = \begin{cases} 0 & \text{if } f(m, n) \leq L_1 \\ 255 & \text{if } f(m, n) > L_1 \end{cases}, \quad (5.1)$$

where $f(m, n)$ is the original image; $g(m, n)$ is the thresholded image to be displayed; and the display range is $[0, 255]$. See also Section 4.4.1.

Methods for the derivation of optimal thresholds are described in Sections 5.4.1, 8.3.2, and 8.7.2.

Example: Figure 5.1 (a) shows a TEM image of a ligament sample demonstrating collagen fibers in cross-section; see Section 1.4. Inspection of the histogram of the image (shown in Figure 2.12) shows that the sections of the

collagen fibers in the image have gray-level values less than about 180; values greater than this level represent the brighter background in the image. The histogram also indicates the fact that the gray-level ranges of the collagen-fiber regions and the background overlap significantly. Figure 5.1 (b) shows a thresholded version of the image in (a), with all pixels less than 180 appearing in black, and all pixels above this level appearing in white. This operation is the same as the thresholding operation given by Equation 5.1, but in the opposite sense. Most of the collagen fiber sections have been detected by the thresholding operation. However, some of the segmented regions are incomplete or contain holes, whereas some parts that appear to be separate and distinct in the original image have been merged in the result. An optimal threshold derived using the methods described in Sections 5.4.1, 8.3.2, and 8.7.2 could lead to better results.

5.2 Detection of Isolated Points and Lines

Isolated points may exist in images due to noise or due to the presence of small particles in the image. The detection of isolated points is useful in noise removal and the analysis of particles. The following convolution mask may be used to detect isolated points [8]:


$$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}. \quad (5.2)$$

The operation computes the difference between the current pixel at the center of the mask and the average of its 8-connected neighbors. (The mask could also be seen as a generalized version of the Laplacian mask in Equation 2.83.) The result of the mask operation could be thresholded to detect isolated pixels where the difference computed would be large.


Straight lines or line segments oriented at 0° , 45° , 90° , and 135° may be detected by using the following 3×3 convolution masks [8]:

$$\begin{aligned} & \begin{bmatrix} -1 & -1 & -1 \\ 2 & 2 & 2 \\ -1 & -1 & -1 \end{bmatrix}, \quad \begin{bmatrix} -1 & -1 & 2 \\ -1 & 2 & -1 \\ 2 & -1 & -1 \end{bmatrix}, \\ & \begin{bmatrix} -1 & 2 & -1 \\ -1 & 2 & -1 \\ -1 & 2 & -1 \end{bmatrix}, \quad \begin{bmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{bmatrix}. \end{aligned} \quad (5.3)$$

A line may be said to exist in the direction for which the corresponding mask provides the largest response.



(a)



(b)

FIGURE 5.1

(a) TEM image of collagen fibers in a scar-tissue sample from a rabbit ligament at a magnification of approximately $\times 30,000$. See also Figure 1.5. Image courtesy of C.B. Frank, Department of Surgery, University of Calgary. See Figure 2.12 for the histogram of the image. (b) Image in (a) thresholded at the gray level of 180.

5.3 Edge Detection

One of the approaches to the detection of an ROI is to detect its edges. The HVS is particularly sensitive to edges and gradients, and some theories and experiments indicate that the detection of edges plays an important role in the detection of objects and analysis of scenes [122, 281, 282].

In Section 2.11.1 on the properties of the Fourier transform, we saw that the first-order derivatives and the Laplacian relate to the edges in the image. Furthermore, we saw that these space-domain operators have equivalent formulations in the frequency domain as highpass filters with gain that is proportional to frequency in a linear or quadratic manner. The enhancement techniques described in Sections 4.6 and 4.7 further strengthen the relationship between edges, gradients, and high-frequency spectral components. We shall now explore how these approaches may be extended to detect the edges or contours of objects or regions.

(*Note:* Some authors consider edge extraction to be a type of image enhancement.)

5.3.1 Convolution mask operators for edge detection

An edge is characterized by a large change in the gray level from one side to the other, in a particular direction dependent upon the orientation of the edge. Gradients or derivatives measure the rate of change, and hence could serve as the basis for the development of methods for edge detection.

The first derivatives in the x and y directions, approximated by the first differences, are given by (using matrix notation)

$$\begin{aligned}f'_{yb}(m, n) &\approx f(m, n) - f(m - 1, n), \\f'_{xb}(m, n) &\approx f(m, n) - f(m, n - 1),\end{aligned}\tag{5.4}$$

where the additional subscript b indicates a backward-difference operation. Because causality is usually not a matter of concern in image processing, the differences may also be defined as

$$\begin{aligned}f'_{yf}(m, n) &\approx f(m + 1, n) - f(m, n), \\f'_{xf}(m, n) &\approx f(m, n + 1) - f(m, n),\end{aligned}\tag{5.5}$$

where the additional subscript f indicates a forward-difference operation. A limitation of the operators as above is that they are based upon the values of only two pixels; this makes the operators susceptible to noise or spurious pixel values. A simple approach to design robust operators and reduce the sensitivity to noise is to incorporate averaging over multiple measurements.

Averaging the two definitions of the derivatives in Equations 5.4 and 5.5, we get

$$\begin{aligned} f'_{ya}(m, n) &\approx 0.5 [f(m+1, n) - f(m-1, n)], \\ f'_{xa}(m, n) &\approx 0.5 [f(m, n+1) - f(m, n-1)], \end{aligned} \quad (5.6)$$

where the additional subscript a indicates the inclusion of averaging.

In image processing, it is also desirable to express operators in terms of odd-sized masks that may be centered upon the pixel being processed. The Prewitt operators take these considerations into account with the following 3×3 masks for the horizontal and vertical derivatives G_x and G_y , respectively:

$$G_x : \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix}. \quad (5.7)$$

$$G_y : \begin{bmatrix} -1 & -1 & -1 \\ 0 & 0 & 0 \\ 1 & 1 & 1 \end{bmatrix}. \quad (5.8)$$

The Prewitt operators use three differences across pairs of pixels in three rows or columns around the pixel being processed. Due to this fact, and due to the scale factor of 0.5 in Equation 5.6, in order to derive the exact gradient, the results of the Prewitt operators should be divided by $3 \times 2 \times \Delta$, where Δ is the sampling interval in x and y ; however, this step could be ignored if the result is scaled for display or thresholded to detect edges.

In order to accommodate the orientation of the edge, a vectorial form of the gradient could be composed as

$$\begin{aligned} \mathbf{G}_f(m, n) &= G_{fx}(m, n) + j G_{fy}(m, n), \\ \|\mathbf{G}_f(m, n)\| &= \sqrt{G_{fx}^2(m, n) + G_{fy}^2(m, n)}, \\ \angle \mathbf{G}_f(m, n) &= \tan^{-1} \left(\frac{G_{fy}(m, n)}{G_{fx}(m, n)} \right), \end{aligned} \quad (5.9)$$

where

$$G_{fx}(m, n) = (f * G_x)(m, n), \quad (5.10)$$

and

$$G_{fy}(m, n) = (f * G_y)(m, n). \quad (5.11)$$

If the magnitude is to be scaled for display or thresholded for the detection of edges, the square-root operation may be dropped, or the magnitude approximated as $|G_{fx}| + |G_{fy}|$ in order to save computation.

The Sobel operators are similar to the Prewitt operators, but include larger weights for the pixels in the row or column of the pixel being processed as

$$G_x : \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}, \quad (5.12)$$

$$G_y : \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}. \quad (5.13)$$

Edges oriented at 45° and 135° may be detected by using rotated versions of the masks as above. The Prewitt operators for the detection of diagonal edges are

$$G_{45^\circ} : \begin{bmatrix} 0 & -1 & -1 \\ 1 & 0 & -1 \\ 1 & 1 & 0 \end{bmatrix}, \quad (5.14)$$

and

$$G_{135^\circ} : \begin{bmatrix} -1 & -1 & 0 \\ -1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}. \quad (5.15)$$

Similar masks may be derived for the Sobel operator.

(*Note:* The positive and negative signs of the elements in the masks above may be interchanged to obtain operators that detect gradients in the opposite directions. This step is not necessary if directions are considered in the range $0^\circ - 180^\circ$ only, or if only the magnitudes of the gradients are required.)

Observe that the sum of all of the weights in the masks above is zero. This indicates that the operation being performed is a derivative or gradient operation, which leads to zero output values in areas of constant gray level, and the loss of intensity information.

The Roberts operator uses 2×2 neighborhoods to compute cross-differences as

$$\begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix} \text{ and } \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}. \quad (5.16)$$

The masks are positioned with the upper-left element placed on the pixel being processed. The absolute values of the results of the two operators are added to obtain the net gradient:

$$g(m, n) = |f(m + 1, n + 1) - f(m, n)| + |f(m + 1, n) - f(m, n + 1)|, \quad (5.17)$$

with the indices in matrix-indexing notation. The individual differences may also be squared, and the square root of their sum taken to be the net gradient. The advantage of the Roberts operator is that it is a forward-looking operator, as a result of which the result may be written in the same array as the input image. This was advantageous when computer memory was expensive and in short supply.

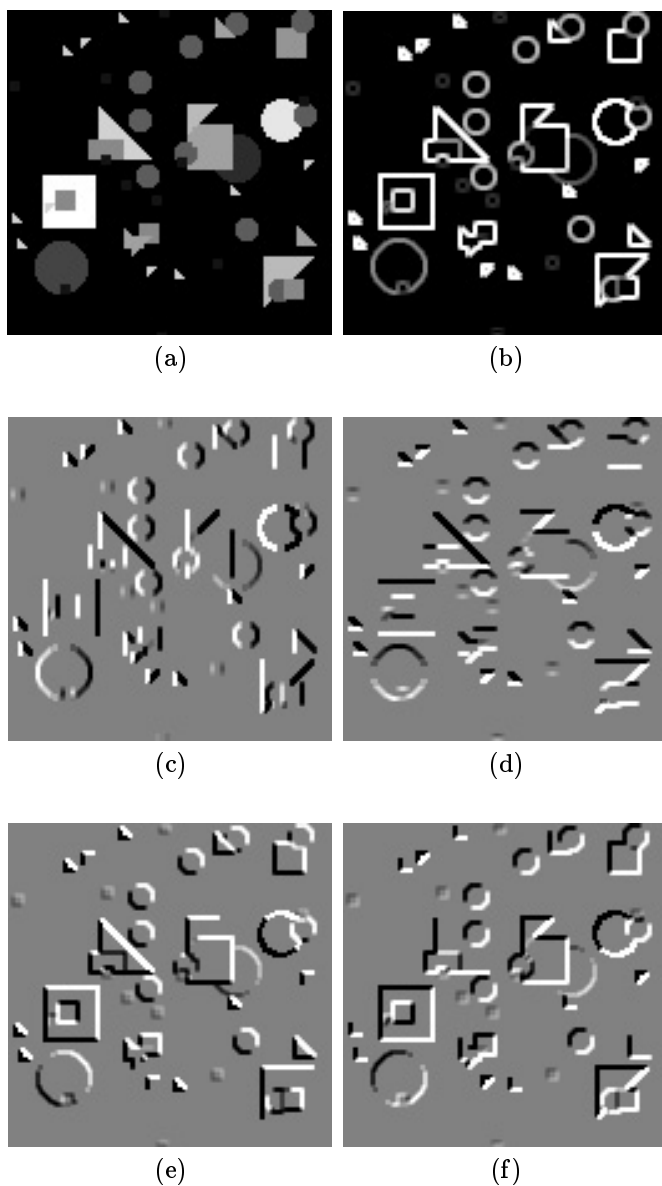
Examples: Figure 5.2 (a) shows the Shapes text image. Part (b) of the figure shows the gradient magnitude image, obtained by combining, as in Equation 5.9, the horizontal and vertical derivatives, shown in parts (c) and (d) of the figure, respectively. The image in part (c) presents high values (positive or negative) at vertical edges only; horizontally oriented edges have been deleted by the horizontal derivative operator. The image in part (d) shows high output at horizontal edges, with the vertically oriented edges having been removed by the vertical derivative operator. The test image has strong edges for most of the objects present, which are clearly depicted in the derivative images; however, the derivative images show the edges of a few objects that are not readily apparent in the original image as well. Parts (e) and (f) of the figure show the derivatives at 45° and 135° , respectively; the images indicate the diagonal edges present in the image.

Figures 5.3, 5.4, and 5.5 show similar sets of results for the clock, the knee MR, and the chest X-ray test images, respectively. In the derivatives of the clock image, observe that the numeral “1” has been obliterated by the vertical derivative operator [Figure 5.3 (d)], but gives rise to high output values for the horizontal derivative [Figure 5.3 (c)]. The clock image has the minute hand oriented at approximately 135° with respect to the horizontal; this feature has been completely removed by the 135° derivative operator, as shown in Figure 5.3 (f), but has been enhanced by the 45° derivative operator, as shown in Figure 5.3 (e). The knee MR image contains sharp boundaries that are depicted well in the derivative images in Figure 5.4. The derivative images of the chest X-ray image in Figure 5.5 indicate large values at the boundaries of the image, but depict the internal details with weak derivative values, indicative of the smooth nature of the image.


5.3.2 The Laplacian of Gaussian

Although the Laplacian is a gradient operator, it should be recognized that it is a *second-order* difference operator. As we observed in Sections 2.11.1 and 4.6, this leads to double-edged outputs with positive and negative values at each edge; this property is demonstrated further by the example in Figure 5.6 (see also Figure 4.26). The Laplacian has the advantage of being omnidirectional, that is, being sensitive to edges in all directions; however, it is not possible to derive the angle of an edge from the result. The operator is also sensitive to noise because there is no averaging included in the operator; the gain in the frequency domain increases quadratically with frequency, causing significant amplification of high-frequency noise components. For these reasons, the Laplacian is not directly useful in edge detection.


The double-edged output of the Laplacian indicates an important property of the operator: the result possesses a zero-crossing in between the positive and negative outputs across an edge; the property holds even when the edge in the original image is significantly blurred. This property is useful in the development of robust edge detectors. The noise sensitivity of the Laplacian

**FIGURE 5.2**

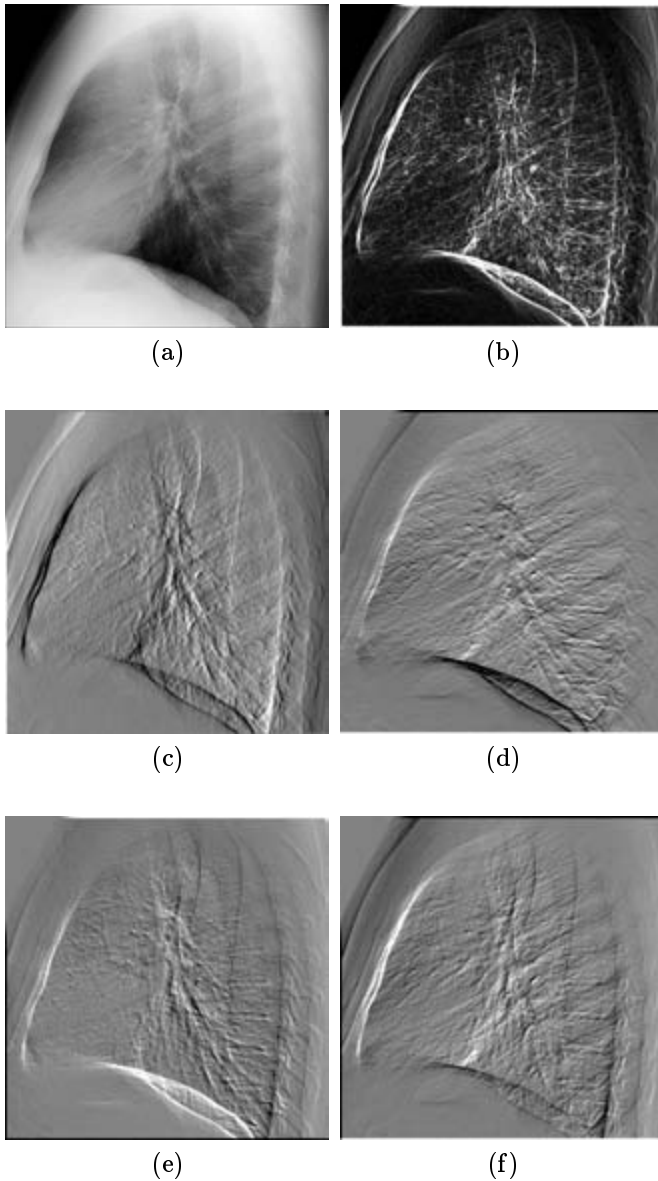
(a) Shapes test image. (b) Gradient magnitude, display range $[0, 400]$ out of $[0, 765]$. (c) Horizontal derivative, display range $[-200, 200]$ out of $[-765, 765]$. (d) Vertical derivative, display range $[-200, 200]$ out of $[-765, 765]$. (e) 45° derivative, display range $[-200, 200]$ out of $[-765, 765]$. (f) 135° derivative, display range $[-200, 200]$ out of $[-765, 765]$.

**FIGURE 5.3**

(a) Clock test image. (b) Gradient magnitude, display range $[0, 100]$ out of $[0, 545]$. (c) Horizontal derivative, display range $[-100, 100]$ out of $[-538, 519]$. (d) Vertical derivative, display range $[-100, 100]$ out of $[-446, 545]$. (e) 45° derivative, display range $[-100, 100]$ out of $[-514, 440]$. (f) 135° derivative, display range $[-100, 100]$ out of $[-431, 535]$.

**FIGURE 5.4**


(a) Knee MR image. (b) Gradient magnitude, display range $[0, 400]$ out of $[0, 698]$. (c) Horizontal derivative, display range $[-200, 200]$ out of $[-596, 496]$. (d) Vertical derivative, display range $[-200, 200]$ out of $[-617, 698]$. (e) 45° derivative, display range $[-200, 200]$ out of $[-562, 503]$. (f) 135° derivative, display range $[-200, 200]$ out of $[-432, 528]$.

**FIGURE 5.5**


(a) Part of a chest X-ray image. (b) Gradient magnitude, display range $[0, 50]$ out of $[0, 699]$. (c) Horizontal derivative, display range $[-50, 50]$ out of $[-286, 573]$. (d) Vertical derivative, display range $[-50, 50]$ out of $[-699, 661]$. (e) 45° derivative, display range $[-50, 50]$ out of $[-452, 466]$. (f) 135° derivative, display range $[-50, 50]$ out of $[-466, 442]$.



profile of edge



first derivative



second derivative

FIGURE 5.6

Top to bottom: A profile of a blurred object showing two edges, the first derivative, and the second derivative (see also Figure 4.26).

may be reduced by including a smoothing operator. A scalable smoothing operator could be defined in terms of a 2D Gaussian function, with the variance controlling the spatial extent or width of the smoothing function. Combining the Laplacian and the Gaussian, we obtain the popular Laplacian-of-Gaussian or LoG operator [8, 122, 281, 282].

Consider the Gaussian specified by the function

$$g(x, y) = -\exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right). \quad (5.18)$$

The usual normalizing scale factor has been left out. Taking partial derivatives with respect to x and y , we obtain

$$\begin{aligned} \frac{\partial^2 g}{\partial x^2} &= -\frac{x^2 - \sigma^2}{\sigma^4} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right) \\ \frac{\partial^2 g}{\partial y^2} &= -\frac{y^2 - \sigma^2}{\sigma^4} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right) \end{aligned} \quad (5.19)$$

which leads to

$$\nabla^2 g(x, y) = \text{LoG}(r) = -\frac{r^2 - 2\sigma^2}{\sigma^4} \exp\left(-\frac{r^2}{2\sigma^2}\right), \quad (5.20)$$

where $r = \sqrt{x^2 + y^2}$. The LoG function is isotropic and has positive and negative values. Due to its shape, it is often referred to as the Mexican hat or sombrero.

Figure 5.7 shows the LoG operator in image and mesh-plot formats; the basic Gaussian used to derive the LoG function is also shown for reference. The Fourier magnitude spectra of the Gaussian and LoG functions are also shown in the figure. It should be observed that, whereas the Gaussian is a lowpass filter (which is also a 2D Gaussian in the frequency domain), the LoG function is a bandpass filter. The width of the filters is controlled by the parameter σ of the Gaussian.

Figure 5.8 shows profiles of the LoG and the related Gaussian for two values of σ . Figure 5.9 shows the profiles of the Fourier transforms of the functions in Figure 5.8. The profiles clearly demonstrate the nature of the functions and their filtering characteristics.

An approximation to the LoG operator is provided by taking the difference between two Gaussians of appropriate variances: this operator is known as the difference-of-Gaussians or DoG operator [282].

Examples: Figure 5.10 shows the Shapes test image, the LoG of the image with $\sigma = 1$ pixel, and the locations of the zero-crossings of the LoG of the image with $\sigma = 1$ pixel and $\sigma = 2$ pixels. The zero-crossings indicate the locations of the edges in the image. The use of a large value for σ reduces the effect of noise, but also causes smoothing of the edges and corners, as well as the loss of the minor details present in the image.






FIGURE 5.7

The Laplacian of Gaussian in (b) image format and (d) as a mesh plot. The related Gaussian functions are shown in (a) and (c). The size of the arrays is 51×51 pixels; standard deviation $\sigma = 4$ pixels. The Fourier magnitude spectra of the functions are shown in (e) and (f).



(a)



(b)

FIGURE 5.8

Profiles of the Laplacian of Gaussian (solid line) and the related Gaussian (dashed line) in Figure 5.7. The functions have been normalized to a maximum value of unity. The unit of r is pixels. (a) $\sigma = 4$ pixels. (b) $\sigma = 2$ pixels.




FIGURE 5.9

Profiles of the Fourier magnitude spectra of the Laplacian of Gaussian (solid line) and the related Gaussian (dashed line) in Figure 5.7. Both functions have been normalized to have a maximum value equal to unity. (a) $\sigma = 4$ pixels. (b) $\sigma = 2$ pixels. The zero-frequency point is at the center of the horizontal axis.

Figure 5.11 shows the clock image, its LoG, and the zero-crossings of the LoG with $\sigma = 1$ pixel and $\sigma = 2$ pixels. The results illustrate the performance of the LoG operator in the presence of noise.

Figures 5.12, 5.13, and 5.14 show similar sets of results for the myocyte image, the knee MR image, and the chest X-ray test images. Comparative analysis of the scales of the details present in the images and the zero-crossings of the LoG for different values of σ indicates the importance of selecting values of the σ parameter in accordance with the scale of the details to be detected.

5.3.3 Scale-space methods for multiscale edge detection

Marr and Hildreth [281, 282] suggested that physical phenomena may be detected simultaneously over several channels tuned to different spatial sizes or *scales*, with an approach known as the spatial coincidence. An intensity change that is due to a single physical phenomenon is indicated by zero-crossing segments present in independent channels over a certain range of scales, with the segments having the same position and orientation in each channel. A significant intensity change indicates the presence of a major event that is registered as a physical boundary, and is recognized as a single physical phenomenon. The boundaries of a significant physical pattern should be present over several channels, suggesting that the use of techniques based on zero-crossings generated from filters of different scales could be more effective than the conventional (single-scale) methods for edge detection; see, for example, Figure 5.13.


Zero-crossings and scale-space: The multichannel model for the HVS [283] and the Marr-Hildreth spatial coincidence assumption [281] led to the development of methods for the detection of edges based upon multiscale analysis performed with filters of different scales. Marr and Hildreth proposed heuristic rules to combine information from the different channels in a multichannel vision model; they suggested the use of a bank of LoG filters with several values of σ , which may be represented as $\{\nabla^2 g(x, y; \sigma)\}$, with $\sigma > 0$.

A method for obtaining information in images across a continuum of scales was suggested by Witkin [284], who introduced the concept of scale-space. The method rapidly gained considerable interest, and has been explored further by several researchers in image processing and analysis [285, 286, 287, 288, 289]. The scale-space $\Psi(x, y; \sigma)$ of an image $f(x, y)$ is defined as the set of all zero-crossings of its LoG:


$$\begin{aligned} \{\Psi(x, y; \sigma)\} &= \{(x, y; \sigma) \mid \zeta(x, y; \sigma) = 0, \\ &\quad \left(\frac{\partial \zeta}{\partial x}\right)^2 + \left(\frac{\partial \zeta}{\partial y}\right)^2 \neq 0, \sigma > 0, \end{aligned} \quad (5.21)$$

where


$$\zeta(x, y; \sigma) = \{\nabla^2 g(x, y; \sigma) * f(x, y)\}. \quad (5.22)$$

**FIGURE 5.10**


(a) The Shapes test image. (b) The LoG of the image in (a) with $\sigma = 1$ pixel. (c) Locations of the zero-crossings in the LoG in (b). (d) Locations of the zero-crossings in the LoG with $\sigma = 2$ pixels.

**FIGURE 5.11**


(a) The clock test image. (b) The LoG of the image in (a) with $\sigma = 1$ pixel. (c) Locations of the zero-crossings in the LoG in (b). (d) Locations of the zero-crossings in the LoG with $\sigma = 2$ pixels.

**FIGURE 5.12**

(a) Image of a myocyte. (b) The LoG of the image in (a) with $\sigma = 2$ pixels. (c) Locations of the zero-crossings in the LoG in (b). (d) Locations of the zero-crossings in the LoG with $\sigma = 4$ pixels.

**FIGURE 5.13**

(a) MR image of a knee. (b) The LoG of the image in (a) with $\sigma = 2$ pixels. (c) – (f) Locations of the zero-crossings in the LoG with $\sigma = 1, 2, 3,$ and 4 pixels.

**FIGURE 5.14**

(a) Part of a chest X-ray image. (b) The LoG of the image in (a) with $\sigma = 2$ pixels; display range $[-150, 150]$ out of $[-231, 956]$. (c) Locations of the zero-crossings in the LoG in (b). (d) Locations of the zero-crossings in the LoG with $\sigma = 4$ pixels.

As the scale σ varies from 0 to ∞ , the set $\{\Psi(x, y; \sigma)\}$ forms continuous surfaces in the $(x, y; \sigma)$ scale-space.

Several important scale-space concepts apply to 1D and 2D signals. It has been shown that the scale-space of almost all signals filtered by a Gaussian determines the signal uniquely up to a scaling constant [285] (except for noise-contaminated signals and some special functions [290]). The importance of this property lies in the fact that, theoretically, for almost all signals, no information is lost by working in the scale-space instead of the image domain. This property plays an important role in image understanding [291], image reconstruction from zero-crossings [285, 292], and image analysis using the scale-space approach [288]. Furthermore, it has also been shown that the Gaussian does not create additional zero-crossings as the scale σ increases beyond a certain limit, and that the Gaussian is the only filter with this desirable scaling behavior [285].

Based on the spatial-coincidence assumption, Witkin [284] proposed a 1D stability analysis method for the extraction of primitive events that occur over a large range of scales. The primitive events were organized into a qualitative signal description representing the major events in the signal. Assuming that zero-crossing curves do not cross one another (which was later proven to be incorrect by Katz [293]), Witkin defined the stability of a signal interval as the scale range over which the signal interval exists; major events could then be captured via stability analysis. However, due to the complex topological nature of spatial zero-crossings, it is often difficult to directly extend Witkin's 1D stability analysis method to 2D image analysis. The following problems affect Witkin's method for stability analysis:

- It has been shown that zero-crossing curves do cross one another [293].
- It has been shown that real (authentic) zero-crossings could turn into false (phantom) zero-crossings as the scale σ increases [294]. Use of the complete scale-space (with σ ranging from 0 to ∞) may introduce errors in certain applications; an appropriate scale-space using only a finite range of scales could be more effective.
- For 2D signals (images), the scale-space consists of zero-crossing surfaces that are more complex than the zero-crossing curves for 1D signals. The zero-crossing surfaces may split and merge as the scale varies (decreases or increases, respectively).
- As a consequence of the above, there is no simple topological region associated with a zero-crossing surface, and tracing a zero-crossing surface across scales becomes computationally difficult.

Liu et al. [295] proposed an alternative definition of zero-crossing surface stability in terms of important spatial boundaries. In this approach, a spatial boundary is defined as a region of steep gradient and high contrast, and is well-defined if it has no neighboring boundaries within a given range. This

definition of spatial boundaries is consistent with the Marr-Hildreth spatial-coincidence assumption. Furthermore, stability maps [288] associated with the scale-space are used. A relaxation algorithm is included in the process to generate zero-crossing maps.

In the method of Liu et al. [295], the discrete scale-space approach is used to construct a representation of a given image in terms of a stability map, which is a measure of pattern boundary persistence over a range of filter scales. For a given image $f(x, y)$, a set of zero-crossing maps is generated by convolving the image with the set of isotropic functions $\nabla^2 g(x, y; \sigma_i)$, $1 \leq i \leq N$. It was indicated that $N = 8$ sampled σ_i values ranging from 1 to 8 pixels were adequate for the application considered. Ideally, one would expect a pattern boundary to be accurately located over all of the scales. However, it has been shown [296, 297] that the accuracy of zero-crossing localization depends upon the width of the central excitatory region of the filter (defined as $w_i = 2\sqrt{2} \sigma_i$ [298]). Chen and Medioni [299] proposed a 1D method for localization of zero-crossings that works well for ideal step edges and image patterns with sharp contrast; however, the method may not be effective for the construction of the spatial scale-space for real-life images with poor and variable contrast. Instead of directly matching all the zero-crossing locations at a point (x, y) over the zero-crossing maps, Liu et al. proposed a criterion $C(\sigma_i)$ that is a function of the scale σ_i to define a neighborhood in which the matching procedure is performed at a particular scale:

$$C(\sigma_i) = \{(x', y') \mid x - \lambda\sigma_i \leq x' \leq x + \lambda\sigma_i, \\ y - \lambda\sigma_i \leq y' \leq y + \lambda\sigma_i, \lambda \leq 1, \quad (5.23)$$

where (x', y') are the actual locations of the zero-crossings, (x, y) is the pixel location at which the filters are being applied, and λ is a constant to be determined experimentally ($\lambda = 1$ was used by Liu et al.). Therefore, if a zero-crossing $\psi(x, y; \sigma_i)$ is found in the neighborhood defined by $C(\sigma_i)$, an arbitrary constant κ is assigned to a function $S_i(x, y)$, which otherwise is assigned a zero, that is,

$$S_i(x, y) = \begin{cases} \kappa & \text{if } \psi(x, y; \sigma_i) \in C(\sigma_i) \\ 0 & \text{otherwise,} \end{cases} \quad (5.24)$$

where the subscript i corresponds to the i^{th} scale σ_i .

Applying Equations 5.23 and 5.24 to the set of zero-crossings detected, a set of adjusted zero-crossing maps $\{S_1(x, y), S_2(x, y), \dots, S_N(x, y)\}$ is obtained, where N is the number of scales. The adjusted zero-crossing maps are used to construct the zero-crossing stability map $\chi(x, y)$ as

$$\chi(x, y) = \sum_{i=1}^N S_i(x, y). \quad (5.25)$$

The values of $\chi(x, y)$ are, in principle, a measure of boundary stability through the filter scales. Marr and Hildreth [281] and Marr and Poggio [300] suggested

that directional detection of zero-crossings be performed after the LoG operator has been applied to the image.

According to the spatial-coincidence assumption, a true boundary should be high in contrast and have relatively large χ values at the corresponding locations. Furthermore, there should be no other edges within a given neighborhood. Thus, if in a neighborhood of $\chi(x, y)$, nonzero stability map values exist only along the orientation of a local segment of the stability map that crosses (x, y) , then $\chi(x, y)$ may be considered to signify a stable edge pixel at (x, y) . On the other hand, if many nonzero stability map values are present at different directions, $\chi(x, y)$ indicates an insignificant boundary pixel at (x, y) . In other words, a consistent stability indexing method (in the sense of the spatial-coincidence assumption) should take neighboring stability indices into account. Based upon this argument, Liu et al. proposed a relative stability index $\mu(x, y)$ computed from the stability map where $\chi(x, y) \neq 0$, as follows.

In a neighborhood of $\chi(x, y)$, if m nonzero values are found, $\chi(x, y)$ is relabeled as l_0 , and the rest of $\chi(x_k, y_k)$ are relabeled as l_k , $k = 1, \dots, m - 1$; see Figure 5.15. In order to avoid using elements in the neighborhood that belong to the same edge, those $\chi(x', y')$ having the same orientation as that of l_0 are not included in the computation of $\mu(x, y)$. Based upon these requirements, the relative stability index $\mu(x, y)$ is defined as

$$\mu(x, y) = \frac{l_0}{\sum_{k=0}^{m-1} \rho_k l_k}, \quad (5.26)$$

where $\rho_k = \exp(-d_k^2)$ and $d_k = \sqrt{(x - x_k)^2 + (y - y_k)^2}$, and (x_k, y_k) are the locations of l_k . It should be noted that $0 < \mu(x, y) \leq 1$, and that the value of $\mu(x, y)$ is governed by the geometrical distribution of the neighboring stability index values.

Stability of zero-crossings: Liu et al. [295] observed that the use of zero-crossings to indicate the presence of edges is reliable as long as the edges are well-separated; otherwise, the problem of false zero-crossings could arise [301]. The problem with using zero-crossings to localize edges is that the zeros of the second derivative of a function localize the extrema in the first derivative of the function; the extrema include the local minima and maxima in the first derivative of the function, whereas only the local maxima indicate the presence of edge points. Intuitively, those zero-crossings that correspond to the minima of the first derivative are not associated with edge points at all. In image analysis based upon the notion of zero-crossings, it is desirable to be able to distinguish real zero-crossings from false ones, and to discard the false zero-crossings. Motivated by the work of Richter and Ullman [301], Clark [294] conducted an extensive study on the problem of false and real zero-crossings, and proposed that zero-crossings may be classified as real if $\xi(x, y) < 0$ and false if $\xi(x, y) > 0$, where $\xi(x, y) = \nabla[\nabla^2 p(x, y)] \bullet \nabla p(x, y)$, where \bullet denotes the dot product, $p(x, y)$ is a smoothed version of the given image [such as $p(x, y) = g(x, y; \sigma) * f(x, y)$], $\nabla p(x, y) = \left[\frac{\partial p}{\partial x}, \frac{\partial p}{\partial y} \right]^T$, and $\nabla[\nabla^2 p(x, y)] =$




FIGURE 5.15

A case where three zero-crossings $\{l_1, l_2, l_3\}$ are found in a neighborhood of a zero-crossing l_0 . d_i indicates the distance from l_i to l_0 . The arrows indicate the directions of the zero-crossings. (a) The neighboring zero-crossings are far apart from l_0 , imposing a low penalty to the zero-crossing associated with l_0 . (b) The neighboring zero-crossings are close to l_0 , imposing a high penalty to the zero-crossing associated with l_0 . Reproduced with permission from Z.-Q. Liu, R.M. Rangayyan, and C.B. Frank, "Statistical analysis of collagen alignment in ligaments by scale-space analysis", *IEEE Transactions on Biomedical Engineering*, 38(6):580–588, 1991. © IEEE.

$\left[\frac{\partial^3 p}{\partial x^3} + \frac{\partial^3 p}{\partial x \partial y^2}, \frac{\partial^3 p}{\partial x^2 \partial y} + \frac{\partial^3 p}{\partial y^3} \right]^T$. Liu et al. included this step in their method to detect true zero-crossings.

Example: Figure 5.16 (a) shows an SEM image of collagen fibers in a ligament scar-tissue sample. Parts (b) – (d) of the figure show the zero-crossing maps obtained with $\sigma = 1, 4$, and 8 pixels. The result in (b) contains several spurious or insignificant zero-crossings, whereas that in (d) contains smoothed edges of only the major regions of the image. Part (e) shows the stability map, which indicates the edges of the major objects in the image. The stability map was used to detect the collagen fibers in the image. See Section 8.5 for details on directional analysis of oriented patterns by further processing of the stability map. Methods for the directional analysis of collagen fibers are described in Section 8.7.1.

See Sections 5.10.2, 8.4, and 8.9 for more examples on multiscale analysis.


5.3.4 Canny's method for edge detection

Canny [302] proposed an approach for edge detection based upon three criteria for good edge detection, multidirectional derivatives, multiscale analysis, and optimization procedures. The three criteria relate to low probabilities of false edge detection and missing real edges, represented in the form of an SNR; good localization, represented by the RMS distance of the detected edge from the true edge; and the production of a single output for a single edge, represented by the distance between the adjacent maxima in the output. A basic filter derived using the criteria mentioned above was approximated by the first derivative of a Gaussian. Procedures were proposed to incorporate multiscale analysis and directional filters to facilitate efficient detection of edges at all orientations and scales including adaptive thresholding with hysteresis.

The LoG filter is nondirectional, whereas Canny's method selectively evaluates a directional derivative across each edge. By avoiding derivatives at other angles that would not contribute to edge detection but increase the effects of noise, Canny's method could lead to better results than the LoG filter.

5.3.5 Fourier-domain methods for edge detection

In Section 4.7, we saw that highpass filters may be applied in the Fourier-domain to extract the edges in the given image. However, the inclusion of all of the high-frequency components present in the image could lead to noisy results. Reduction of high-frequency noise suggests the use of bandpass filters, which may be easily implemented as a cascade of a lowpass filter with a highpass filter. In the frequency domain, such a cascade of filters results in the multiplication of the corresponding transfer functions. Because edges are often weak or blurred in images, some form of enhancement of the corresponding frequency components would also be desirable. This argument leads us to the LoG filter: a combination of the Laplacian, which is a high-frequency-

**FIGURE 5.16**

(a) SEM image of a ligament scar-tissue sample. (b) – (d) Zero-crossing locations detected using the LoG operator with $\sigma = 1, 4,$ and 8 pixels, respectively. (e) The stability map, depicting the major edges present in the image. Reproduced with permission from Z.-Q. Liu, R.M. Rangayyan, and C.B. Frank, “Statistical analysis of collagen alignment in ligaments by scale-space analysis”, *IEEE Transactions on Biomedical Engineering*, 38(6):580–588, 1991. © IEEE.

emphasis filter with its gain quadratically proportional to frequency, with a Gaussian lowpass filter. The methods and results presented in Section 5.3.2 demonstrate the edge-detection capabilities of the LoG filter, which may be easily implemented in the frequency domain. Frequency-domain implementation using the FFT may be computationally advantageous when the LoG function is specified with a large spatial array, which would be required in the case of large values of σ .

Several other line-detection and edge-detection methods, such as Gabor filters (see Sections 5.10, 8.4, 8.9, and 8.10) and fan filters (see Section 8.3) may also be implemented in the frequency domain with advantages.

5.3.6 Edge linking

The results of most methods for edge detection are almost always discontinuous, and need to be processed further to link disjoint segments and obtain complete representations of the boundaries of ROIs. Two principal properties that may be used to establish the similarity of edge pixels from gradient images are the following [8]:

- The strength of the gradient — a point (x', y') in a neighborhood of (x, y) is similar in gradient magnitude to the point (x, y) if

$$\|\mathbf{G}(x, y) - \mathbf{G}(x', y')\| \leq T, \quad (5.27)$$

where $\mathbf{G}(x, y)$ is the gradient vector of the given image $f(x, y)$ at (x, y) and T is a threshold.

- The direction of the gradient — a point (x', y') in a neighborhood of (x, y) is similar in gradient direction to the point (x, y) if

$$\begin{aligned} |\alpha(x, y) - \alpha(x', y')| &\leq A, \\ \alpha(x, y) = \angle \mathbf{G}(x, y) &= \tan^{-1} \left\{ \frac{\partial f(x, y) / \partial y}{\partial f(x, y) / \partial x} \right\}, \end{aligned} \quad (5.28)$$

where A is a threshold.

3×3 or 5×5 neighborhoods may be used for checking pixels for similarity in their gradients as above. Further processing steps may include linking of edge segments separated by small breaks and deleting isolated short segments.

See Section 5.10.2 (page 493) for the details of an edge analysis method known as edge-flow propagation.

5.4 Segmentation and Region Growing

Dividing an image into regions that could correspond to structural units, objects of interest, or ROIs is an important prerequisite for most techniques for image analysis. Whereas a human observer may, by merely looking at a displayed image, readily recognize its structural components, computer analysis of an image requires algorithmic analysis of the array of image pixel values before arriving at conclusions about the content of the image. Computer analysis of images usually starts with segmentation, which reduces pixel data to region-based information about the objects and structures present in the image [303, 304, 305, 306, 307].

Image segmentation techniques may be classified into four main categories:

- thresholding techniques [8, 306],
- boundary-based methods [303, 308],
- region-based methods [8, 123, 274, 309, 310, 311, 312], and
- hybrid techniques [313, 314, 315, 316, 317] that combine boundary and region criteria.

Thresholding methods are based upon the assumption that all pixels whose values lie within a certain range belong to the same class; see Section 5.1. The threshold may be determined based upon the valleys in the histogram of the image; however, identifying thresholds to segment objects is not easy even with optimal thresholding techniques [8, 306]. Moreover, because thresholding algorithms are solely based upon pixel values and neglect all of the spatial information in the image, their accuracy of segmentation is limited; furthermore, thresholding algorithms do not cope well with noise or blurring at object boundaries.

Boundary-based techniques make use of the property that, usually, pixel values change rapidly at the boundaries between regions. The methods start by detecting intensity discontinuities lying at the boundaries between objects and their backgrounds, typically through a gradient operation. High values of the output provide candidate pixels for region boundaries, which must then be processed to produce closed curves representing the boundaries between regions, as well as to remove the effects of noise and discontinuities due to nonuniform illumination and other effects. Although edge-linking algorithms have been proposed to assemble edge pixels into a meaningful set of object boundaries (see Section 5.3.6), such as local similarity analysis, Hough-transform-based global analysis, and global processing via graph-theoretic techniques [8], the accurate conversion of disjoint sets of edge pixels to closed-loop boundaries of ROIs is a difficult task.

Region-based methods, which are complements of the boundary-based approach, rely on the postulate that neighboring pixels within a region have similar values. Region-based segmentation algorithms may be divided into two groups: *region splitting and merging* and *region growing*.

Segmentation techniques in the region splitting and merging category initially subdivide the given image into a set of arbitrary, disjoint regions, and then merge and/or split the regions in an attempt to satisfy some prespecified conditions.

Region growing is a procedure that groups pixels into regions. The simplest of region-growing approaches is pixel aggregation, which starts with a seed pixel and grows a region by appending spatially connected neighboring pixels that meet a certain homogeneity criterion. Different homogeneity criteria will lead to regions with different characteristics. It is important, as well as difficult, to select an appropriate homogeneity criterion in order to obtain regions that are appropriate for the application on hand.

Typical algorithms in the group of hybrid techniques refine image segmentation by integration of boundary and region information; proper combination of boundary and region information may produce better segmentation results than those obtained by either method on its own. For example, the morphological watershed method [315] is generally applied to a gradient image, which can be viewed as a topographic map with boundaries between regions as ridges. Consequently, segmentation is equivalent to flooding the topography from the seed pixels, with region boundaries being erected to keep water from the different seed pixels from merging. Such an algorithm is guaranteed to produce closed boundaries, which is known to be a major problem with boundary-based methods. However, because the success of this type of an algorithm relies on the accuracy of the edge-detection procedure, it encounters difficulties with images in which regions are both noisy and have blurred or indistinct boundaries. Another interesting method within this category, called variable-order surface fitting [313], starts with a coarse segmentation of the given image into several surface-curvature-sign primitives (for example, pit, peak, and ridge), which are then refined by an iterative region-growing method based on variable-order surface fitting. This method, however, may only be suitable to the class of images where the image contents vary considerably.

The main difficulty with region-based segmentation schemes lies in the selection of a homogeneity criterion. Region-based segmentation algorithms have been proposed using statistical homogeneity criteria based on regional feature analysis [312], Bayesian probability modeling of images [318], Markov random fields [319], and seed-controlled homogeneity competition [311]. Segmentation algorithms could also rely on homogeneity criteria with respect to gray level, color, texture, or surface measures.

5.4.1 Optimal thresholding

Suppose it is known *a priori* that the given image consists of only two principal brightness levels with the prior probabilities P_1 and P_2 . Consider the situation where natural variations or noise modify the two gray levels to distributions represented by Gaussian PDFs $p_1(x)$ and $p_2(x)$, where x represents the gray level. The PDF of the image gray levels is then [8]

$$\begin{aligned}
 p(x) &= P_1 p_1(x) + P_2 p_2(x) \\
 &= \frac{P_1}{\sqrt{2\pi}\sigma_1} \exp\left[-\frac{(x - \mu_1)^2}{2\sigma_1^2}\right] + \frac{P_2}{\sqrt{2\pi}\sigma_2} \exp\left[-\frac{(x - \mu_2)^2}{2\sigma_2^2}\right], \quad (5.29)
 \end{aligned}$$

where μ_1 and μ_2 are the means of the two regions, and σ_1 and σ_2 are their standard deviations. Let $\mu_1 < \mu_2$.

Suppose that the dark regions in the image correspond to the background, and the bright regions to the objects of interest. Then, all pixels below a threshold T may be considered to belong to the background, and all pixels above T may be considered as pixels belonging to the object of interest. The probability of erroneous classification is then

$$P_e(T) = P_1 \int_T^\infty p_1(x) dx + P_2 \int_{-\infty}^T p_2(x) dx. \quad (5.30)$$

To find the optimal threshold, we may differentiate $P_e(T)$ with respect to T and equate the result to zero, which leads to

$$P_1 p_1(T) = P_2 p_2(T). \quad (5.31)$$

Applying this result to the Gaussian PDFs gives (after taking logarithms and some simplification) the quadratic equation [8]

$$AT^2 + BT + C = 0, \quad (5.32)$$

where

$$\begin{aligned}
 A &= \sigma_1^2 - \sigma_2^2, \\
 B &= 2(\mu_1\sigma_2^2 - \mu_2\sigma_1^2), \\
 C &= \sigma_1^2\mu_2^2 - \sigma_2^2\mu_1^2 + 2\sigma_1^2\sigma_2^2 \ln\left(\frac{\sigma_2 P_1}{\sigma_1 P_2}\right). \quad (5.33)
 \end{aligned}$$

The possibility of two solutions indicates that it may require two thresholds to obtain the optimal threshold.

If $\sigma_1^2 = \sigma_2^2 = \sigma^2$, a single threshold may be used, given by

$$T = \frac{\mu_1 + \mu_2}{2} + \frac{\sigma^2}{\mu_1 - \mu_2} \ln \left(\frac{P_2}{P_1} \right). \quad (5.34)$$

Furthermore, if the two prior probabilities are equal, that is, $P_1 = P_2$, or if the variance is zero, that is, $\sigma = 0$, the optimal threshold is equal to the average of the two means.

Thresholding using boundary characteristics: The number of pixels covered by the objects of interest to be segmented from an image is almost always a small fraction of the total number of pixels in the image: the gray-level histogram of the image is then likely to be almost unimodal. The histogram may be made closer to being bimodal if only the pixels on or near the boundaries of the object regions are considered.

The selection and characterization of the edge or boundary pixels may be achieved by using gradient and Laplacian operators as follows [8]:

$$b(x, y) = \begin{cases} 0 & \text{if } \nabla_f(x, y) < T, \\ L_+ & \text{if } \nabla_f(x, y) \geq T \text{ and } \nabla_f^2(x, y) \geq 0, \\ L_- & \text{if } \nabla_f(x, y) \geq T \text{ and } \nabla_f^2(x, y) < 0, \end{cases} \quad (5.35)$$

where $\nabla_f(x, y)$ is a gradient and $\nabla_f^2(x, y)$ is the Laplacian of the given image $f(x, y)$; T is a threshold; and 0 , L_+ , L_- represent three distinct gray levels. In the resulting image, the pixels that are not on an edge are set to zero, the pixels on the darker sides of edges are set to L_+ , and the pixels on the lighter sides of edges are set to L_- . This information may be used not only to detect objects and edges, but also to identify the leading and trailing edges of objects (with reference to the scanning direction).

See Section 8.3.2 for a description of Otsu's method of deriving the optimal threshold for binarizing a given image; see also Section 8.7.2 for discussions on a few other methods to derive thresholds.

5.4.2 Region-oriented segmentation of images

Let R represent the region spanning the entire space of the given image. Segmentation may be viewed as a process that partitions R into n subregions R_1, R_2, \dots, R_n such that [8]

- $\cup_{i=1}^n R_i = R$, that is, the union of all of the regions detected spans the entire image (then, every pixel must belong to a region);
- R_i is a connected region, $i = 1, 2, \dots, n$;
- $R_i \cap R_j = \emptyset \forall i, j, i \neq j$ (that is, the regions are disjoint);

- $\mathcal{P}(R_i) = TRUE$, for $i = 1, 2, \dots, n$ (for example, all pixels within a region have the same intensity);
- $\mathcal{P}(R_i \cup R_j) = FALSE \quad \forall i, j, i \neq j$ (for example, the intensities of the pixels in different regions are different);

where $\mathcal{P}(R_i)$ is a logical predicate defined over the points in the set R_i , and \emptyset is the null set.

A simple algorithm for region growing by pixel aggregation based upon the similarity of a local property is as follows:

- Start with a *seed* pixel (or a set of seed pixels).
- Append to each pixel in the region those of its 4-connected or 8-connected neighbors that have properties (gray level, color, etc.) that are similar to those of the seed.
- Stop when the region cannot be grown any further.

The results of an algorithm as above depend upon the procedure used to select the seed pixels and the measures of similarity or inclusion criteria used. The results may also depend upon the method used to traverse the image; that is, the sequence in which neighboring pixels are checked for inclusion.

5.4.3 Splitting and merging of regions

Instead of using seeds to grow regions or global thresholds to separate an image into regions, one could initially consider to divide the given image arbitrarily into a set of disjoint regions, and then to split and/or merge the regions using conditions or predicates \mathcal{P} .

A general split/merge procedure is as follows [8]: Assuming the image to be square, subdivide the entire image R successively into smaller and smaller quadrant regions such that, for any region R_i , $\mathcal{P}(R_i) = TRUE$. In other words, if $\mathcal{P}(R) = FALSE$, divide the image into quadrants; if \mathcal{P} is *FALSE* for any quadrant, subdivide that quadrant into subquadrants. Iterate the procedure until no further changes are made, or a stopping criterion is reached. The splitting technique may be represented as a quadtree. Difficulties could exist in selecting an appropriate predicate \mathcal{P} .

Because the splitting procedure could result in adjacent regions that are similar, a merging step would be required, which may be specified as follows: Merge two adjacent regions R_i and R_k if $\mathcal{P}(R_i \cup R_k) = TRUE$. Iterate until no further merging is possible.

5.4.4 Region growing using an additive tolerance

A commonly used region-growing scheme is pixel aggregation [8, 123]. The method compares the properties of spatially connected neighboring pixels with

those of the seed pixel; the properties used are determined by homogeneity criteria. For intensity-based image segmentation, the simplest property is the pixel gray level. The term “additive tolerance level” stands for the permitted absolute gray-level difference between the neighboring pixels and the seed pixel: a neighboring pixel $f(m, n)$ is appended to the region if its absolute gray-level difference with respect to the seed pixel is within the additive tolerance level T :

$$|f(m, n) - \text{seed}| \leq T. \quad (5.36)$$

Figure 5.17 shows a simple example of additive-tolerance region growing using different seed pixels in a 5×5 image. The additive tolerance level used in the example is $T = 3$. Observe that two different regions are obtained by starting with two seeds at different locations as shown in Figure 5.17 (b) and Figure 5.17 (c). In order to overcome this dependence of the region on the seed pixel selected, the following modified criterion could be used to determine whether a neighboring pixel should be included in a region or not: instead of comparing the incoming pixel with the gray level of the seed, the gray level of a neighboring pixel is compared with the mean gray level, called the *running mean* μ_{R_c} , of the region being grown at its current stage, R_c . This criterion may be represented as

$$|f(m, n) - \mu_{R_c}| \leq T, \quad (5.37)$$

where

$$\mu_{R_c} = \frac{1}{N_c} \sum_{(m,n) \in R_c} f(m, n), \quad (5.38)$$

where N_c is the number of pixels in R_c . Figure 5.17 (d) shows the result obtained with the running-mean algorithm by using the same additive tolerance level as before ($T = 3$). With the running-mean criterion, no matter which pixel is selected as the seed, the same final region is obtained in the present example, as long as the seed pixel is within the region, which is the central highlighted area in Figure 5.17 (d).

In the simple scheme described above, the seed pixel is always used to check the incoming neighboring pixels, even though most of them are not spatially connected or close to the seed. Such a region-growing procedure may fail when a seed pixel is inappropriately located at a noisy pixel. Shen [320, 321, 322] suggested the use of the “current center” pixel as the reference instead of the seed pixel that was used to commence the region-growing procedure. For example, the shaded area shown in Figure 5.18 represents a region being grown. After the pixel C is appended to the region, its 4-connected neighbors (labeled as Ni , $i = 1, 2, 3, 4$) or 8-connected neighbors (marked as Ni , $i = 1, 2, \dots, 8$) would be checked for inclusion in the region, using

$$|Ni - C| \leq T. \quad (5.39)$$

The pixel C is called the current center pixel. However, because some of the neighboring pixels ($N1$ and $N5$ in the illustration in Figure 5.18) are

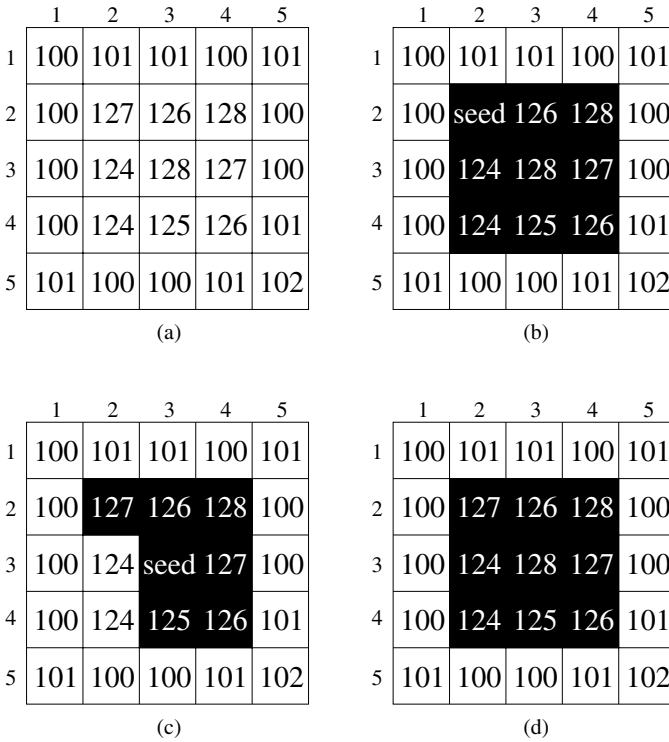


FIGURE 5.17

Example of additive-tolerance region growing using different seed pixels ($T = 3$). (a) Original image. (b) The result of region growing (shaded in black) with the seed pixel at (2, 2). (c) The result of region growing with the seed pixel at (3, 3). (d) The result of region growing with the running-mean algorithm or the “current center pixel” method using any seed pixel within the highlighted region. Figure courtesy of L. Shen [320].

already included in the region shown, only $N2$, $N3$, and $N4$ in the case of 4-connectivity, or $N2$, $N3$, $N4$, $N6$, $N7$, and $N8$ in the case of 8-connectivity are compared with their current center pixel C for region growing, rather than with the original seed pixel. For the example shown in Figure 5.17, this procedure generates the same result as shown in Figure 5.17 (d) independent of the location of the seed pixel (within the ROI) when using the same additive tolerance level ($T = 3$).

seed							
		$N5$	$N1$	$N6$			
		$N2$	C	$N3$			
		$N7$	$N4$	$N8$			

FIGURE 5.18

Illustration of the concept of the “current center pixel” in region growing. Figure courtesy of L. Shen [320].

5.4.5 Region growing using a multiplicative tolerance

In addition to the sensitivity of the region to seed pixel selection with additive-tolerance region growing, the additive tolerance level or absolute difference in gray level T is not a good criterion for region growing: an additive tolerance level of 3, while appropriate for a seed pixel value or running mean of, for example, 127, may not be suitable when the seed pixel gray level or running mean is at a different level, such as 230 or 10. In order to address this problem, a relative difference, based upon a *multiplicative tolerance level* τ , could be employed. Then, the criterion for region growing could be defined as

$$\frac{|f(m, n) - \mu_{R_c}|}{\mu_{R_c}} \leq \tau \quad (5.40)$$

or

$$2 \frac{|f(m, n) - \mu_{R_c}|}{f(m, n) + \mu_{R_c}} \leq \tau, \quad (5.41)$$

where $f(m, n)$ is the gray level of the current pixel being checked for inclusion, and μ_{R_c} could stand for the original seed pixel value, the current center pixel value, or the running-mean gray level. Observe that the two equations above are comparable to the definitions of simultaneous contrast in Equations 2.7 and 2.8.

The additive and multiplicative tolerance levels both determine the maximum gray-level deviation allowed within a region, and any deviation less than this level is considered to be an intrinsic property of the region, or to be noise. Multiplicative tolerance is meaningful when related to the SNR of a region (or image), whereas additive tolerance has a direct connection with the standard deviation of the pixels within the region or a given image.

5.4.6 Analysis of region growing in the presence of noise

In order to analyze the performance of region-growing methods in the presence of noise, let us assume that the given image \mathbf{g} may be modeled as an ideal image \mathbf{f} plus a noise image $\boldsymbol{\eta}$, where \mathbf{f} consists of a series of strictly uniform, disjoint, or nonoverlapping regions R_i , $i = 1, 2, \dots, k$, and $\boldsymbol{\eta}$ includes their corresponding noise parts η_i , $i = 1, 2, \dots, k$. Mathematically, the image may be expressed as

$$\mathbf{g} = \mathbf{f} + \boldsymbol{\eta}, \tag{5.42}$$

where

$$\mathbf{f} = \bigcup_i R_i, \quad i = 1, 2, \dots, k, \tag{5.43}$$

and

$$\boldsymbol{\eta} = \bigcup_i \eta_i, \quad i = 1, 2, \dots, k. \tag{5.44}$$

A strictly uniform region R_i is composed of a set of connected pixels $f(m, n)$ at positions (m, n) whose values equal a constant κ_i , that is,

$$R_i = \{ (m, n) \mid f(m, n) = \kappa_i \}. \tag{5.45}$$

The set of regions R_i , $i = 1, 2, \dots, k$, is what we expect to obtain as the result of segmentation. Suppose that the noise parts η_i , $i = 1, 2, \dots, k$, are composed of white noise with zero mean and standard deviation σ_i ; then, we have

$$\mathbf{g} = \bigcup_i (R_i + \eta_i), \quad i = 1, 2, \dots, k, \tag{5.46}$$

and

$$\mathbf{f} = \bigcup_i R_i = \mathbf{g} - \bigcup_i \eta_i, \quad i = 1, 2, \dots, k. \tag{5.47}$$

As a special case, when all the noise components have the same standard deviation σ , that is,

$$\sigma_1 = \sigma_2 = \dots = \sigma_k = \sigma \tag{5.48}$$

and

$$\eta_1 \simeq \eta_2 \simeq \dots \simeq \eta_k \simeq \eta, \quad (5.49)$$

where the symbol \simeq represents statistical similarity, the image \mathbf{f} may be described as

$$\mathbf{g} \simeq \bigcup_i R_i + \eta, \quad i = 1, 2, \dots, k, \quad (5.50)$$

and

$$\mathbf{f} = \bigcup_i R_i \simeq \mathbf{g} - \eta; \quad i = 1, 2, \dots, k. \quad (5.51)$$

Additive-tolerance region growing is well-suited for segmentation of this special type of image, and an additive tolerance level solely determined by σ may be used globally over the image. However, such special cases are rare in real images. A given image generally has to be modeled, as in Equation 5.46, where multiplicative-tolerance region growing may be more suitable, with the expectation that a global multiplicative tolerance level can be derived for all of the regions in the given image. Because the multiplicative tolerance level could be made a function of $\frac{\sigma_i}{\kappa_i}$ that is directly related to the SNR, which can be defined as $10 \log_{10} \frac{\kappa_i^2}{\sigma_i^2} \text{ dB}$ for each individual region R_i , such a global tolerance level can be found if

$$\frac{\sigma_1}{\kappa_1} = \frac{\sigma_2}{\kappa_2} = \dots = \frac{\sigma_k}{\kappa_k}. \quad (5.52)$$

5.4.7 Iterative region growing with multiplicative tolerance

Simple region growing with fixed additive or multiplicative tolerance may provide good segmentation results with images satisfying either Equations 5.48 and 5.49, or Equation 5.52. However, many images do not meet these conditions, and one may have to employ different additive or multiplicative tolerance levels at different locations in the given image. This leads to the problem of finding the appropriate tolerance level for each individual region, which the iterative multitolerance region-growing approach proposed by Shen et al. [274, 320] attempts to solve.

When human observers attempt to recognize an object in an image, they are likely to use the information available in both an object of interest and its surroundings. Unlike most of the reported boundary detection methods, the HVS detects object boundaries based on not only the information around the boundary itself, such as the pixel variance and gradient around or across the boundary, but also on the characteristics of the object. Shen et al. [274, 320] proposed using the information contained within a region as well as its relationship with the surrounding background to determine the appropriate tolerance level to grow a region. Such information could be represented by a set of features characterizing the region and its background. With increasing tolerance levels obtained using a certain step size, it could be expected that the

values of the feature set at successive tolerance levels will be either the same or similar (which means that the corresponding regions are similar) when the region-growing procedure has identified an actual object or region. Suppose that the feature set mentioned above includes M features; then, the feature vector \mathbf{V}_k at the tolerance level k may be expressed as

$$\mathbf{V}_k = [V_{k,1}, V_{k,2}, \dots, V_{k,M}]^T. \quad (5.53)$$

The minimum normalized distance d_{\min} between the feature vectors at successive tolerance levels could be utilized to select the final region:

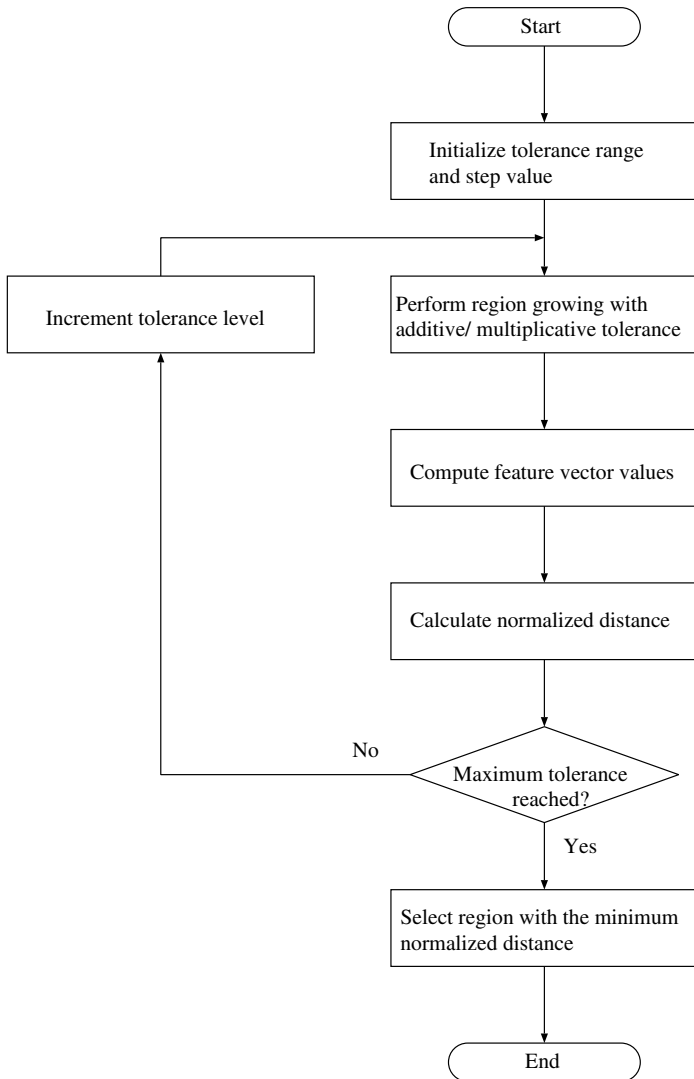
$$d_{\min} = \min_k [d_k] = \min_k \left[\sum_{m=1}^M \frac{(V_{k+1,m} - V_{k,m})^2}{\left(\frac{V_{k+1,m} + V_{k,m}}{2}\right)^2} \right]^{\frac{1}{2}}, \quad (5.54)$$

where $k = 1, 2, \dots, K - 1$, and K is the number of tolerance values used. As an example, the feature set could be chosen to include the coordinates of the centroid of the region, the number of pixels in the region, a region shape descriptor such as compactness, and the ratio of the mean pixel value of the foreground region to the mean pixel value of a suitably defined background.

Figure 5.19 shows the flowchart of the iterative multitolerance region-growing algorithm. The algorithm starts with a selected pixel, called the seed pixel, as the first region pixel. Then, the pixel value $f(m, n)$ of every 4-connected neighbor of the seed (or all pixels belonging to the region, at later stages of the algorithm) is checked for the condition described in Equation 5.40. If the condition is satisfied, the pixel is included in the region. This recursive procedure is continued until no spatially connected pixel meets the condition for inclusion in the region. The outermost layer of connected pixels of the region grown is then treated as the boundary or contour of the region.

In order to permit the selection of the most appropriate tolerance level for each region, the iterative multitolerance region-growing procedure proposed by Shen et al. uses a range of the fractional tolerance value τ , from $\frac{2}{\text{seed}}$ to 0.40 with a step size of $\frac{1}{\text{seed}}$. The specified feature vector is computed for the region obtained at each tolerance level. The normalized distance between the feature vectors for successive tolerance levels is calculated. The feature vector with the minimum distance, as defined in Equation 5.54, is selected and the corresponding region is retained as the final object region.

Figure 5.20 (a) represents an ideal image, composed of a series of strictly uniform regions, with two ROIs: a relatively dark, small triangle, and a relatively bright, occluded circle. Two versions of the image with white noise added are shown in Figure 5.20 (b) (SNR = 30 dB) and Figure 5.20 (c) (SNR = 20 dB). As expected according to the model defined in Equation 5.50, the fixed multiplicative-tolerance region-growing method detects the two ROIs with properly selected tolerance values τ : 0.02 for SNR = 30 dB and 0.05 for SNR = 20 dB; the regions are shown highlighted in black and white in Figures

**FIGURE 5.19**

Flowchart of the iterative, multitolerance, region-growing algorithm. Figure courtesy of L. Shen [320].

5.20 (d) and (e). However, with the composite images shown in Figure 5.20 (f) and (g), which were obtained by appending the two noisy images in parts (b) and (c) of the figure, the fixed-tolerance region-growing approach fails, as stated above. The detected regions in the lower part are incorrect when the growth tolerance level is set to be 0.02, as shown in Figure 5.20 (f), whereas the detection of the occluded circle in the upper part fails with the tolerance value of 0.05, as shown in Figure 5.20 (g). The iterative region-growing technique automatically determines the correct growth tolerance level, and has thereby successfully detected both of the ROIs in the composite image as shown in Figure 5.20 (h).

5.4.8 Region growing based upon the human visual system

Despite its strength in automatically determining an appropriate tolerance level for each region within an image, the iterative multitolerance region-growing algorithm faces two limitations: the iterative procedure is time-consuming, and difficulties exist in the selection of the range of the tolerance value as well as the increment step size. Shen et al. [274, 320, 322] used the range of $\frac{2}{\text{seed}}$ to 0.40 with a step size of $\frac{1}{\text{seed}}$. The motivation for defining tolerance as being inversely related to the seed pixel value was to avoid the possibilities of no change between regions grown at successive tolerance levels due to the use of too small a tolerance value, and negligible increment in terms of the general intensity of the region. A possible solution to avoid searching for the appropriate growth tolerance level is to make use of some of the characteristics of the HVS.

The multitolerance region-growing algorithm focuses mainly on the properties of the image to be processed without consideration of the observer's characteristics. The main purpose of an image processing method, such as segmentation, is to obtain a certain result to be presented to human observers for further analysis or for further computer analysis. In either case, the result should be consistent with a human observer's assessment. Therefore, effective segmentation should be achievable by including certain basic properties of the HVS.

The HVS is a nonlinear system with a large dynamic range and a band-pass filter behavior [122, 282]. The filtering property is characterized by the reciprocal of the threshold contrast C_T that is a function of both the frequency u and background luminance L . The smallest luminance difference that a human observer can detect when an object of a certain size appears with a certain background luminance level is defined as the JND, which can be quantified as [256]

$$JND = L C_T. \quad (5.55)$$

A typical variation of threshold contrast as a function of the background luminance, known as the Weber-Fechner relationship [323], is graphically depicted in Figure 5.21. The curve can be typically divided into two asymptotic re-

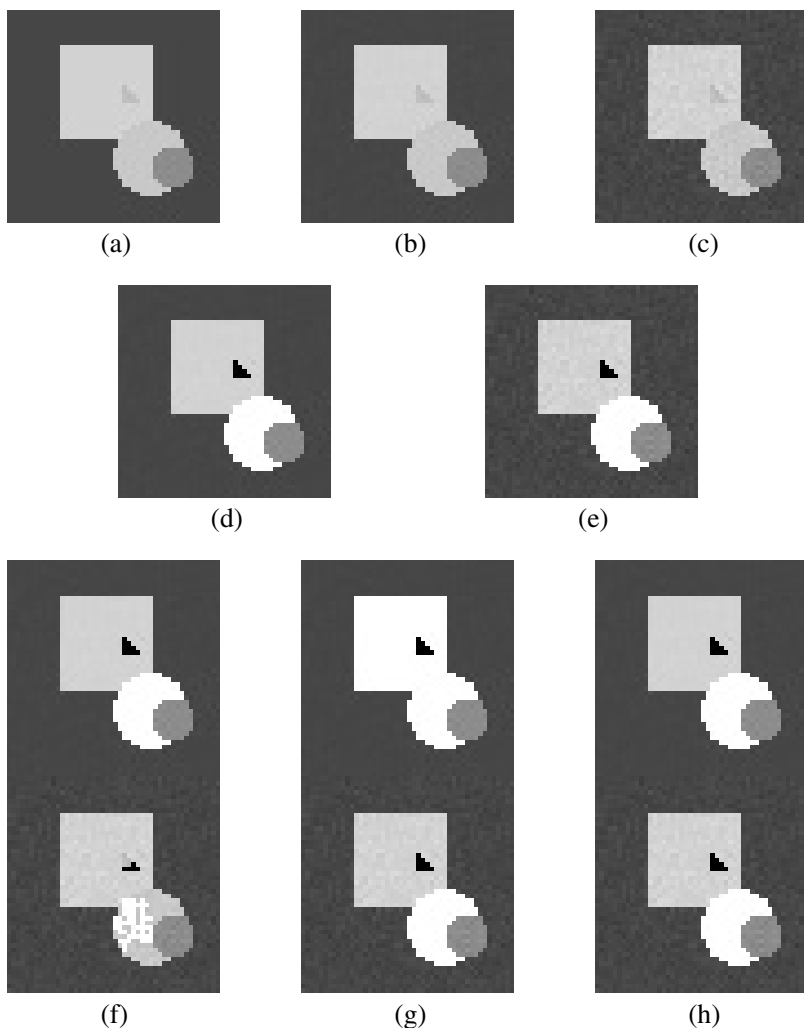


FIGURE 5.20

Demonstration of the multiplicative-tolerance region-growing algorithm. (a) An ideal test image. (b) Test image with white noise ($\text{SNR} = 30 \text{ dB}$). (c) Test image with white noise ($\text{SNR} = 20 \text{ dB}$). (d) Regions grown with fixed $\tau = 0.02$ for the image in (b). (e) Regions grown with fixed $\tau = 0.05$ for the image in (c). (f) Regions grown with fixed $\tau = 0.02$. (g) Regions grown with fixed $\tau = 0.05$. (h) Regions grown with the iterative algorithm with adaptive τ . The original composite images in (f) – (h) were obtained by combining the images in (b) at the top and (c) at the bottom. The detected regions are highlighted in black for the triangle and in white for the occluded circle in figures (d) – (h). Figure courtesy of L. Shen [320].

gions: the Rose – de Vries region where C_T decreases when the background luminance increases, with the relationship described as

$$C_T \propto \frac{1}{L^{0.5}}, \tag{5.56}$$

and the Weber region where C_T is independent of the background luminance, and the relationship obeys Weber’s law:

$$C_T = \frac{JND}{L} = C_0 = \text{constant}. \tag{5.57}$$




FIGURE 5.21

A typical threshold contrast curve, known as the Weber-Fechner relationship. Figure courtesy of L. Shen [320].

It is possible to determine the JND as a function of the background gray level from psychophysical experiments. In one such study conducted by Shen [320], a test was set up with various combinations of foreground and background levels using an image containing a series of square boxes, with the width ranging from 1 pixel to 16 pixels and a fixed space of 7 pixels in between the boxes. Also included was a series of groups of four, vertical, 64-pixel lines, with the width ranging from 1 pixel to 6 pixels and a spacing of

the same number of pixels between any two adjacent lines, and a fixed gap of 12 pixels in between the groups of lines. Figure 5.22 shows one such test image with the background gray level set to be 100 and the foreground at 200. Based upon the visibility or detection of up to the 2-pixel-wide square and line group on a monitor, a relationship between the JND and the background gray level was obtained as shown in Figure 5.23. In order to obtain a general JND relation, a large number of trials involving the participation of a large number of subjects is necessary, along with strict control of the experimental environment. Regardless, Shen [320] used the JND relationship illustrated in Figure 5.23 to develop a region-growing algorithm based upon the characteristics of the HVS.





FIGURE 5.22

Visual test image for determination of the JND as a function of background gray level (foreground is 200 and background is 100 in the displayed image). Figure courtesy of L. Shen [320].

The HVS-based region-growing algorithm starts with a 4-connected neighbor-pixel grouping based upon the JND relationships of adjacent pixel gray

**FIGURE 5.23**

A relationship between the JND and the background gray level based upon a psychophysical experiment. Figure courtesy of L. Shen [320].

levels. The JND condition is defined as

$$|p_1 - p_2| \leq \min\{JND(p_1), JND(p_2)\}, \quad (5.58)$$

where p_1 and p_2 are two connected pixels. This step is followed by the removal of small regions (defined as regions having fewer than five pixels in the study of Shen [320]) by merging with a connected region with the minimum mean gray-level difference. Then, merging of connected regions is performed if any of two neighboring regions meet the JND condition, with p_1 and p_2 representing the regions' mean values. The procedure is iterated until no neighboring region satisfies the JND condition.

Figure 5.24 shows the results of region growing with the same test image as in the test of the tolerance-based region-growing algorithms in Figure 5.20 (f). The HVS-based region-growing algorithm has successfully segmented the regions at the two SNR levels. The method is not time-consuming because a JND table is used to determine the parameters required.




FIGURE 5.24

Results of the HVS-based region-growing algorithm with the same test image as in Figure 5.20 (f). Figure courtesy of L. Shen [320].

5.4.9 Application: Detection of calcifications by multitolerance region growing

Microcalcifications are the most important and sometimes the only mammographic sign in early, curable breast cancer [52, 324]. Due to their subtlety, detection and classification (as benign or malignant) are two major problems. Several researchers in the field of mammographic image analy-

sis [325, 326, 327, 328, 329, 330, 331, 332, 333] have focused attention on the detection of calcifications. Shen et al. [274] reported on a method to detect and classify mammographic calcifications based upon a multitolerance region-growing procedure, shape analysis, and neural networks. The flowchart of the system is shown in Figure 5.25. The calcification detection algorithm consists of three steps:

1. selection of seed pixels;
2. detection of potential calcification regions; and
3. confirmation of calcification regions.

Selection of seed pixels: One of the common characteristics of calcifications in mammograms is that they are relatively bright due to the higher X-ray attenuation coefficient (or density) of calcium as compared with other normal breast tissues. Hence, a simple criterion to select seed pixels to search for calcifications could be based on the median or the mean value of the mammogram. The scheme employed by Shen et al. [274] is as follows:

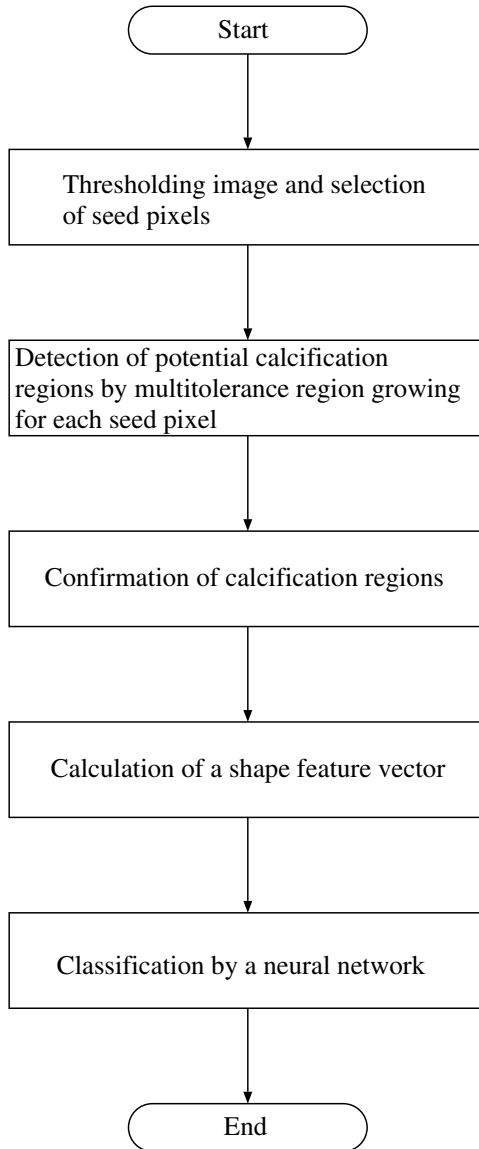
Every pixel with a value greater than the median gray level of the mammogram is identified as a potential seed pixel for the next two steps of calcification detection. The pixels identified as above are processed in sequence by selecting the highest intensity pixel remaining, in raster-scan order, as long the pixel has not been included in any of the regions already labeled as calcifications.

The selection scheme is simple and effective in most cases, but faces limitations in analyzing mammograms of dense breasts. Removal of the parts of the mammographic image outside the breast boundary (see Section 5.9) and the high-density area of the pectoral muscle (see Section 5.10) could be useful preprocessing steps.

Detection of potential calcification regions: Calcifications could appear in regions of varying density in mammograms. Calcifications present within dense masses or superimposed by dense tissues in the process of acquisition of mammograms could present low gray-level differences or contrast with respect to their local background. On the other hand, calcifications present against a background of fat or low-density tissue would possess higher differences and contrast. The iterative multitolerance region-growing method presented in Section 5.4.7 could be expected to perform well in this application, by adapting to variable conditions as above.

In the method proposed by Shen et al. [274] for the detection of calcifications, the algorithm starts a region-growing procedure with each seed pixel selected as above. Every 4-connected neighbor $f(m, n)$ of the pixels belonging to the region is checked for the following condition:

$$0.5(1 + \tau)(R_{\max} + R_{\min}) \geq f(m, n) \geq 0.5(1 - \tau)(R_{\max} + R_{\min}), \quad (5.59)$$

**FIGURE 5.25**

Flowchart of a method for the detection and classification of mammographic calcifications. Figure courtesy of L. Shen [320].

where R_{\max} and R_{\min} are the current maximum and minimum pixel values of the region being grown, and τ is the growth tolerance. The fractional tolerance value τ for region growing is increased from 0.01 to 0.40 with a step size determined as the inverse of the seed-pixel's gray level. A feature vector including compactness, defined as $c = 1 - 4\pi \frac{\text{area}}{\text{perimeter}^2}$ (see Section 6.2.1 for details), the (x, y) coordinates of the centroid, and the size or area in number of pixels, is calculated for the region obtained at each tolerance level. The normalized distance between the feature vectors for successive tolerance levels is computed, as given by Equation 5.54. The feature set with the minimum distance is selected as the final set, and the corresponding region considered to be a potential calcification region.

Confirmation of calcification regions: Each potential calcification region detected is treated as a calcification region only if the size S in pixels and contrast C , computed as in Equation 2.7, of the region at the final level meet the following conditions:

$$5 < S < 2,500 \quad (5.60)$$

and

$$C > 0.20. \quad (5.61)$$

The upper limit on the area corresponds to about 6.25 mm^2 with a pixel resolution of $50 \mu\text{m}$. The background region required to compute C is formed by using pixels circumscribing the region contour to a thickness of 3 pixels. The contrast threshold of 0.2 was selected based upon another study on segmentation and analysis of calcifications [334].

Examples: Two examples of the detection of calcifications by the method described above are presented in Figures 5.26 and 5.27. Figure 5.26 (a) and Figure 5.27 (a) show two sections of mammograms of size 512×512 pixels with benign calcifications and malignant calcifications, respectively. Figure 5.26 (b) and Figure 5.27 (b) show the same mammogram sections with the contours of the calcification regions extracted by the algorithm as described above.

Sections of size $1,024 \times 768$, 768×512 , 512×768 , and 512×768 pixels of four typical mammograms from complete images of up to $2,560 \times 4,096$ pixels with biopsy-proven calcifications were used in the study by Shen et al. [274]. Two of the sections had a total of 58 benign calcifications whereas the other two contained 241 ± 10 malignant calcifications. Based upon visual inspection by a radiologist, the detection rates of the multitolerance region-growing algorithm were 81% with 0 false calcifications and $85 \pm 3\%$ with 29 false calcifications for the benign and malignant mammograms, respectively.

Bankman et al. [335] compared their hill-climbing segmentation algorithm for detecting microcalcifications with the multitolerance region-growing algorithm described above. Their results showed that the two algorithms have similar discrimination powers based on an ROC analysis conducted with six mammograms (containing 15 clusters with a total of 124 calcifications). Bankman

et al. stated that “The multitolerance algorithm provides a good solution to avoid the use of statistical models, local statistic estimators, and the manual selection of thresholds. However, the cost is multiple segmentations of the same structure and computation of features during the segmentation of each structure. . . . The segmented regions were comparable . . . in many cases.”


Details on further analysis of the calcifications detected as above are provided in Sections 6.6 and 12.7. See Section 5.4.10 for another method for the detection of calcifications.

5.4.10 Application: Detection of calcifications by linear prediction error


The simple seed selection method used by Shen et al. [274] and described in Section 5.4.9 encounters limitations in the case of calcifications present in or superimposed by dense breast tissue. Serrano et al. [336] proposed a method to detect seed pixels for region growing based upon the error of linear prediction. The 2D linear prediction error computation method proposed by Kuduvalli and Rangayyan [174, 337, 338] was used to compute the prediction error directly from the image data without the need to compute the prediction coefficients; see Section 11.8.1 for details.

The method proposed by Serrano et al. [336] starts with a prefiltering step. Considering the small size of microcalcifications, a lowpass filter with a wide kernel could be expected to remove them from the image while conserving a background of high density in the image. Conversely, a highpass filter may be used to detect microcalcifications. Serrano et al. employed a highpass filter specified as $h(m, n) = 1 - g(m, n)$, where $g(m, n)$ was a lowpass Gaussian function with a variance of 2.75 pixels; a filter kernel size of 21×21 pixels was chosen. In the next step, a 2D linear prediction error filter [174, 337, 338] was applied to the output of the highpass filter. A pixel was selected as a seed for the multitolerance region-growing algorithm if its prediction error was greater than an experimentally determined threshold. This was based upon the observation that a microcalcification can be seen as a point of nonstationarity in an approximately homogeneous region or neighborhood in a mammogram; such a pixel cannot be predicted well by the linear predictor, and hence leads to a high error.

The detection algorithm was tested with three mammograms containing a total of 428 microcalcifications of different nature and diagnosis. The results obtained with the algorithm were examined by a radiologist who determined the accuracy of the detection. Figure 5.28 shows a segment of a mammogram with several calcifications, the seed pixels identified by thresholding the prediction error, and the regions obtained by application of the multitolerance region-growing algorithm to the seed pixels. In comparison with the algorithm of Shen et al. [274], the detection accuracy of the method of Serrano et al. was higher with a smaller number of false detections for the images tested,




(a)




(b)

FIGURE 5.26

Mammogram section with benign calcifications. (a) Original image. (b) Image with the contours of the calcification regions detected. The section shown is of size 512×512 pixels (approximately $2.25 \text{ cm} \times 2.25 \text{ cm}$), out of the full matrix of $1,536 \times 4,096$ pixels of the complete mammogram. Reproduced with permission from L. Shen, R.M. Rangayyan, and J.E.L. Desautels, "Detection and classification of mammographic calcifications", *International Journal of Pattern Recognition and Artificial Intelligence*, 7(6): 1403–1416, 1993. © World Scientific Publishing Co.



(a)



(b)

FIGURE 5.27

Mammogram section with malignant calcifications. (a) Original image. (b) Image with the contours of the calcification regions detected. The section shown is of size 512×512 pixels (approximately $2.25 \text{ cm} \times 2.25 \text{ cm}$), out of the full matrix of $1,792 \times 4,096$ pixels of the complete mammogram. Reproduced with permission from L. Shen, R.M. Rangayyan, and J.E.L. Desautels, "Detection and classification of mammographic calcifications", *International Journal of Pattern Recognition and Artificial Intelligence*, 7(6): 1403–1416, 1993. © World Scientific Publishing Co.

although the detection capability was diminished; that is, more calcifications were missed by the prediction-error method.




FIGURE 5.28

(a) Mammogram section with malignant calcifications; 234×137 pixels with a resolution of $160 \mu\text{m}$. (b) Seed pixels detected by thresholding the prediction error (marked in black). (c) Image with the contours of the calcification regions detected by region growing from the seed pixels in (b). Reproduced with permission from C. Serrano, J.D. Trujillo, B. Acha, and R.M. Rangayyan, "Use of 2D linear prediction error to detect microcalcifications in mammograms", *CDROM Proceedings of the II Latin American Congress on Biomedical Engineering*, Havana, Cuba, 23–25 May 2001. © Cuban Society of Bioengineering.

5.5 Fuzzy-set-based Region Growing to Detect Breast Tumors

Although mammography is being used for breast cancer screening, the analysis of masses and tumors on mammograms is, at times, difficult because developing signs of cancer may be minimal or masked by superimposed tis-

sues. Additional diagnostic procedures may be recommended when the original mammogram is equivocal.

Computer-aided image analysis techniques have the potential to improve the diagnostic accuracy of mammography and reduce the use of adjunctive procedures and morbidity, as well as health-care costs. Computer analysis can facilitate the enhancement, detection, characterization, and quantification of diagnostic features such as the shapes of calcifications and masses, growth of tumors into surrounding tissues, and the distortion caused by developing densities. The annotation of mammograms with objective measures may assist radiologists in diagnosis.

Computer-aided detection of breast masses is a challenging problem requiring sophisticated techniques due to the low contrast and poor definition of their boundaries. Classical segmentation techniques attempt to define precisely the ROI, such as a calcification or a mass. Shen et al. [274] proposed thresholding and multitolerance region growing methods for the detection of potential calcification regions and extraction of their contours; see Sections 5.4.7 and 5.4.9. Karssemeijer [339], Laine et al. [340], and Miller and Ramsey [341] proposed methods for tumor detection based on scale-space analysis. Zhang et al. [342] proposed an automated detection method for the initial identification of spiculated lesions based on an analysis of mammographic texture patterns. Matsubara et al. [343] described an algorithm based on an adaptive thresholding technique for mass detection. Kupinski and Giger [344] presented two methods for segmenting lesions in digital mammograms: a radial-gradient-index-based algorithm that considers both the gray-level information and a geometric constraint, and a probabilistic approach. However, defining criteria to realize precisely the boundaries of masses in mammograms is difficult. The problem is compounded by the fact that most malignant tumors possess fuzzy boundaries with a slow and extended transition from a dense core region to the surrounding tissues. (For detailed reviews on the detection and analysis of breast masses, refer to Rangayyan et al. [163, 345] and Mudigonda et al. [165, 275]. See Sections 6.7, 7.9, 12.11, and 12.12 for related discussions.)

An alternative to address the problem of detecting breast masses is to represent tumor or mass regions by fuzzy sets [307]. The most popular algorithm that uses the fuzzy-set approach is the fuzzy C-means algorithm [346, 347, 348]. The fuzzy C-means algorithm uses iterative optimization of an objective function based on weighted similarity measures between the pixels in the image and each cluster center. The segmentation method of Chen and Lee [348] uses fuzzy C-means as a preprocessing step in a Bayesian learning paradigm realized via the expectation-maximization algorithm for edge detection and segmentation of calcifications and masses in mammograms. However, their final result is based on classical segmentation to produce crisp boundaries. Sameti and Ward [349] proposed a lesion segmentation algorithm using fuzzy sets to partition a given mammogram. Their method divides a mammogram into two crisp regions according to a fuzzy membership function and an iterative optimization procedure to minimize an objective function. If more

than two regions are required, the algorithm can be applied to each region obtained in the preceding step using the same procedure. The authors presented results of application of the method to mammograms with four levels of segmentation.

Guliatto et al. [276] proposed two segmentation methods that incorporate fuzzy concepts. The first method determines the boundary of a tumor or mass by region growing after a preprocessing step based on fuzzy sets to enhance the ROI. The second segmentation method is a fuzzy region-growing method that takes into account the uncertainty present around the boundaries of tumors. These methods are described in the following sections.

5.5.1 Preprocessing based upon fuzzy sets

A mass or tumor typically appears on a mammogram as a relatively dense region, whose properties could be characterized using local density, gradient, texture, and other measures. A set of such local properties could be used to define a feature vector of a mass ROI and/or a pixel belonging to the ROI. Given a feature vector, a pixel whose properties are similar to those represented by the feature vector of the mass could be assigned a high intensity. If the properties do not match, the pixel intensity could be made low. At the end of such a process, the pixels in and around the ROI will be displayed according to their degree of similarity with respect to the features of the mass ROI.

A fuzzy set may be defined by assigning to each element considered from the universal set Ω a value representing its grade of membership in the fuzzy set [350, 351]. The grade corresponds to the degree with which the element is similar to or compatible with the concept represented by the fuzzy set. Let $\Gamma : \Omega \rightarrow L$ be a membership function that maps Ω into L , where L denotes any set that is at least partially ordered. The most commonly used range of values for membership functions is the unit real interval $[0, 1]$. Crisp sets can be seen as a particular case of fuzzy sets where $\Gamma : \Omega \rightarrow \{0, 1\}$; that is, the range includes only the discrete values 0 and 1.

The enhancement, and subsequent detection, of an ROI may be achieved by defining an appropriate membership function that evaluates the similarity between the properties of the pixel being considered and those of the ROI itself, given by the feature vector. In this procedure, the original image is mapped to a fuzzy set according to the membership function, which:

- assigns a membership degree equal to 1 to those pixels that possess the same properties as the mass ROI;
- represents the degree of similarity between the features of the mass ROI and those of the pixel being considered;
- exhibits symmetry with respect to the difference between the features of the ROI and those of the pixel being considered; and

- decreases monotonically from 1 to 0.

Guliato et al. [276] considered the mean intensity of a seed region, identified by the user, as the ROI feature. A membership function with the characteristics cited above, illustrated in Figure 5.29, is given by the function

$$\Gamma(p) = \frac{1}{1 + \beta |\mathbf{A} - \mathbf{B}|}, \quad (5.62)$$

where p is the pixel being processed, \mathbf{A} is the feature vector of the mass (gray level), \mathbf{B} is the feature vector of the pixel being analyzed, and β defines the opening of the membership function. For large β the opening is narrow and the function's behavior is strict; for small β the opening is wide, and the function presents a more permissive behavior.




FIGURE 5.29

Fuzzy membership function for preprocessing. Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, “Segmentation of breast tumors in mammograms using fuzzy sets”, *Journal of Electronic Imaging*, 12(3): 369 – 378, 2003. © SPIE and IS&T.

The fuzzy set obtained by the method described above represents pixels whose properties are close to those of the mass with a high membership degree; the opposite case results in a low membership degree. The membership degree may be used as a scale factor to obtain gray levels and display the result as an image. The contrast of the ROI in the resulting image depends upon the parameter β .

Figures 5.30 (a) and (b) show a 700×700 -pixel portion of a mammogram with a spiculated malignant tumor and the result of fuzzy-set-based preprocessing with $\beta = 0.007$, respectively. It is seen from the image in Figure

5.30 (b) that the pixels in the tumor region (the bright area in the upper-left part of the image) have higher values than the pixels in other parts of the image, indicating a higher degree of similarity with respect to the ROI or seed region. The membership values decrease gradually across the boundary of the tumor, as expected, due to the malignant nature of the tumor in this particular case. Note, however, that a few other spatially disconnected regions on the right-hand side of the image also have high values; these regions can be eliminated by further processing, as described in Section 5.5.2.




Figure 5.30 (a)

5.5.2 Fuzzy segmentation based upon region growing

Region growing is an image segmentation technique that groups pixels or subregions into larger regions according to a similarity criterion. Statistical measures provide good tools for defining homogeneous regions. The success of image segmentation is directly associated with the choice of the measures and a suitable threshold. In particular, mean and standard deviation measures are often used as parameters to control region growing; however, these measures are influenced by extreme pixel values. As a consequence, the final shape of the region grown depends upon the strategy used to traverse the image.




Figure 5.30 (b)





Figure 5.30 (c)



(d)

FIGURE 5.30

(a) A 700×700 -pixel portion of a mammogram with a spiculated malignant tumor. Pixel size = $62.5 \mu\text{m}$. (b) Fuzzy-set-based ROI enhancement with $\beta = 0.007$. (c) Contour extracted (white line) by region growing with the result in (b). The black line represents the boundary drawn by a radiologist (shown for comparison). $\beta = 0.007$, threshold = 0.63. (d) Result of fuzzy region growing with the image in (a) with $\Delta\mu_{\text{max}} = 45$, $\Delta CV_{\text{max}} = 0.01$, $\beta = 0.07$. The contour drawn by the radiologist is superimposed for comparison. Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, "Segmentation of breast tumors in mammograms using fuzzy sets", *Journal of Electronic Imaging*, 12(3): 369 – 378, 2003. © SPIE and IS&T.

Furthermore, the algorithm could present unstable behavior; for example, different pixels with the same values that are rejected at an earlier stage may be accepted later on in the region-growing method. It is also possible that the stopping condition is not reached when the gray level in the image increases slowly in the same direction as that of the traversal strategy. Besides, traditional region-growing methods represent the ROI by a classical set, defining precisely the region's boundary. In such a case, the transition information is lost, and the segmentation task becomes a critical stage in the image analysis system. In order to address these concerns, Guliato et al. [276] presented two image segmentation methods: the first based on classical region growing with the fuzzy-set preprocessed image (described in the following paragraphs), and the second based on fuzzy region growing using statistical measures in homogeneity criteria, described in Section 5.5.3.

The pixel values in the fuzzy-set preprocessed image represent the membership degrees of pixels with respect to the ROI as defined by the seed region. To perform contour extraction, the region-growing algorithm needs a threshold value and a seed region that lies inside the ROI. The region-growing process starts with the seed region. Four-connected neighboring pixels that are above the threshold are labeled as zero, the neighbors of the pixels labeled as zero are inspected, and the procedure continued. If the connected pixel is less than the threshold, it is labeled as one, indicating a contour pixel, and its neighborhood is not processed. The recursive process continues until all spatially connected pixels fail the test for inclusion in the region. A post-processing step is included to remove isolated pixels and regions that lie within the outermost contour.

The algorithm is simple and easy to implement, and will always produce closed contours. The method was evaluated with a number of synthetic test images as well as medical images such as CT and nuclear medicine images, and produced good results even in the presence of high levels of noise [352].

Examples: Figure 5.31 shows the results of the method with a synthetic image for three representative combinations of parameters. The three results exhibit a good degree of similarity and illustrate the robustness of the method in the presence of noise.

Figure 5.30 (c) shows the contour extracted for the mammogram in part (a) of the same figure. Figure 5.32 (a) shows a part of a mammogram with a circumscribed benign mass; part (b) of the figure shows the corresponding enhanced image. Figure 5.32 (c) shows the contour obtained: the image is superimposed with the contour obtained by region growing in white; the contour in black is the boundary drawn independently by an experienced radiologist, shown for comparison.

Results of application to mammograms: Guliato et al. [276] tested their method with 47 mammograms including 25 malignant tumors and 22 benign masses, and observed good agreement between the contours given by the method and those drawn independently by a radiologist. The seed region and threshold value were selected manually for each case; the threshold


**FIGURE 5.31**

Illustration of the effects of seed pixel and threshold selection on fuzzy-set preprocessing and region growing. (a) Original image (128×128 pixels) with additive Gaussian noise, with $\sigma = 12$ and $\text{SNR} = 2.66$. Results with (b) seed pixel $(60, 60)$ and threshold = 0.82; (c) seed pixel $(68, 60)$ and threshold = 0.85; (d) seed pixel $(68, 80)$ and threshold = 0.85. Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, "Segmentation of breast tumors in mammograms using fuzzy sets", *Journal of Electronic Imaging*, 12(3): 369 – 378, 2003. © SPIE and IS&T.




Figure 5.32 (a)

values varied between 0.57 and 0.90 for the images used. The same value of the membership function parameter $\beta = 0.007$ was used to process all of the images in the study. It was observed that the result of segmentation depended upon the choice of the seed to start region growing and the threshold. Automatic selection of the seed pixel or region and the threshold is a difficult problem that was not addressed in the study. It was observed that the threshold could possibly be derived as a function of the statistics (such as the mean and standard deviation) of the fuzzy-set preprocessed image.

Measure of fuzziness: In order to compare the results obtained by segmentation with the contours of the masses drawn by the radiologist, Guliato et al. [276, 277] developed a method to aggregate the segmented region with the reference contour. The procedure can aggregate not only two contours but also a contour with a fuzzy region, and hence is more general than classical intersection. The method uses a fuzzy fusion operator that generalizes classical intersection of sets, producing a fuzzy set that represents the agreement present among the two inputs; see Section 5.11 for details. The result of fusion was evaluated by a measure of fuzziness computed as

$$f(X) = \frac{\sum_{p \in X} [1 - |2\Gamma(p) - 1|]}{|X|}, \quad (5.63)$$




Figure 5.32 (b)





Figure 5.32 (c)



(d)

FIGURE 5.32

(a) A $1,024 \times 1,024$ -pixel portion of a mammogram with a circumscribed benign mass. Pixel size = $50 \mu\text{m}$. (b) Fuzzy-set-based ROI enhancement with $\beta = 0.007$. (c) Contour extracted (white line) by region growing with the result in (b). The black line represents the boundary drawn by a radiologist (shown for comparison). $\beta = 0.007$, threshold = 0.87. (d) Result of fuzzy region growing with the image in (a) with $\Delta\mu_{\text{max}} = 15$, $\Delta CV_{\text{max}} = 0.01$, $\beta = 0.07$. The contour drawn by the radiologist is superimposed for comparison. Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, "Segmentation of breast tumors in mammograms using fuzzy sets", *Journal of Electronic Imaging*, 12(3): 369 – 378, 2003. © SPIE and IS&T.

where X is the result of aggregation, and $\Gamma(p)$ is the degree of membership of the pixel p . The denominator in the expression above normalizes the measure with respect to the area of the result of fusion, resulting in a value in the range $[0, 1]$, with zero representing perfect agreement and unity indicating no intersection between the two inputs.

The values of the measure of fuzziness obtained for the 47 mammograms in the study were in the range $(0.13, 0.85)$, with the mean and standard deviation being 0.42 and 0.17, respectively. The measure of fuzziness was less than 0.5 for 34 out of the 47 cases. In most cases where the measure of fuzziness was greater than 0.5, the segmented region was smaller than, but contained within, the region indicated by the contour drawn by the radiologist. Regardless of the agreement in terms of the measure of fuzziness, it was argued that, for a spiculated lesion, there is no definite number of spicules that characterizes the lesion as malignant. The method captured the majority of the spicules in the cases analyzed, providing sufficient information for diagnosis (according to the analysis of the results performed by an expert radiologist).

Assessment of the results by pattern classification: In order to derive a parameter for discriminating between benign masses and malignant tumors, the following procedure was applied by Guliato et al. [276, 353]. A morphological erosion procedure with a square structuring element of size equal to 25% of the shorter dimension of the smallest rectangle containing the contour was applied to the contour, so that the core of the ROI was separated from the boundary. A parameter labeled as DCV was computed from the fuzzy-set preprocessed image, by taking the difference between the coefficient of variation (CV) of the entire ROI and that of the core of the ROI. A high value of DCV represents an inhomogeneous ROI, which could be indicative of a malignant tumor. The probability of malignancy based upon DCV was computed using the logistic regression method (see Section 12.5 for details); the result is illustrated in Figure 5.33. Several cut points were analyzed with the curve; the cut point of 0.02 resulted in all 22 benign masses and 16 out of the 25 malignant tumors being correctly classified, yielding a high specificity of 1.0 but a low sensitivity of 0.64.

5.5.3 Fuzzy region growing

Guliato et al. [276] also proposed a fuzzy region-growing algorithm to obtain mass regions in mammograms. In this method, an adaptive similarity criterion is used for region growing, with the mean and the standard deviation of the pixels in the region being grown as control parameters. The region is represented by a fuzzy set to preserve the transition information around boundary regions.

The algorithm starts with a seed region that lies inside the ROI and spreads by adding to the region 8-connected pixels that have similar properties. The homogeneity of the region is evaluated by calculating the mean (μ), standard deviation (σ), and the coefficient of variation $CV = \frac{\sigma}{\mu}$.




FIGURE 5.33

The probability of malignancy (vertical axis) derived from the parameter DCV (horizontal axis). Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, “Segmentation of breast tumors in mammograms using fuzzy sets”, *Journal of Electronic Imaging*, 12(3): 369 – 378, 2003. © SPIE and IS&T.

Let $\Delta\mu_{\max}$, ΔCV_{\max} , and β be the control parameters for region growing. $\Delta\mu_{\max}$ specifies the maximum allowed difference between the value of the pixel being analyzed and the mean of the subregion already grown. ΔCV_{\max} indicates the desired degree of homogeneity between two subregions; β defines the opening of the membership function. Let p be the next pixel to be analyzed and $I(p)$ be the value of p . The segmentation algorithm is executed in two steps:

1. $|I(p) - \mu| \leq \Delta\mu_{\max}$. If this condition is not satisfied, then the pixel is labeled as rejected. If the condition is satisfied, p is temporarily added to the subregion and μ_{new} and σ_{new} are calculated.
2. $|\frac{\sigma}{\mu} - \frac{\sigma_{\text{new}}}{\mu_{\text{new}}}| \leq \Delta CV_{\max}$. If the condition is satisfied, then p must definitely be added to the subregion and labeled as accepted, and μ and σ must be updated, that is, $\mu = \mu_{\text{new}}$ and $\sigma = \sigma_{\text{new}}$. If the condition is not satisfied, p is added to the subregion with the label *accepted with restriction*, and μ and σ are not modified.

The second step given above analyzes the distortion that the pixel p can produce if added to the subregion. At the beginning of the process, the region includes all the pixels in the seed region, and the standard deviation is set to zero. While the standard deviation of the region being grown is zero, a specific procedure is executed in the second step: $|\frac{\sigma}{\mu} - \frac{\sigma_{\text{new}}}{\mu_{\text{new}}}| \leq 2 \Delta CV_{\max}$. The parameter ΔCV_{\max} works as a filter that avoids the possibility that the

mean and standard deviation measures suffer undesirable modification during the region-growing process. Furthermore, the algorithm processes pixels in expanding concentric squares around the seed region, evaluating each pixel only once. These steps provide stability to the algorithm.

The membership function that maps the pixel values of the region resulting from the preceding procedure to the unit interval $[0, 1]$ could be based upon the mean of the region. Pixels that are close to the mean will have a high membership degree, and in the opposite case, a low membership degree. The desirable characteristics of the membership function are:

- the membership degree of the seed pixel or region must be 1;
- the membership degree of a pixel labeled as rejected must be 0;
- the membership function must be as independent of the seed pixel or region as possible;
- the membership degree must represent the proximity between a pixel labeled as accepted or accepted with restriction and the mean of the resulting region;
- the function must be symmetric with respect to the difference between the mean and the pixel value; and
- the function must decrease monotonically from 1 to 0.

The membership function Γ used by Guliato et al. [276] is illustrated in Figure 5.34, where $a = |\text{mean_seed_region} - \mu|$ and $b = \Delta\mu_{\max}$. The value of a pixel p is mapped to the fuzzy membership degree $\Gamma(p)$ as follows:

if $|I(p) - \mu| \leq a$ then $\Gamma(p) = 1$

else if $|I(p) - \mu| > b$ then $\Gamma(p) = 0$

$$\text{else } \Gamma(p) = \frac{1}{1 + \beta |I(p) - \mu|}.$$

The method was tested on several synthetic images with various levels of noise. Figure 5.35 illustrates three representative results of the method with a synthetic image and different seed pixels. The results do not differ significantly, indicating the low effect of noise on the method.

Results of application to mammograms: The fuzzy region for the malignant tumor shown in Figure 5.30 (a) is illustrated in part (d) of the same figure. Figure 5.32 (d) shows the fuzzy region obtained for the benign mass shown in part (a) of the same figure.

An interactive graphical interface was developed by Guliato et al. [353], using an object-oriented architecture with controller classes. Some of the features of the interface are fast and easy upgradability, portability, and threads

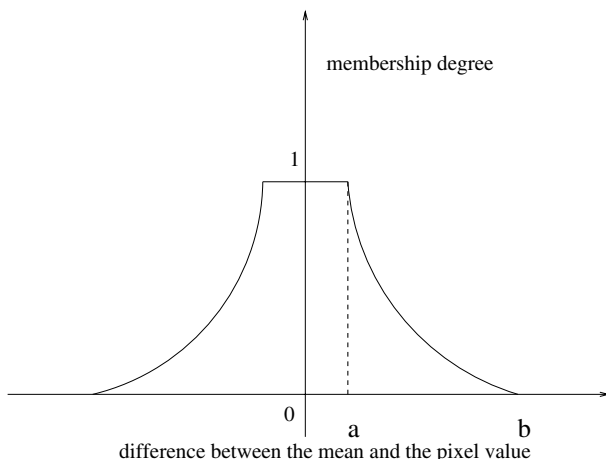


FIGURE 5.34

Fuzzy membership function for region growing, where $a = |\text{mean_seed_region} - \mu|$, and $b = \Delta\mu_{\max}$. Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, “Segmentation of breast tumors in mammograms using fuzzy sets”, *Journal of Electronic Imaging*, 12(3): 369 – 378, 2003. © SPIE and IS&T.

to support parallelism between tasks. The interface integrates procedures to detect contours using fuzzy preprocessing and region growing, extract fuzzy regions using fuzzy region growing, compute statistical parameters, and classify masses and tumors as benign or malignant. The interface also provides access to basic image processing procedures including zooming in or out, filters, histogram operations, the Bezier method to manipulate contours, and image format conversion.

Guliato et al. [276] applied the fuzzy region-growing method to 47 test images maintaining the same values of $\beta = 0.07$ and $\Delta CV_{\max} = 0.01$, and varying only the parameter $\Delta\mu_{\max}$. The values of the parameters were selected by comparing the results of segmentation with the contours drawn by a radiologist. The $\Delta\mu_{\max}$ parameter ranged from 5 to 48 for the 47 masses and tumors analyzed.

The fuzzy regions obtained for the 47 mammograms were compared objectively with the corresponding contours drawn by the radiologist, by computing the measure of fuzziness as in Equation 5.63. The values were distributed over the range (0.098, 0.82), with the mean and standard deviation being 0.46 and 0.19, respectively. The measure of fuzziness was smaller than 0.5 in 27 of the 47 cases analyzed. Regardless of this measure of agreement, it was found that the fuzzy regions segmented contained adequate information to facilitate discrimination between benign masses and malignant tumors, as described next.


**FIGURE 5.35**

Illustration of the effects of seed pixel selection on fuzzy region growing. (a) Original image (128×128 pixels) with Gaussian noise, with $\sigma = 12$ and $\text{SNR} = 2.66$. Results with (b) seed pixel $(60, 60)$, $\Delta\mu_{\max} = 18$, $\Delta CV_{\max} = 0.007$, $\beta = 0.01$; (c) seed pixel $(68, 60)$, $\Delta\mu_{\max} = 18$, $\Delta CV_{\max} = 0.007$, $\beta = 0.01$; (d) seed pixel $(68, 80)$, $\Delta\mu_{\max} = 18$, $\Delta CV_{\max} = 0.007$, $\beta = 0.01$. Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, "Segmentation of breast tumors in mammograms using fuzzy sets", *Journal of Electronic Imaging*, 12(3): 369 – 378, 2003. © SPIE and IS&T.

Assessment of the results by pattern classification: In order to derive parameters for pattern classification, Guliato et al. [276] analyzed the characteristics of a *fuzzy ribbon*, defined as the connected region whose pixels possess membership degrees less than unity and separate the tumor core from the background, as illustrated in Figure 5.36. Shape factors of mass contours as well as measures of edge sharpness and texture have been proposed for the purpose of classification of breast masses [163, 165, 275, 345, 354]; see Sections 6.7, 7.9, 12.11, and 12.12 for related discussions. However, important information is lost in analysis based on crisply defined contours: the uncertainty present in and/or around the ROI is not considered. Guliato et al. evaluated the potential use of statistical measures of each segmented fuzzy region and of its fuzzy ribbon as tools to classify masses as benign or malignant. Observe that the fuzzy ribbon of the malignant tumor in Figure 5.36 (a) contains more pixels with low values than that of the benign mass in part (b) of the same figure. This is due to the fact that, in general, malignant tumors possess ill-defined boundaries, whereas benign masses are well-circumscribed. Based upon this observation, Guliato et al. computed the coefficient of variation CV_{fr} of the membership values of the pixels lying only within the fuzzy ribbon, and the ratio ν_{fr} of the number of pixels with membership degree less than 0.5 to the total number of pixels within the fuzzy ribbon. It was expected that the fuzzy ribbons of malignant tumors would possess higher CV_{fr} and ν_{fr} than those of benign masses.

In pattern classification experiments, discrimination between benign masses and malignant tumors with the parameter ν_{fr} had no statistical significance. The probability of malignancy curve based upon CV_{fr} , computed using the logistic regression method (see Section 12.5 for details), is illustrated in Figure 5.37. The cut point of 0.18 resulted in the correct classification of 20 out of 25 malignant tumors and 20 out of 22 benign masses processed, leading to a sensitivity of 0.8 and a specificity of 0.9.

The fuzzy segmentation techniques described above represent the ROI by fuzzy sets instead of crisp sets as in classical segmentation. The results of the fuzzy approach agree well with visual perception, especially at transitions around boundaries. The methods allow the postponement of the crisp decision to a higher level of image analysis. For further theoretical notions related to fuzzy segmentation and illustrations of application to medical images, see Udupa and Samarasekera [355] and Saha et al. [356, 357].

5.6 Detection of Objects of Known Geometry

Occasionally, images contain objects that may be represented in an analytical form, such as straight-line segments, circles, ellipses, and parabolas. For




Figure 5.36 (a)


example, the edge of the pectoral muscle appears as an almost-straight line in MLO mammograms; benign calcifications and masses appear as almost-circular or oval objects in mammograms; several types of cells in pathology specimens have circular or elliptic boundaries, and some may have nearly rectangular shapes; and parts of the boundaries of malignant breast tumors may be represented using parabolas. The detection, modeling, and characterization of objects as above may be facilitated by prior knowledge of their shapes. In this section, we shall explore methods to detect straight lines and circles using the Hough transform.

5.6.1 The Hough transform

Hough [358] proposed a method to detect straight lines in images based upon the representation of straight lines in the image (x, y) space using the slope-intercept equation

$$y = m x + c, \quad (5.64)$$

where m is the slope and c is the position where the line intercepts the y axis; see Figure 5.38. In the Hough domain or space, straight lines are characterized by the pair of parameters (m, c) ; the Hough space is also known as the parameter space. A disadvantage of this representation is that both m and c have unbounded ranges, which creates practical difficulties in the computa-



(b)

FIGURE 5.36

The fuzzy ribbons of (a) the malignant tumor in Figure 5.30 (a) and (b) the benign mass in Figure 5.32 (a). Reproduced with permission from D. Gu-liato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, "Segmentation of breast tumors in mammograms using fuzzy sets", *Journal of Electronic Imaging*, 12(3): 369 – 378, 2003. © SPIE and IS&T.




FIGURE 5.37

The probability of malignancy (vertical axis) derived from the parameter CV_{fr} (horizontal axis). Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, “Segmentation of breast tumors in mammograms using fuzzy sets”, *Journal of Electronic Imaging*, 12(3): 369 – 378, 2003. © SPIE and IS&T.

tional representation of the (m, c) space. In order to overcome this limitation, Duda and Hart [359] proposed the representation of straight lines using the normal parameters (ρ, θ) as

$$\rho = x \cos \theta + y \sin \theta; \quad (5.65)$$

see Figure 5.38. This representation has the advantage that θ is limited to the range $[0, \pi]$ (or $[0, 2\pi)$) and ρ is limited by the size of the given image. The origin may be chosen to be at the center of the given image or at any other convenient point; the limits of the parameters (ρ, θ) are affected by the choice of the origin.

A procedure for the detection of straight lines using this parametric representation is described in the next section. The Hough transform may be extended for the detection of any curve that may be represented in a parametric form; see Rangayyan and Krishnan [360] for an application of the Hough transform for the identification of linear, sinusoidal, and hyperbolic frequency-modulated components of signals in the time-frequency plane.

5.6.2 Detection of straight lines

Suppose we are given a digital image that contains a straight line. Let the pixels along the line be represented as $\{x(n), y(n)\}$, $n = 0, 1, 2, \dots, N - 1$, where N is the number of pixels along the line. It is assumed that the image has been binarized, such that the pixels that belong to the line have the value




FIGURE 5.38

Parametric representation of a straight line in three coordinate systems: (x, y) , (m, c) , and (ρ, θ) .

1, and all other pixels have the value 0. It is advantageous if the line is one-pixel thick; otherwise, several lines could exist within a thick line.

If the normal parameters of the line are (ρ_0, θ_0) , all pixels along the line satisfy the relationship

$$\rho_0 = x(n) \cos \theta_0 + y(n) \sin \theta_0. \quad (5.66)$$

For a given pixel $\{x(n), y(n)\}$, this represents a sinusoidal curve in the (ρ, θ) parameter space; it follows that the curves for all the N pixels intersect at the point (ρ_0, θ_0) .

The following properties of the above representation follow [359]:

- A point in the (x, y) space corresponds to a sinusoidal curve in the (ρ, θ) parameter space.
- A point in the (ρ, θ) space corresponds to a straight line in the (x, y) space.
- Points lying on the same straight line in the (x, y) space correspond to curves through a common point in the parameter space.
- Points lying on the same curve in the parameter space correspond to lines through a common point in the (x, y) space.

Based upon the discussion above, a procedure to detect straight lines is as follows:

1. Discretize the (ρ, θ) parameter space into bins by quantizing ρ and θ as ρ_k , $k = 0, 1, 2, \dots, K - 1$, and θ_l , $l = 0, 1, 2, \dots, L - 1$; the bins are commonly referred to as accumulator cells. Suitable limits may be imposed on the ranges of the parameters (ρ, θ) .
2. For each point in the given image that has a value of 1, increment by 1 each accumulator cell in the (ρ, θ) space that satisfies the relationship $\rho = x(n) \cos \theta + y(n) \sin \theta$. Note that exact equality needs to be translated to a range of acceptance depending upon the discretization step size of the parameter space.
3. The coordinates of the point of intersection of all the curves in the parameter space provide the parameters of the line. This point will have the highest count in the parameter space.

The procedure given above assumes the existence of a single straight line in the image. If several lines exist, there will be the need to search for all possible points of intersection of several curves (or the local maxima). Note that the count in a given accumulator cell represents the number of pixels that lie on a straight line or several straight-line segments that have the corresponding (ρ, θ) parameters. A threshold may be applied to detect only lines that have a certain minimum length (number of pixels). All cells in the parameter space that have counts above the threshold may be taken to represent straight lines (or segments) with the corresponding (ρ, θ) values and numbers of pixels.

Examples: Figure 5.39 (a) shows an image with a single straight line, represented by the parameters $(\rho, \theta) = (20, 30^\circ)$. The limits of the x and y axes are ± 50 , with the origin at the center of the image. The Hough transform of the image is shown in part (b) of the figure in the (ρ, θ) parameter space. The maximum value in the parameter space occurs at $(\rho, \theta) = (20, 30^\circ)$.

An image containing two lines with $(\rho, \theta) = (20, 30^\circ)$ and $(-50, 60^\circ)$ is shown in Figure 5.40 (a), along with its Hough transform in part (b) of the figure. (The value of ρ was considered to be negative for normals to lines extending below the horizontal axis $x = 0$ in the image, with the origin at the center of the image; the range of θ was defined to be $[0, 180^\circ]$. It is also possible to maintain ρ to be positive, with the range of θ extended to $[0, 360^\circ]$.) The parameter space clearly demonstrates the expected sinusoidal patterns, as well as two peaks at the locations corresponding to the parameters of the two lines present in the image. Observe that the intensity of the point at the intersection of the sinusoidal curves for the second line (the lower of the two bright points in the parameter space) is less than that for the first line, reflecting its shorter length.

The application of the Hough transform to detect the pectoral muscle in mammograms is described in Section 5.10. See Section 8.6 for further discus-

sion on the Hough transform, and for a modification of the Hough transform by inclusion of the Radon transform.




FIGURE 5.39

(a) Image with a straight line with $(\rho, \theta) = (20, 30^\circ)$. The limits of the x and y axes are ± 50 , with the origin at the center of the image. (b) Hough transform parameter space for the image. The display intensity is $\log(1 + \text{accumulator cell value})$. The horizontal axis represents $\theta = [0, 180^\circ]$; the vertical axis represents $\rho = [-75, 75]$.


5.6.3 Detection of circles

For example, all points along the perimeter of a circle of radius c centered at $(x, y) = (a, b)$ satisfy the relationship

$$(x - a)^2 + (y - b)^2 = c^2. \quad (5.67)$$

Any circle is represented by a single point in the 3D (a, b, c) parameter space. The points along the perimeter of a circle in the (x, y) plane describe a right-circular cone in the (a, b, c) parameter space. The algorithm for the detection of straight lines (described in Section 5.6.2) may be easily extended for the detection of circles using this representation.

Example: A circle of radius 10 pixels, centered at $(x, y) = (15, 15)$ in a 30×30 image, is shown in Figure 5.41. The Hough parameter (a, b, c) space

**FIGURE 5.40**

(a) Image with two straight lines with $(\rho, \theta) = (20, 30^\circ)$ and $(-50, 60^\circ)$. The limits of the x and y axes are ± 50 , with the origin at the center of the image. (b) Hough transform parameter space for the image. The display intensity is $\log(1 + \text{accumulator cell value})$. The horizontal axis represents $\theta = [0, 180^\circ]$; the vertical axis represents $\rho = [-75, 75]$.

of the circle is shown in Figure 5.42. The parameter space demonstrates a clear peak at $(a, b, c) = (15, 15, 10)$ as expected.




FIGURE 5.41

A 30×30 image with a circle of radius 10 pixels, centered at $(x, y) = (15, 15)$.

An image derived from the scar-tissue collagen fiber image in Figure 1.5 (b) is shown in Figure 5.43 (a). The image was binarized with a threshold equal to 0.8 of its maximum value. In order to detect the edges of the objects in the image, an image was created with the value zero at all pixels having the value zero and also having all of their 4-connected pixels equal to zero in the binarized image. The same step was applied to all pixels with the value of one. All remaining pixels were assigned the value of one. Figure 5.43 (b) shows the result that depicts only the edges of the fibers, which are nearly circular in cross-section. Observe that some of the object boundaries are incomplete and not exactly circular. The Hough parameter (a, b, c) space of the image is shown in Figure 5.44 for circles of radius in the range 1 – 12 pixels. Several distinct peaks are visible in the planes for radius values of 4, 5, 6, and 7 pixels; the locations of the peaks give the coordinates of the centers of the circles that are present in the image and their radii. The parameter space could be further processed and thresholded to detect automatically the peaks present, which relate to the circles in the image. Prior knowledge of the range of possible radius values could assist in the process.

Figure 5.43 (c) shows the parameter space plane for a radius of 5 pixels, superimposed on a reversed version of the edge image in Figure 5.43 (b); the edges are shown in black. The composite image demonstrates clearly that the peaks in the parameter space coincide with the centers of nearly circular objects, with an estimated radius of 5 pixels; no peaks or high values are present within smaller or larger objects. A similar composite image is shown in part (d) of the figure for a radius of 7 pixels. Similar results are shown in Figures 5.45 and 5.46 for another TEM image of a normal ligament sample.




FIGURE 5.42

Hough parameter (a, b, c) space of the circle image in Figure 5.41. Each image is of size 30×30 pixels, and represents the range of the (a, b) coordinates of the center of a potential circle. The series of images represents the various planes of the (a, b, c) parameter space with $c = 1, 2, \dots, 36$, going left to right and top to bottom, representing the radius of potential circles. The intensity of the parameter space values has been enhanced with the log operation. The maximum value in the Hough space is located at the center of the plane for $c = 10$.

Observe that the Hough transform works well even when the given image has incomplete and slightly distorted versions of the pattern being represented.

Frank et al. [33] performed an analysis of the diameter distribution of collagen fibers in rabbit ligaments. Scar-tissue samples from injured and healing ligaments were observed to have an almost-uniform distribution of fiber diameter in the range 60 – 70 *nm*, whereas normal samples were observed to have a wider range of diameter, with an average value of about 150 *nm*.


5.7 Methods for the Improvement of Contour or Region Estimates

It is often the case that the contour of an ROI provided by an image processing technique does not satisfy the user. The user may wish to impose conditions, such as smoothness, on the contour. It may also be desirable to have the result authenticated by an expert in the field of application. In such cases, the need may arise to modify the contour. A few techniques that may assist in this task are summarized below.


Polygonal and parabolic models: The contour on hand may be segmented into significant parts by selecting control points or by detecting its points of inflection (see Section 6.1.3). The contour in its entirety, or its segments, may then be approximated by parametric curves, such as a polygon [245, 361] (see Section 6.1.4) or parabolas (see Section 6.1.5). Minor artifacts, details, or errors in the contour are removed by the parametric models, to the extent permitted by the specified tolerance or error.

B-spline and Bezier curves: Bezier polynomials may be used to fit smooth curves between guiding points that are specified in an interactive manner [245]. This approach may be used to modify a part of a contour without affecting the rest of the contour. In the case where several guiding or critical points are available around an ROI and a smooth curve is desired to connect them, B-splines may be used to derive the interpolating points [245]. The guiding points may also be used for compact representation of the contour.


Active contour models or “snakes”: Kass et al. [362] proposed active contour models or “snakes” that allow an initial contour to reshape and mold itself to a desired ROI based upon constraints related to the derivatives of the contour, image gradient, and energy functionals. A snake is an energy-minimizing spline that is guided by external constraint forces, influenced by image-based forces that pull it toward features such as edges, and constrained by internal spline forces that impose piecewise smoothness. The initial contour may be provided by an image processing technique, or could be provided by the user in the form of a simple (polygonal, circular, or oval) contour within the ROI. The active contour model determines the “true” boundary of the

**FIGURE 5.43**


(a) TEM image showing collagen fibers in cross-section [a part of the image in Figure 1.5 (b)]. The image is of size 85×85 pixels. (b) Edges extracted from the image in (a). (c) Negative version of the image in (b), overlaid with 10 times the $c = 5$ plane of the Hough transform parameter space. (d) Same as in (c) but with the $c = 7$ plane. See also Figure 5.44.

**FIGURE 5.44**

Hough parameter (a, b, c) space of the image in Figure 5.43 (b). Each image is of size 85×85 pixels, and represents the range of the (a, b) coordinates of the center of a potential circle. The series of images represents the various planes of the (a, b, c) parameter space with $c = 1, 2, \dots, 12$, going left to right and top to bottom, representing the radius of potential circles. The intensity of the parameter space values has been enhanced with the log operation.

**FIGURE 5.45**

(a) TEM image showing collagen fibers in cross-section [a part of the image in Figure 1.5 (a)]. The image is of size 143×157 pixels. (b) Edges extracted from the image in (a). (c) Negative version of the image in (b), overlaid with 10 times the $c = 13$ plane of the Hough transform parameter space. (d) Same as in (c) but with the $c = 20$ plane. See also Figure 5.46.

**FIGURE 5.46**

Hough parameter (a, b, c) space of the image in Figure 5.45 (b). Each image is of size 143×157 pixels, and represents the range of the (a, b) coordinates of the center of a potential circle. The series of images represents the various planes of the (a, b, c) parameter space with $c = 1, 2, \dots, 20$, going left to right and top to bottom, representing the radius of potential circles. The intensity of the parameter space values has been enhanced with the log operation.

ROI that satisfies the conditions imposed. The use of active contour models to obtain the breast boundary in mammograms is described in Section 5.9.

The “live wire”: Falcão et al. [363, 364] argued that there will continue to be situations where automatic image segmentation methods fail, and proposed a user-steered segmentation paradigm labeled as the “live wire”. Their guiding principles were to provide effective control to the user over the segmentation process during execution, and to minimize the user’s time required in the process of segmentation. In the live-wire approach, the user initially specifies a point on the boundary of the ROI using a cursor. Then, as the user moves the cursor, an optimal path connecting the initial point to the current cursor position is computed and displayed in real time. Optimal paths are determined by computing a number of features and their associated costs to every boundary element, and finding the minimum-cost paths via dynamic programming. When the cursor moves close to the true boundary, the live wire snaps on to the boundary. If the result is satisfactory, the user deposits the cursor, at which point the location of the cursor becomes the starting point of a new live-wire segment. The entire boundary of the ROI is specified as a set of live-wire segments. The features used include the image intensity values on the inside and outside of the boundary, several gradient measures, and the distance from the boundary in the preceding slice (in the case of segmentation of 3D images). Falcão et al. indicated that the method could assist in fast and accurate segmentation of ROIs in 2D and 3D medical images after some training.

Fusion of multiple results of segmentation: The segmentation of regions such as masses and tumors in mammograms is complicated by several factors, such as poor contrast, superimposition of tissues, imaging geometry, and the invasive nature of cancer. In such cases, it may be desirable to apply several image processing techniques based upon different principles and image characteristics, and then to combine or fuse the multiple results obtained into a single region or contour. It could be expected that the result so obtained would be better than any of the individual results. A fuzzy fusion operator that can fuse a contour and a fuzzy region to produce another fuzzy region is described in Section 5.11.

5.8 Application: Detection of the Spinal Canal

In an application to analyze CT images of neuroblastoma [365, 366] (see Section 9.9), the spinal canal was observed to interfere with the segmentation of the tumor using the fuzzy connectivity algorithm [355]. In order to address this problem, a method was developed to detect the center of the spinal canal in each CT slice, grow the 3D region containing the spinal canal, and remove

the structure. The initializing seeds for the region-growing procedure were automatically obtained with the following procedure.

The outer region in the CT volume containing materials outside the patient, the skin, and peripheral fat was first segmented and removed [365, 366]. The CT volume was then thresholded at $+800$ HU to detect the high-density bone structures. All voxels not within 8 mm from the inner boundary of the peripheral fat layer were rejected. Regions were grown using each remaining voxel, and all of the resulting regions were merged to form the bone volume. The inclusion criteria were in terms of the CT values being within $+800 \pm 2\sigma$ HU, with $\sigma = 103$ HU being the standard deviation of bone, and spatial connectivity.

The resulting CT volume was cropped to limit the scope of further analysis, as follows. The width of the image was divided into three equal parts, and the outer thirds were rejected. The height of the image was divided into six equal parts, and the lower fourth and fifth parts were included in the cropped region. In the interslice direction, the first 13% of the slices were removed, and the subsequent 20 slices were included in the cropped volume. Figure 5.47 (b) and Figure 5.48 (b) show the effect of cropping on two CT slices.

The cropped, binarized bone volume was subjected to a 3D derivative operator to produce the edges of the bone structures. The vertebral column is not continuous, but made up of interlocking elements. As a result, the bone-edge map could be sparse. Figures 5.47 (c) and 5.48 (c) show the bone-edge maps of the binarized bone volume related to the corresponding CT images in parts (b) of the same figures.

The Hough transform for the detection of circles, as described in Section 5.6.1 and Equation 5.67, was applied to each slice of the bone-edge map. The radius in the Hough space was limited to the range $6 - 10$ mm. Because of the possibility of partial structures and edges in a given image, the global maximum in the Hough space may not relate to the inner circular edge of the spinal canal, as desired. In order to obtain the center and radius of the ROI, the CT values of bone marrow ($\mu = +142$ HU and $\sigma = 48$ HU) and the spinal canal ($\mu = +30$ HU and $\sigma = 8$ HU) were used as constraints. If the center of the circle corresponding to the Hough-space maximum was not within the specified HU range, the circle was rejected, and the next maximum in the Hough space was evaluated. This process was continued until a suitable circle was detected.

Figure 5.47 (d) shows the best-fitting circle detected, drawn on the original CT image. When the bone structure is clearly delineated, the best-fitting circle approximates well the spinal canal boundary. Figure 5.48 (d) shows four circles related to the maximum in the corresponding Hough space and the subsequent three peaks; the related Hough-space slices are shown in Figure 5.49. The best-fitting circle (which was not given by the global maximum in the Hough space) was obtained by applying the constraints defined above.

The centers of the circles detected as above were used as the seed voxels in a fuzzy connectivity algorithm to segment the spinal canal. The mean and

standard deviation required for this procedure were estimated using a $7 \times 7 \times 2$ neighborhood around each seed voxel. The spinal canal detected over all of the CT slices available for the case illustrated in Figures 5.47 and 5.48 is shown in Figure 5.50. The spinal canal volume was then removed from the CT volume, resulting in improved segmentation of the tumor volume.

5.9 Application: Detection of the Breast Boundary in Mammograms

Identification of the breast boundary is important in order to demarcate the breast region on a mammogram. The inclusion of this preliminary procedure in CAD systems can avoid useless processing time and data storage. By identifying the boundary of the breast, it is possible to remove any artifact present outside the breast area, such as patient markings (often high-intensity regions) and noise, which can affect the performance of image analysis and pattern recognition techniques. Identification and extraction of the effective breast region is also important in PACS and telemammography systems [367].

The profile of the breast has been used as additional information in different tasks in mammography. Bick et al. [368] and Byng et al. [369], for example, used the skin-air boundary information to perform density correction of peripheral breast tissue on digital mammograms, which is affected by the compression procedure applied during imaging. Chandrasekhar and Attikiouzel [370] discussed the importance of the skin-air boundary profile as a constraint in searching for the nipple location, which is often used as a reference point for registering mammograms taken at different times of the same subject. Other groups have used the breast boundary to perform registration between left and right mammograms in the process of detection of asymmetry [371, 372].

Most of the works presented in the literature to identify the boundary of the breast are based upon histogram analysis [367, 368, 369, 371, 372, 373], which may be critically dependent upon the threshold-selection process and the noise present in the image. Such techniques, as discussed by Bick et al. [374], may not be robust for a screening application. Ferrari et al. [279, 375] proposed active contour models, especially designed to be locally adaptive, for identification of the breast boundary in mammograms. The methods, including several preprocessing steps, are described in the following sections.





FIGURE 5.47

(a) A 512×512 CT image slice of a patient with neuroblastoma. Pixel width = 0.41 mm . The tumor boundary drawn by a radiologist is shown in black. (b) Cropped area for further processing. (c) Edge map of the binarized bone volume. (d) Best-fitting circle as determined by Hough-space analysis. The circle has a radius of 6.6 mm . Figures courtesy of the Alberta Children's Hospital and H.J. Deglinc [365, 366].

**FIGURE 5.48**

(a) A 512×512 CT image slice of a patient with neuroblastoma. Pixel width = 0.41 mm . (b) Cropped area for further processing. (c) Edge map of the binarized bone volume. (d) The four circles corresponding to the first four peaks in the Hough space. The radii of the circles range from 6.6 mm to 9.8 mm . See also Figure 5.49. Figures courtesy of the Alberta Children's Hospital and H.J. Deglint [365, 366].




FIGURE 5.49

(a) Edge map of the binarized bone volume related to the ROI of the CT image in Figure 5.48. Hough-space slices related to the detection of circles of radius (b) 15, (c) 17, and (d) 19 pixels, with the pixel width being 0.41 mm. Figures courtesy of H.J. Deglint [365, 366].




FIGURE 5.50

3D rendition of the spinal canal detected for the case related to the CT slices shown in Figures 5.47 and 5.48. The spinal canal is shown in a relatively dark shade of gray against the bone volume for reference.

5.9.1 Detection using the traditional active deformable contour model

In an initial study, Ferrari et al. [375] used the traditional active deformable contour model or snakes [362] for the detection of the breast boundary. The method, summarized in Figure 5.51, is composed of six main stages, as described in the following paragraphs.

Stage 1: The image contrast is enhanced by using a simple logarithmic operation [8]. A contrast-correction step using a simple logarithmic operation as

$$g(x, y) = \log[1 + f(x, y)] \quad (5.68)$$

is applied to the original image $f(x, y)$; $g(x, y)$ is the transformed image. This operation for dynamic range compression, although applied to the whole image, significantly enhances the contrast of the regions near the breast boundary in mammograms, which are characterized by low density and poor definition of details [368, 369]. The rationale behind the application of this procedure to the image is to determine an approximate breast contour as close as possible to the true breast boundary. The effect of this procedure can be seen by comparing the original and the enhanced images in Figures 5.52 (a) and 5.52 (b).

Stage 2: A binarization procedure using the Lloyd–Max quantization algorithm is applied to the image [118]; see Section 2.3.2 for details. Figure 5.52 (c) shows the binarized version of the image in part (b) of the same figure.

Stage 3: Spurious details generated by the binarization step are removed by using a morphological opening operator [8] with a circular structuring element with a diameter of 7 pixels. Figures 5.52 (c)–(d) show the result of the binarization procedure for the mammogram in Figure 5.52 (a), before and after the application of the morphological opening operator, respectively.

Stage 4: After the binarization procedure, an approximate contour C_{appr} of the breast is extracted by using the chain-code method [8]; see Section 6.1.2 for details. The starting point of C_{appr} is obtained by following the horizontal path that starts at the centroid of the image and is directed toward the chest wall until a background pixel is found. This procedure avoids selecting an initial boundary from artifacts or patient labels that may be present in the image. Four control points [see Figure 5.52 (e)] are automatically determined and used to limit the breast boundary. The points are defined as N1: the top-left corner pixel of the boundary loop; N2: the farthest point on the boundary from N3 (in terms of the Euclidean distance through the breast); N3: the lowest pixel on the left-hand edge of the boundary loop; and N4: the farthest point on the skin-air boundary loop from N1.

Stage 5: Pixels along lines of length 40 pixels (length = 0.8 cm at a sampling resolution of 200 μm) are identified at each point of the approximate

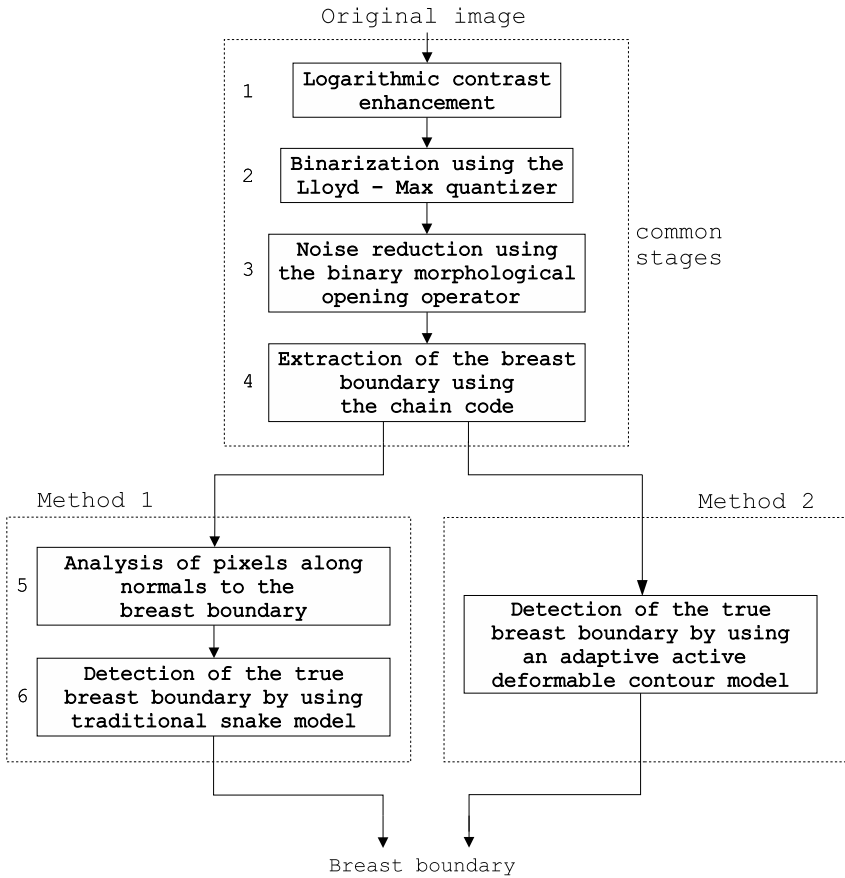


FIGURE 5.51

Flowchart of the procedures for identification of the skin-air boundary of the breast in mammograms. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Identification of the breast boundary in mammograms using active contour models”, *Medical and Biological Engineering and Computing*, 42: 201 – 208, 2004. © IFMBE.




Figure 5.52 (a)




Figure 5.52 (b)




Figure 5.52 (c)




Figure 5.52 (d)




Figure 5.52 (e)




Figure 5.52 (f)





Figure 5.52 (g)

skin-air boundary in the original image in the direction normal to the boundary [see Figure 5.52 (f)]. The gray-level histogram of the pixels along each normal line is computed, and the skin-air intersection is defined as the first pixel, while traversing along the normal line from inside the breast toward the outside, that has the gray level associated with the maximum value in the histogram, as illustrated in Figure 5.53. This procedure was designed in order to provide a close estimate to the true skin-air boundary [see Figure 5.52 (g)], and thereby reduce the chances of the active contour (used in the next stage) converging to a wrong contour.

Stage 6: The traditional snakes model is applied to define the true breast boundary. The contour determined in the previous stage is used as the input to a traditional parametric active contour or snakes model [362]. The contour is moved through the spatial domain of the image in order to minimize the energy functional

$$E = \int_0^1 \left[\frac{1}{2} \left\{ \alpha |v'(s)|^2 + \beta |v''(s)|^2 \right\} + E_{\text{ext}} \{v(s)\} \right] ds, \quad (5.69)$$


where α and β are weighting parameters that control, respectively, the tension and rigidity of the snake. The $v'(s)$ and $v''(s)$ values denote the first and second derivatives of $v(s)$ with respect to s , where $v(s)$ indicates the continuous representation of the contour, and s represents distance along the contour




(h)

FIGURE 5.52

Results of each stage of the method for identification of the breast boundary. (a) Image mdb042 from the Mini-MIAS database [376]. (b) Image after the logarithmic operation. (c)–(d) Binary image before and after applying the binary morphological opening operator. (e) Control points N1 to N4 (automatically determined) used to limit the breast boundary. (f) Normal lines computed from each pixel on the skin-air boundary. (g) Boundary after histogram-based analysis of the normal lines. (h) Final boundary. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Identification of the breast boundary in mammograms using active contour models”, *Medical and Biological Engineering and Computing*, 42: 201 – 208, 2004. © IFMBE.



(a)



(b)

FIGURE 5.53

(a) Profile of a sample normal line used to determine an approximate skin-air boundary. The symbol \times indicates the skin-air intersection determined in Stage 5 of the method. (b) Histogram computed from (a). Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Identification of the breast boundary in mammograms using active contour models”, *Medical and Biological Engineering and Computing*, 42: 201 – 208, 2004. © IFMBE.

(normalized to the range $[0, 1]$). The external energy function $E_{\text{ext}} \{v(s)\}$ is derived from the original image $f(x, y)$ as

$$E_{\text{ext}}(x, y) = - \|\nabla f(x, y)\|^2, \quad (5.70)$$

where ∇ is the gradient operator. The values $\alpha = 0.001$ and $\beta = 0.09$ were experimentally derived by Ferrari et al., based upon the approximate boundary obtained in the previous stage, the quality of the external force derived from the original image, and the final contours obtained.

Results of application to mammograms: Sixty-six images from the Mini-MIAS database [376] were used to assess the performance of the method. The results were subjectively analyzed by an expert radiologist. According to the opinion of the radiologist, the method accurately detected the breast boundary in 50 images, and reasonably well in 11 images. The final contour for the mammogram in Figure 5.52 (a) is given in part (h) of the same figure.

The method failed in five images because of distortions and artifacts present near the breast boundary (see Figure 5.54). Limitations of the method exist mainly in Stages 5 and 6. Although Stage 5 helps in obtaining a good breast contour, it is time-consuming, and may impose a limitation in practical applications. The traditional snakes model used in Stage 6 is not robust (the method is not locally adaptive) in the presence of noise and artifacts, and has a short range of edge capture.

In order to address these limitations, Ferrari et al. proposed an improved method, described in the following section, by replacing the traditional snakes algorithm with an adaptive active deformable contour model specifically designed for the application to detect breast boundaries. The algorithm includes a balloon force in an energy formulation that minimizes the influence of the initial contour on the convergence of the algorithm. In the energy formulation, the external energy is also designed to be locally adaptive.

5.9.2 Adaptive active deformable contour model

The improved method to identify the breast boundary is summarized in the flowchart in Figure 5.51. Stages 1 – 4 of the initial method described in the preceding section are used to find an approximate breast boundary. The approximate contour $V = v_1, v_2, \dots, v_N$, with an ordered collection of N points $v_i = (x_i, y_i)$, $i = 1, 2, \dots, N$, is obtained from Stage 4 of the previous method by sampling the approximate contour C_{appr} [see Figure 5.55 (a)], and used as the initial contour in the adaptive active deformable contour model (AADCM); see Figure 5.55 (b). Only 10% of the total number of points present in C_{appr} are used in the sampled contour.

In the AADCM, which combines several characteristics from other known active contour models [377, 378, 379], the contour is moved through the spatial domain of the image in order to minimize the following functional of energy:




Figure 5.54 (a)

$$E_{\text{total}} = \sum_{i=1}^N [\alpha E_{\text{internal}}(v_i) + \beta E_{\text{external}}(v_i)] , \quad (5.71)$$


where α and β are weighting parameters that control the internal and external energies E_{internal} and E_{external} , respectively, at each point v_i . The internal energy is composed of two terms:

$$E_{\text{internal}}(v_i) = a E_{\text{continuity}}(v_i) + b E_{\text{balloon}}(v_i) . \quad (5.72)$$

This energy component ensures a stable shape for the contour and constrains to keep constant the distance between the points in the contour. In the work of Ferrari et al., the weighting parameters a and b were initially set to unity ($a = b = 1$), because the initial contours present smooth shapes and are close to the true boundary in most cases. For each element (m, n) in a neighborhood of 7×7 pixels of v_i , the continuity term $e_{c(m,n)}(v_i)$ is computed as

$$e_{c(m,n)}(v_i) = \frac{1}{l(V)} |p_{(m,n)}(v_i) - \rho(v_{i-1} + v_{i+1})|^2 , \quad (5.73)$$

where $l(V) = \frac{1}{N} \sum_{i=1}^N |v_{i+1} - v_i|^2$ is a normalization factor that makes the continuity energy independent of the size, location, and orientation of V ;



(b)

FIGURE 5.54

Result of the segmentation algorithm showing wrong convergence of the breast contour to a region of high gradient value. (a) Breast boundary detected, superimposed on the original image mdb006 from the Mini-MIAS database. (b) Details of the breast contour attracted to the image identification marker, corresponding to the boxed region in (a). Compare with Figure 5.59. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Identification of the breast boundary in mammograms using active contour models”, *Medical and Biological Engineering and Computing*, 42: 201 – 208, 2004. © IFMBE.




Figure 5.55 (a)

$p_{(m,n)}(v_i)$ is the point in the image at the position (m, n) in the 7×7 neighborhood of v_i ; and $\rho = [2 \cos(\frac{2\pi}{N})]^{-1}$ is a constant factor to keep the location of the minimum energy lying on the circle connecting v_{i-1} and v_{i+1} , in the case of closed contours; see Figure 5.56.


The balloon force is used to force the expansion of the initial contour toward the breast boundary. In the work of Ferrari et al., the balloon force was made adaptive to the magnitude of the image gradient, causing the contour to expand faster in homogeneous regions and slower near the breast boundary. The balloon energy term $e_{b(m,n)}(v_i)$ is defined as

$$e_{b(m,n)}(v_i) = \mathbf{n}_i \bullet \{v_i - p_{(m,n)}(v_i)\}, \tag{5.74}$$

where \mathbf{n}_i is the outward unit normal vector of V at the point v_i , and the symbol \bullet indicates the dot product. \mathbf{n}_i is computed by rotating the vector $\mathbf{t}_i = \frac{v_i - v_{i-1}}{|v_i - v_{i-1}|} + \frac{v_{i+1} - v_i}{|v_{i+1} - v_i|}$, which is the tangent vector at the point v_i , by 90° ; see Figure 5.57.

The external energy is based upon the magnitude and direction of the image gradient, and is intended to attract the contour to the breast boundary. It is defined as

$$e_{e(m,n)}(v_i) = -\mathbf{n}_i \bullet \nabla f\{p_{(m,n)}(v_i)\}, \tag{5.75}$$



(b)

FIGURE 5.55

(a) Approximate breast contour obtained from Stage 4 of the method described in Section 5.9.1, for the image mdb042. (b) Sampled breast contour used as the input to the AADCM. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Identification of the breast boundary in mammograms using active contour models”, *Medical and Biological Engineering and Computing*, 42: 201 – 208, 2004. © IFMBE.




FIGURE 5.56

Characteristics of the continuity energy component in the adaptive active deformable contour model. Figure adapted with permission from B.T. Mackiewicz [377].




FIGURE 5.57

Characteristics of the balloon energy component in the adaptive active deformable contour model. Figure adapted with permission from B.T. Mackiewicz [377].

where $\nabla f\{p_{(m,n)}(v_i)\}$ is the image gradient vector at (m, n) in the 7×7 neighborhood of v_i ; see Figure 5.58. The direction of the image gradient is used to avoid the attraction of the contour by edges that may be located near the true breast boundary, such as identification marks and small artifacts; see Figure 5.59. In this situation, the gradient direction at the position (m, n) on an edge near the breast boundary and the direction of the unit normal of the contour will have opposite signs, which makes the functional of energy present a large value at (m, n) .

Minimization of the energy functionals: In order to allow comparison between the various energy components described above, in the work of Ferrari et al., each energy parameter was scaled to the range $[0, 1]$ as follows:

$$E_{\text{continuity}}(v_i) = \frac{e_{c(m,n)}(v_i) - e_{c \min}(v_i)}{e_{c \max}(v_i) - e_{c \min}(v_i)}; \quad (5.76)$$

$$E_{\text{balloon}}(v_i) = \frac{e_{b(m,n)}(v_i) - e_{b \min}(v_i)}{e_{b \max}(v_i) - e_{b \min}(v_i)} \left(1 - \frac{\|\nabla f(v_i)\|}{\|\nabla f\|_{\max}} \right); \quad (5.77)$$

$$E_{\text{external}}(v_i) = \frac{e_{e(m,n)}(v_i) - e_{e \min}(v_i)}{\max[e_{e \max}(v_i) - e_{e \min}(v_i), \|\nabla f\|_{\max}]}. \quad (5.78)$$

Here, e_{\min} and e_{\max} indicate the minimum and maximum of the corresponding energy component in the 7×7 neighborhood of v_i . $\|\nabla f\|_{\max}$ is the maximum gradient magnitude in the entire image.

Ferrari et al. used the Greedy algorithm, proposed by Williams and Shah [379], to perform minimization of the functional of energy in Equation 5.71.

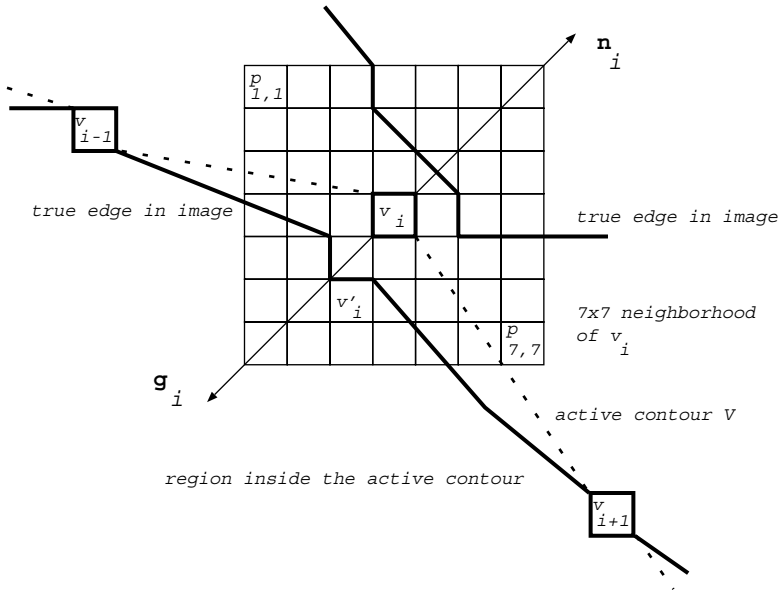


FIGURE 5.58

Characteristics of the external energy component in the adaptive active deformable contour model. Figure adapted with permission from B.T. Mackiewicz [377].




Figure 5.59 (a)

Although this algorithm has the drawback of not guaranteeing a global-minimum solution, it is faster than the other methods proposed in the literature such as dynamic programming, variational calculus, and finite elements. It also allows the insertion of hard constraints, such as curvature evaluation, as described below.

Convergence of the AADCM is achieved in two stages by smoothing the original image with two different Gaussian kernels defined with $\sigma_x = \sigma_y = 3$ pixels, and $\sigma_x = \sigma_y = 1.5$ pixels. At each stage, the iterative process is stopped when the total energy of the contour increases between consecutive iterations. This coarse-to-fine representation is expected to provide more stability to the contour.

In order to allow the deformable contour to adjust to corner regions, such as the upper-right limit of the breast boundary, a constraint was inserted at the end of each iteration by Ferrari et al., to relax the continuity term defined in Equation 5.73. The curvature value $C(v_i)$ at each point v_i of the contour was computed as

$$C(v_i) = 2 \sin \left(\frac{\theta}{2} \right) = \left\| \frac{\mathbf{u}_i}{\|\mathbf{u}_i\|} - \frac{\mathbf{u}_{i-1}}{\|\mathbf{u}_{i-1}\|} \right\|^2, \quad (5.79)$$

where $\mathbf{u}_i = (v_{i+1} - v_i)$ is the vector joining two neighboring contour elements and θ is the external angle between such vectors sharing a common contour



(b)

FIGURE 5.59

Application of the gradient direction information to avoid the attraction of the boundary to objects near the true boundary. (a) Breast boundary detected automatically, superimposed on the original image mdb006 from the Mini-MIAS database. (b) Details of the detected breast boundary close to the image identification marker, corresponding to the boxed region in the original image. Compare with Figure 5.54. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, "Identification of the breast boundary in mammograms using active contour models", *Medical and Biological Engineering and Computing*, 42: 201 – 208, 2004. © IFMBE.

element. This definition of curvature, as discussed by Williams and Shah [379], has three important advantages over other curvature measures: it requires only simple computation, gives coherent values, and depends solely on relative direction.

At each v_i , the weight values for the continuity term and the external energy are set, respectively, to zero ($a = 0$) and to twice the initial value ($\beta \leftarrow 2\beta$) if $[C(v_i) > C(v_{i-1})]$ and $[C(v_i) > C(v_{i+1})]$ and $[C(v_i) > T]$. The threshold value T was set equal to 0.25, which corresponds to an external angle of approximately 29° . According to Williams and Shah [379], this value of the threshold has been proven experimentally to be sufficiently large to differentiate between corners and curved lines. Figures 5.60 (b) and (c) illustrate an example without and with the curvature constraint to correct corner effects.





Figure 5.60 (a)

In the work of Ferrari et al., the weighting parameters α and β in Equation 5.71 were initialized to 0.2 and 1.0, respectively, for each contour element. This set of weights was determined experimentally by using a set of 20 images randomly selected from the Mini-MIAS database [376], not including any image in the test set used to evaluate the results. A larger weight was given to the gradient energy to favor contour deformation toward the breast boundary rather than smoothing due to the internal force.

**FIGURE 5.60**

Example of the constraint used in the active contour model to prevent smoothing effects at corners. (a) Original image; the box indicates the region of concern. (b) – (c) Details of the breast contour without and with the constraint for corner correction, respectively. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Identification of the breast boundary in mammograms using active contour models”, *Medical and Biological Engineering and Computing*, 42: 201 – 208, 2004. © IFMBE.

5.9.3 Results of application to mammograms

Ferrari et al. applied their methods to 84 images randomly chosen from the Mini-MIAS database [376]. All images were MLO views with $200\ \mu\text{m}$ sampling interval and 8-bit gray-level quantization. For reduction of processing time, all images were downsampled with a fixed sampling distance so that the original images corresponding to a matrix size of $1,024 \times 1,024$ pixels were transformed to 256×256 pixels. The results obtained with the downsampled images were mapped to the original mammograms for subsequent analysis and display. The results were evaluated in consultation with two radiologists experienced in mammography.

The test images were displayed on a computer monitor with a diagonal size of $47.5\ \text{cm}$ and dot pitch of $0.27\ \text{mm}$. By using the Gimp program [380], the contrast and brightness of each image were manually enhanced so that the breast contour could be easily visualized. The breast boundary was manually drawn under the supervision of a radiologist, and the results printed on paper by using a laser printer with $600\ \text{dpi}$ resolution. The zoom option of the Gimp program was used to aid in drawing the contours. The breast boundaries of all images were visually checked by a radiologist using the printed images (hardcopy) along with the displayed images (softcopy); the assessment was recorded for analysis.

The segmentation results related to the breast contours detected by image processing were evaluated based upon the number of false-positive (FP) and false-negative (FN) pixels identified and normalized with reference to the corresponding areas demarcated by the manually drawn contours. The reference area for the breast boundary was defined as the area of the breast image delimited by the hand-drawn breast boundary. The FP and FN average percentages and the corresponding standard deviation values obtained for the 84 images were $0.41 \pm 0.25\%$ and $0.58 \pm 0.67\%$, respectively. Thirty-three images presented both FP and FN percentages less than 0.5% ; 38 images presented FP and FN percentages between 0.5% and 1% ; the FP and FN percentages were greater than 1% for 13 images. The most common cause of FN pixels was related to missing the nipple region, as illustrated by the example in Figure 5.61 (c). By removing Stage 5 used to approximate the initial contour to the true breast boundary (see Figure 5.51), the average time for processing an image was reduced from $2.0\ \text{min}$ to $0.18\ \text{min}$. However, the number of images where the nipple was not identified increased. The AADCM performed successfully in cases where a small bending deformation of the contour was required to detect the nipple; see Figure 5.62.

The method for the detection of the breast boundary was used as a pre-processing step in the analysis of bilateral asymmetry by Ferrari et al. [381] (see Section 8.9). The method may also be used in other applications, such as image compression by using only the effective area of the breast, and image registration.




Figure 5.61 (a)





Figure 5.61 (b)



(c)

FIGURE 5.61

Results obtained for the image mdb003 from the Mini-MIAS database. (a) Original image. (b) Hand-drawn boundary superimposed on the histogram-equalized image. (c) Breast boundary detected, superimposed on the original image. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Identification of the breast boundary in mammograms using active contour models”, *Medical and Biological Engineering and Computing*, 42: 201 – 208, 2004. © IFMBE.




Figure 5.62 (a)





Figure 5.62 (b)



(c)

FIGURE 5.62

Results obtained for the image mdb114 from the Mini-MIAS database. (a) Original image. (b) Hand-drawn boundary superimposed on the histogram-equalized image. (c) Breast boundary detected, superimposed on the original image. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Identification of the breast boundary in mammograms using active contour models”, *Medical and Biological Engineering and Computing*, 42: 201 – 208, 2004. © IFMBE.

5.10 Application: Detection of the Pectoral Muscle in Mammograms

The pectoral muscle represents a predominant density region in most MLO views of mammograms, and can affect the results of image processing methods. Intensity-based methods, for example, can present poor performance when applied to differentiate dense structures such as the fibroglandular disc or small suspicious masses, because the pectoral muscle appears at approximately the same density as the dense tissues of interest in the image. The inclusion of the pectoral muscle in the image data being processed could also bias the detection procedures. Another important need to identify the pectoral muscle lies in the possibility that the local information of its edge, along with internal analysis of its region, could be used to identify the presence of abnormal axillary lymph nodes, which may be the only manifestation of occult breast carcinoma in some cases [54].

Karssemeijer [382] used the Hough transform and a set of threshold values applied to the accumulator cells in order to detect the pectoral muscle. Aylward et al. [383] used a gradient-magnitude ridge-traversal algorithm at small scale, and then resolved the resulting multiple edges via a voting scheme in order to segment the pectoral muscle. Ferrari et al. [278, 375] proposed a technique to detect the pectoral muscle based upon the Hough transform [8, 10], which was a modification of the method proposed by Karssemeijer [382]. However, the hypothesis of a straight line for the representation of the pectoral muscle is not always correct, and may impose limitations on subsequent stages of image analysis. Subsequently, Ferrari et al. [278] proposed another method based upon directional filtering using Gabor wavelets; this method overcomes the limitation of the straight-line representation mentioned above. The Hough-transform-based and Gabor-wavelet-based methods are described in the following sections.

5.10.1 Detection using the Hough transform

The initial method to identify the pectoral muscle proposed by Ferrari et al. [375], summarized in the flowchart in Figure 5.63, starts by automatically identifying an appropriate ROI containing the pectoral muscle, as shown in Figure 5.64. To begin with, an approximate breast contour delimiting the control points is obtained by using a method for the detection of the skin-air boundary [279, 375], described in Section 5.9. The six control points N1–N6 used to define the ROI are determined as follows, in order:

1. N1: the top-left corner pixel of the boundary loop;
2. N5: the lowest pixel on the left-hand edge of the boundary loop;

3. N3: the mid-point between N1 and N5;
4. N2: the farthest point on the boundary from N5 in terms of the Euclidean distance through the breast (if this point is not located on the upper edge of the mammogram, it is projected vertically to the upper edge);
5. N4: the point that completes a rectangle with N1, N2, and N3 (not necessarily on the boundary loop);
6. N6: the farthest point on the boundary loop from N1. In the case of the mammogram in Figure 5.64, the points N5 and N6 have coincided.

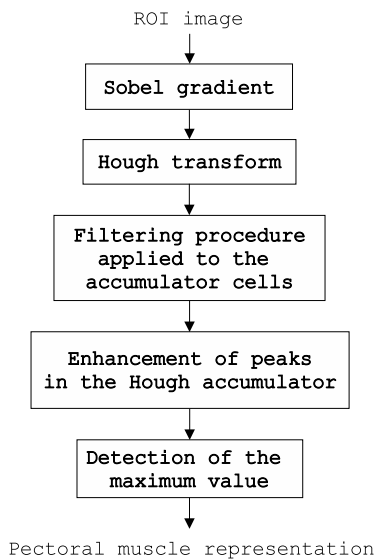
The ROI is defined by the rectangular region delimited by the points N1, N2, N3, and N4, as illustrated in Figure 5.64. Although, in some cases, this region may not include the total length of the pectoral muscle, the portion of the muscle present is adequate to define a straight line to represent its edge. By limiting the size of the ROI as described above, the bias that could be introduced by other linear structures that may be present in the fibroglandular disc is minimized. Differing from the method of Karssemeijer, Ferrari et al. [375] did not use any threshold value in order to reduce the number of unlikely pectoral lines. Instead, geometric and anatomical constraints were incorporated into the method, as follows:

1. The pectoral muscle is considered to be a straight line limited to an angle θ between 120° and 170° , with the angle computed as indicated in Figure 5.65. Mammograms of right breasts are flipped (mirrored) before performing pectoral muscle detection.
2. The pectoral line intercepts the line segment N1 – N2, as indicated in Figure 5.64.
3. The pectoral line is present, in partial or total length, in the ROI defined as the rectangular region delimited by the points N1, N2, N3, and N4, as illustrated in Figure 5.64.
4. The pectoral muscle appears on mammograms as a dense region with homogeneous gray-level values.

After selecting the ROI, a Gaussian filter with $\sigma_x = \sigma_y = 4$ pixels is used to smooth the ROI in order to remove the high-frequency noise in the image. The Hough transform is then applied to the Sobel gradient of the ROI [10] to detect the edge of the pectoral muscle. The representation of a straight line for the Hough transform computation is specified as

$$\rho = (x - x_0) \cos \theta + (y - y_0) \sin \theta, \quad (5.80)$$

where (x_0, y_0) is the origin of the coordinate system defined as the center of the image, and ρ and θ represent, respectively, the distance and angle between

**FIGURE 5.63**

Procedure for the identification of the pectoral muscle by using the Hough transform. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, "Automatic identification of the pectoral muscle in mammograms", *IEEE Transactions on Medical Imaging*, 23: 232 – 245, 2004. © IEEE.





Figure 5.64 (a)

(x_0, y_0) and the coordinates (x, y) of the pixel being analyzed, as illustrated in Figure 5.65. The Hough accumulator is quantized to 45 bins of 4° each by using the constraint $|\phi_{x,y} - \theta| < 2^\circ$, where $\phi_{x,y}$ is the orientation of the Sobel gradient at the pixel (x, y) . In the work of Ferrari et al. [278, 375], the accumulator cells were incremented using the magnitude of the gradient instead of unit increments; thus, pixels with a strong gradient have larger weights. Only values of θ in the range $[120^\circ, 170^\circ]$ were considered in the analysis, because the pectoral muscle of the mammogram was positioned on the left-hand side of the image before computing the Hough transform (see Figure 5.65).

After computing the accumulator cell values in the Hough space, a filtering procedure is applied to eliminate all lines (pairs of parameters ρ and θ) that are unlikely to represent the pectoral muscle. In this procedure, all lines intercepting the top of the image outside the N1 – N2 line segment (see Figure 5.64) or with slopes outside the range $[120^\circ, 170^\circ]$ are removed. (In the present notation, the x axis corresponds to 0° , and the chest wall is positioned on the left-hand side; see Figure 5.65.) Each remaining accumulator cell is also multiplied by the factor

$$\alpha = \frac{\mu}{\sigma^2} A \Big|_{\theta, \rho}, \quad (5.81)$$



(b)

FIGURE 5.64

(a) Image mdb042 from the Mini-MIAS database. (b) Approximate boundary of the breast along with the automatically determined control points N1 – N6 used to limit the ROI (rectangle marked) for the detection of the pectoral muscle. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Automatic identification of the pectoral muscle in mammograms”, *IEEE Transactions on Medical Imaging*, 23: 232 – 245, 2004. © IEEE.




FIGURE 5.65

Coordinate system used to compute the Hough transform. The pectoral muscle line detected is also shown. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Automatic identification of the pectoral muscle in mammograms”, *IEEE Transactions on Medical Imaging*, 23: 232 – 245, 2004. © IEEE.

where μ and σ^2 are, respectively, the mean and the variance of the gray-level values in the area A of the pectoral muscle (see Figure 5.65), defined by the straight line specified by the parameters θ and ρ . This procedure was applied in order to enhance the Hough transform peaks that define regions with the property stated in Item 4 of the list provided earlier in this section (page 482). The weight related to the area was designed to differentiate the true pectoral muscle from the pectoralis minor; the latter could present a higher contrast than the former in some cases, albeit enclosing a smaller area than the former.

Finally, the parameters ρ and θ of the accumulator cell with the maximum value are taken to represent the pectoral muscle line. Figure 5.66 shows the Hough accumulator cells at the different stages of the procedure described above for the mammogram in Figure 5.64 (a). The pectoral muscle line detected for the mammogram in Figure 5.64 (a) is shown in Figure 5.65.

5.10.2 Detection using Gabor wavelets

In an improved method to detect the pectoral muscle edge, summarized by the flowchart in Figure 5.67, Ferrari et al. [278] designed a bank of Gabor filters to enhance the directional, piecewise-linear structures that are present in an ROI containing the pectoral muscle. Figure 5.68 (a) illustrates a mammogram from the Mini-MIAS database; Figure 5.68 (b) shows the ROI — defined automatically as the rectangle formed by the chest wall as the left-hand edge, and a vertical line through the upper-most point on the skin-air boundary drawn along the entire height of the mammogram as the right-hand edge — to be used for the detection of the pectoral muscle. Differing from the Hough-transform-based method described in Section 5.10.1, the ROI here is defined to contain the entire pectoral muscle region, because the straight-line representation is not used. Decomposition of the ROI into components with different scale and orientation is performed by convolution of the ROI image with a bank of tunable Gabor filters. The magnitude and phase components of the filtered images are then combined and used as input to a post-processing stage, as described in the following paragraphs.

Gabor wavelets: A 2D Gabor function is a Gaussian modulated by a complex sinusoid, which can be specified by the frequency of the sinusoid W and the standard deviations σ_x and σ_y of the Gaussian envelope as [381, 384]

$$\psi(x, y) = \frac{1}{2\pi\sigma_x\sigma_y} \exp \left[-\frac{1}{2} \left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} \right) + j2\pi Wx \right]. \quad (5.82)$$

(See also Sections 8.4, 8.9, and 8.10.) Gabor wavelets are obtained by dilation and rotation of $\psi(x, y)$ as in Equation 5.82 by using the generating function


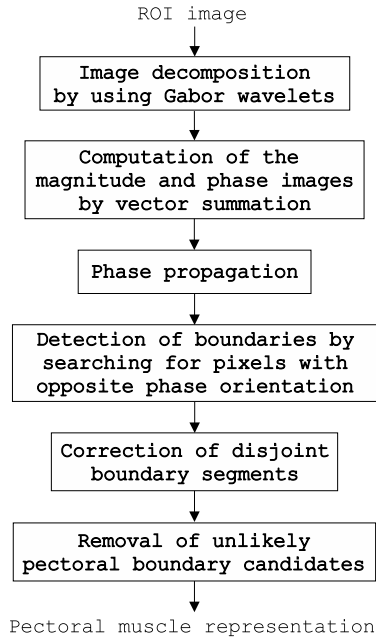


FIGURE 5.66

Hough accumulator cells obtained at three stages of the procedure to detect the pectoral muscle. The contrast of the images has been modified for improved visualization. (a) Accumulator cells obtained by using the constraint $|\phi_{x,y} - \theta| < 2^\circ$ and $120^\circ \leq \theta \leq 170^\circ$. (b) After removing the lines intercepting the top of the image outside the region defined by the control points N1 – N2 (see Figure 5.65). (c) After applying the multiplicative factor $\alpha = \frac{\mu}{\sigma_x} A \Big|_{\theta, \rho}$. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Automatic identification of the pectoral muscle in mammograms”, *IEEE Transactions on Medical Imaging*, 23: 232 – 245, 2004. © IEEE.

**FIGURE 5.67**

Flowchart of the procedure for the identification of the pectoral muscle by using Gabor wavelets. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, "Automatic identification of the pectoral muscle in mammograms", *IEEE Transactions on Medical Imaging*, 23: 232 – 245, 2004. © IEEE.




FIGURE 5.68

(a) Image mdb028 from the Mini-MIAS database. (b) The ROI used to search for the pectoral muscle region, defined by the chest wall and the upper limit of the skin-air boundary. The box drawn in (a) is not related to the ROI in (b). Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Automatic identification of the pectoral muscle in mammograms”, *IEEE Transactions on Medical Imaging*, 23: 232 – 245, 2004. © IEEE.

$$\psi_{m,n}(x, y) = a^{-m} \psi(x', y'), \quad a > 1, \quad m, n = \text{integers}, \quad (5.83)$$

$$x' = a^{-m} [(x - x_0) \cos \theta_n + (y - y_0) \sin \theta_n],$$

$$y' = a^{-m} [-(x - x_0) \sin \theta_n + (y - y_0) \cos \theta_n],$$

where (x_0, y_0) is the center of the filter in the spatial domain; $\theta_n = \frac{n\pi}{K}$, $n = 1, 2, \dots, K$; K is the number of orientations desired; and m and n indicate the scale and orientation, respectively. The Gabor filter used by Ferrari et al. [278] is expressed in the frequency domain as

$$\Psi(u, v) = \frac{1}{2\pi\sigma_u\sigma_v} \exp \left\{ -\frac{1}{2} \left[\frac{(u - W)^2}{\sigma_u^2} + \frac{v^2}{\sigma_v^2} \right] \right\}, \quad (5.84)$$

where $\sigma_u = \frac{1}{2\pi\sigma_x}$ and $\sigma_v = \frac{1}{2\pi\sigma_y}$. The design strategy used by Ferrari et al. was to project the filters so as to ensure that the half-peak magnitude supports of the filter responses in the frequency spectrum touch one another, as shown in Figure 5.69. By doing this, it can be ensured that the filters will capture the maximum information with minimum redundancy. In order for the designed bank of Gabor filters to be a family of admissible 2D Gabor wavelets [385], the filters $\psi(x, y)$ must satisfy the admissibility condition of finite energy [386], which implies that their Fourier transforms are pure bandpass functions having zero response at DC. This condition is achieved by setting the DC gain of each filter as $\Psi(0, 0) = 0$, which causes the filters not to respond to regions with constant intensity.

The following formulas provide the filter parameters σ_u and σ_v :

$$a = \left(\frac{U_h}{U_l} \right)^{\frac{1}{S-1}}, \quad (5.85)$$

$$\sigma_u = \frac{(a - 1)U_h}{(a + 1)\sqrt{2 \ln 2}}, \quad (5.86)$$

$$\sigma_v = \frac{\tan\left(\frac{\pi}{2K}\right) \left[U_h - \left(\frac{\sigma_u^2}{U_h} \right) 2 \ln 2 \right]}{\left[2 \ln 2 - \frac{(2 \ln 2)^2 \sigma_u^2}{U_h^2} \right]^{\frac{1}{2}}}, \quad (5.87)$$

where U_l and U_h denote the lower and upper center frequencies of interest. The K and S parameters are, respectively, the number of orientations and the number of scales in the desired multiresolution decomposition procedure. The sinusoid frequency W is set equal to U_h , and $m = 0, 1, \dots, S - 1$.

In the application being considered, interest lies only in image analysis, without the requirement of exact reconstruction or synthesis of the image from




FIGURE 5.69

Bank of Gabor filters designed in the frequency domain. Each ellipse represents the range of the corresponding filter response from 0.5 to 1.0 in squared magnitude (only one half of the response is shown for each filter). The sampling of the frequency spectrum can be adjusted by changing the U_l , U_h , S , and K parameters of the Gabor wavelets. Only the filters shown shaded are used to enhance the directional piecewise-linear structures present in the ROI images. The frequency axes are normalized. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Automatic identification of the pectoral muscle in mammograms”, *IEEE Transactions on Medical Imaging*, 23: 232 – 245, 2004. © IEEE.

the filtered components. Therefore, instead of using the wavelet coefficients, Ferrari et al. [278] used the magnitude of the filter response, computed as

$$a_{m,n}(x, y) = |f(x, y) * \psi_{m,n}^{\text{even}}(x, y)|, \quad (5.88)$$

where $\psi_{m,n}^{\text{even}}(x, y)$ indicates the even-symmetric part of the complex Gabor filter, $f(x, y)$ is the ROI being filtered, and $*$ represents 2D convolution. The phase and magnitude images, indicating the local orientation, were composed by vector summation of the K filtered images ([387], chapter 11); see also Figure 8.10.

The area of each ellipse indicated in Figure 5.69 represents the frequency spectrum covered by the corresponding Gabor filter. Once the range of the frequency spectrum is adjusted, the choice of the number of scales and orientations is made in order to cover the range of the spectrum as required. The choice of the number of scales (S) and orientations (K) used by Ferrari et al. for detecting the pectoral muscle was based upon the resolution required for detecting oriented information with high selectivity [388, 389]. The spatial-frequency bandwidths of the simple and complex cells in mammalian visual systems have been found to range from 0.5 to 2.5 octaves, clustering around 1.2 octaves and 1.5 octaves, and their angular bandwidth is expected to be smaller than 30° [389, 390]. By selecting $U_l = 0.05$, $U_h = 0.45$, $S = 4$, and $K = 12$ for processing mammographic images, Ferrari et al. [278, 381] indirectly adjusted the Gabor wavelets to have a frequency bandwidth of approximately one octave and angular bandwidth of 15° .

In the work of Ferrari et al., all images were initially oriented so that the chest wall was always positioned on the left-hand side; then, the pectoral muscle edge in correctly acquired MLO views will be located between 45° and 90° ([391], p. 34). (Here, the orientation of the pectoral muscle edge is defined as the angle between the horizontal line and an imaginary straight line representing the pectoral muscle edge.) For this reason, Ferrari et al. used only the Gabor filters with the mean orientation of their responses in the image domain at 45° , 60° , and 75° ; the corresponding frequency-domain responses are shown shaded in Figure 5.69.

Post-processing and pectoral muscle edge detection: In the method of Ferrari et al., after computing the phase and magnitude images by vector summation, the relevant edges in the ROI are detected by using an algorithm proposed by Ma and Manjunath [392] for edge-flow propagation, as described below. The magnitude $A(x, y)$ and phase $\phi(x, y)$ at each image location (x, y) are used to represent the edge-flow vector instead of using a predictive coding model as initially proposed by Ma and Manjunath. The phase at each point in the image is propagated until it reaches a location where two opposite directions of flow encounter each other, as follows:

1. Set $n = 0$ and $\mathbf{E}_0(x, y) = [A(x, y) \cos \phi(x, y), A(x, y) \sin \phi(x, y)]$.
2. Set the edge-flow vector $\mathbf{E}_{n+1}(x, y)$ at iteration $n + 1$ to zero.

3. At each image location (x, y) , identify the neighbor (x', y') that has the same direction θ as that of the edge-flow vector $\mathbf{E}_n(x, y)$. The direction θ is computed as $\theta = \tan^{-1} \frac{(y'-y)}{(x'-x)}$.
4. If $\mathbf{E}_n(x', y') \bullet \mathbf{E}_n(x, y) > 0$
 then $\mathbf{E}_{n+1}(x', y') = \mathbf{E}_{n+1}(x', y') + \mathbf{E}_n(x, y)$
 else $\mathbf{E}_{n+1}(x, y) = \mathbf{E}_{n+1}(x, y) + \mathbf{E}_n(x, y)$,
 where the symbol \bullet indicates the dot-product operation.
5. If nothing has been changed
 then stop iterating
 else go to Step 2 and repeat the procedure.

Figures 5.70 (b) and (c) illustrate an example of the orientation map before and after applying the edge-flow propagation procedure to the image shown in part (a) of the same figure.




FIGURE 5.70

(a) Region indicated by the box in Figure 5.68 (a) containing a part of the pectoral muscle edge. (b) and (c) Edge-flow map before and after propagation. Each arrowhead represents the direction of the edge-flow vector at the corresponding position in the image. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Automatic identification of the pectoral muscle in mammograms”, *IEEE Transactions on Medical Imaging*, 23: 232 – 245, 2004. © IEEE.


After propagating the edge-flow vector, the boundary candidates for the pectoral muscle are obtained by identifying the locations that have nonzero edge-flow arriving from two opposite directions. Weak edges are eliminated by thresholding the ROI image with a threshold value of 10% of the maximum gray-level value in the ROI. Figure 5.71 shows images resulting from each stage of the procedure for pectoral muscle detection.

In order to connect the disjoint boundary segments that are usually present in the image after the edge-flow propagation step, a half-elliptical neighborhood is defined with its center located at each boundary pixel being processed. The half-ellipse is adjusted to be proportional to the length of the contour, with $R_1 = 0.2 C_{\text{len}}$ and $R_2 = 5$ pixels (where R_1 and R_2 are, respectively, the major and minor axes of the half-ellipse, and C_{len} is the length of the boundary), with its major axis oriented along the direction of the contour line. If an ending or a starting pixel of an unconnected line segment is found in the defined neighborhood, it is connected by linear interpolation. The iterative method stops when all disjoint lines are connected; this procedure took 5 to 20 iterations with the images used in the work of Ferrari et al. Next, the false edges that may result in the filtered images due to structures inside the fibroglandular disc or due to the filtering process [see Figure 5.71 (c)–(d)] are removed by checking either if their limiting points are far away from the upper and left-hand side of the ROI, or if the straight line having the same slope as the pectoral muscle edge candidate intercepts outside the upper and left-hand limits of the ROI. Finally, the largest line in the ROI is selected to represent the pectoral muscle edge; see Figure 5.71 (e).


5.10.3 Results of application to mammograms

Ferrari et al. tested their methods with 84 images randomly chosen from the Mini-MIAS database [376]. All images were MLO views with $200 \mu\text{m}$ sampling interval and 8-bit gray-level quantization. For reduction of processing time, all images were downsampled with a fixed sampling distance so that the original images with the matrix size of $1,024 \times 1,024$ pixels were transformed to 256×256 pixels. The results obtained with the downsampled images were mapped to the original $1,024 \times 1,024$ mammograms for subsequent analysis and display.

The results obtained with both of the Hough-transform-based and Gabor-wavelet-based methods were evaluated in consultation with two radiologists experienced in mammography. The test images were displayed on a computer monitor with diagonal size of 47.5 cm and dot pitch of 0.27 mm . By using the Gimp program [380], the contrast and brightness of each image were manually enhanced so that the pectoral muscle edge could be easily visualized. Then, the pectoral muscle edges were manually drawn and the results printed on paper by using a laser printer with 600 dpi resolution. The zoom option of the Gimp program was used to aid in drawing the edges. The pectoral muscle edges of all images were checked by a radiologist using the printed images



(a)




(b)



Figure 5.71 (c)



(d)



(e)

FIGURE 5.71

Result of each stage of the Gabor-wavelet-based method: (a) ROI used. (b) Image magnitude after filtering and vector summation, enhanced by gamma correction ($\gamma = 0.7$). (c)–(d) Results before and after the post-processing stage. (e) Final boundary. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Automatic identification of the pectoral muscle in mammograms”, *IEEE Transactions on Medical Imaging*, 23: 232 – 245, 2004. © IEEE.

(hardcopy) along with the displayed images (softcopy); the assessment was recorded for further analysis.

The segmentation results were evaluated based upon the number of FP and FN pixels normalized with reference to the corresponding numbers of pixels in the regions demarcated by the manually drawn edges. The reference region for the pectoral muscle was defined as the region contained between the left-hand edge of the image and the hand-drawn pectoral muscle edge. An FP pixel was defined as a pixel outside the reference region that was included in the pectoral region segmented. An FN pixel was defined as a pixel in the reference region that was not present within the segmented region. Table 5.1 provides a summary of the results.

TABLE 5.1

Average False-positive and False-negative Rates in the Detection of the Pectoral Muscle by the hough-transform-based and Gabor-wavelet-based Methods.

Method	Hough	Gabor
FP $\pm \sigma$	1.98 \pm 6.09%	0.58 \pm 4.11%
FN $\pm \sigma$	25.19 \pm 19.14%	5.77 \pm 4.83%
Number of images with		
(FP and FN) < 5%	10	45
5% < (FP and FN) < 10%	8	22
(FP and FN) > 10%	66	17

Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, "Automatic identification of the pectoral muscle in mammograms", *IEEE Transactions on Medical Imaging*, 23: 232 – 245, 2004. © IEEE.

In the example illustrated in Figure 5.72, detection using the Hough transform resulted in an underestimated pectoral region due to the limitation imposed by the straight-line hypothesis used to represent the pectoral muscle edge; this translated to a high FN rate. The segmentation result of the Gabor-wavelet-based method is closer to the pectoral muscle edge drawn by the radiologist.




Figure 5.72 (a)




Figure 5.72 (b)





Figure 5.72 (c)

Figure 5.73 shows a case where the pectoral muscle appears as an almost-straight line. Even in this case, the result obtained using Gabor wavelets is more accurate than the result obtained using the Hough transform. The Gabor-wavelet-based method provided good results even in cases where the pectoralis minor was present in the mammogram (see Figure 5.74).

Detection of the pectoral muscle was used as a preprocessing step in segmentation of the fibroglandular discs in mammograms for the analysis of bilateral asymmetry by Ferrari et al. [381] (see Section 8.9).

5.11 Application: Improved Segmentation of Breast Masses by Fuzzy-set-based Fusion of Contours and Regions

Given the difficult nature of the problem of the detection of masses and tumors in a mammogram, the question arises “Can the problem benefit from the use of multiple approaches?” Guliato et al. [276] proposed two approaches to the detection problem: one based upon contour detection, and the other based



(d)

FIGURE 5.72

Results obtained for the image mdb003 from the Mini-MIAS database. (a) Original image. (b) Hand-drawn pectoral muscle edge superimposed on the histogram-equalized image. (c) and (d) Pectoral muscle edges detected by the Hough-transform-based and Gabor-wavelet-based methods, respectively, superimposed on the original image. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Automatic identification of the pectoral muscle in mammograms”, *IEEE Transactions on Medical Imaging*, 23: 232 – 245, 2004. © IEEE.




Figure 5.73 (a)




Figure 5.73 (b)





Figure 5.73 (c)

upon a fuzzy region-growing method. The former method is simple and easy to implement, always produces closed contours, and yields good results even in the presence of high levels of noise (see Section 5.5.2); the latter produces a fuzzy representation of the ROI, and preserves the uncertainty around the boundaries of tumors (see Section 5.5.3). As a follow-up, Guliato et al. [277] considered the following question: How may we combine the results of the two approaches — which may be considered to be complementary — so as to obtain a possibly better result?

In generic terms, the process of image segmentation may be defined as a procedure that groups the pixels of an image according to one or more local properties. A property of pixels is said to be *local* if it depends only on a pixel or its immediate neighborhood (for example, gray level, gradient, and local statistical measures). Techniques for image segmentation may be divided into two main categories: those based on discontinuity of local properties, and those based on similarity of local properties [305]. The techniques based on discontinuity are simple in concept, but generally produce segmented regions with disconnected edges, requiring the application of additional methods (such as contour following). Techniques based on similarity, on the other hand, depend on a seed pixel (or a seed subregion) and on a strategy to traverse the image for region growing. Because different segmentation methods explore distinct, and sometimes complementary, characteristics of the given image



(d)

FIGURE 5.73

Results obtained for the image mdb008 from the Mini-MIAS database. (a) Original image. (b) Hand-drawn pectoral muscle edge superimposed on the histogram-equalized image. (c) and (d) Pectoral muscle edges detected by the Hough-transform-based and Gabor-wavelet-based methods, respectively, superimposed on the original image. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Automatic identification of the pectoral muscle in mammograms”, *IEEE Transactions on Medical Imaging*, 23: 232 – 245, 2004. © IEEE.


**FIGURE 5.74**

Image mdb110 from the Mini-MIAS database, showing the result of the detection of the pectoral muscle in the presence of the pectoralis minor. (a) Edge candidates after the post-processing stage. (b) Final boundary detected by the Gabor-wavelet-based method. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, R.A. Borges, and A.F. Frère, “Automatic identification of the pectoral muscle in mammograms”, *IEEE Transactions on Medical Imaging*, 23: 232 – 245, 2004. © IEEE.

(such as contour detection and region growing), it is natural to consider combinations of techniques that could possibly produce better results than any one technique on its own.

Although cooperative combination of the results of segmentation procedures can offer good results, there are only a few publications devoted to this subject [314, 316, 393, 394, 395, 396, 397]. This is partly due to the difficulty in simultaneously handling distinct local properties, and due to the limitations of the commonly used Boolean-set operations in combining multiple results of image segmentation. Using the theory of fuzzy sets, it is possible to define several classes of fusion operators that generalize Boolean operators. Guliato et al. [277] proposed a general fusion operator, oriented by a finite automaton, to combine information from different sources. The following paragraphs provide descriptions of the method and present results of the fusion operator applied to tumor regions in mammographic images.

Elementary concepts of fusion operators: A *fusion operator* over fuzzy sets is formally defined as a function $h : [0, 1]^n \rightarrow [0, 1]$, where $n \geq 2$ represents the number of sources of input information. Fusion operators may be classified according to their behavior into three classes: conjunctive, disjunctive, and compromise operators [351, 398], as follows:

- An operator is said to be *conjunctive* if $h(a_1, a_2, \dots, a_n) \leq \min \{a_1, a_2, \dots, a_n\}$, where $a_i \in [0, 1]$. Conjunctive operators are those that represent a consensus between the items of information being combined. They generalize classical intersection, and agree with the source that offers the smallest measure while trying to obtain simultaneous satisfaction of its criteria. We can say that conjunctive operators present a *severe* behavior.
- An operator is said to be *disjunctive* if $h(a_1, a_2, \dots, a_n) \geq \max \{a_1, a_2, \dots, a_n\}$. Disjunctive operators generalize classical union. They agree with the source that offers the greatest measure, and express redundancy between criteria. We can say that they present a *permissive* behavior.
- An operator is said to be a *compromise operator* if $\min\{a_1, a_2, \dots, a_n\} \leq h(a_1, a_2, \dots, a_n) \leq \max\{a_1, a_2, \dots, a_n\}$. Compromise operators produce an intermediate measure between items of information obtained from several sources. They present *cautious* behavior.

Bloch [399] presented a classification scheme that describes a fusion operator in more refined terms not only as conjunctive, disjunctive, or compromise, but also according to its behavior with respect to the information values being combined (input values): context-independent constant-behavior operators that maintain the same behavior independent of the input variables; context-independent variable-behavior operators, whose behavior varies according to the input variables; and context-dependent operators, whose behavior varies as in the previous case, also taking into account the agreement between the

sources and their reliability. The following paragraphs provide the description of a class of fusion operators that generalize context-dependent operators, taking into consideration different degrees of confidence in the sources, specific knowledge, and spatial context while operating with conceptually distinct sources.

Considerations in the fusion of the results of complementary segmentation techniques: Figure 5.75 illustrates an overlay of two segmentation results obtained by two complementary techniques — region growing represented by a fuzzy set S_r , and closed-contour detection represented by a fuzzy set S_c — for the same ROI. The straight line within S_r indicates a possible artifact. The results are not the same: different segmentation algorithms may produce different results for the same ROI. A fusion operator designed to aggregate such entities should produce a third entity that takes into consideration the inputs and is better than either input on its own. In order to realize this, the fusion operator must be able to identify regions of certainty and uncertainty during its execution.




FIGURE 5.75

Superimposition of the results of two complementary segmentation techniques. The circular region S_r represents the result of region growing. The square box S_c represents the result of contour detection. Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, “Fuzzy fusion operators to combine results of complementary medical image segmentation techniques”, *Journal of Electronic Imaging*, 12(3): 379 – 389, 2003. © SPIE and IS&T.

Considering a pixel p being analyzed, let $\Gamma_{S_r}(p)$ be the membership degree of p , such that $S_r = \Gamma_{S_r} : I \rightarrow [0, 1]$, where I is the original image. Also, let $\Gamma_{S_c}(p)$ be the membership degree of p , such that $S_c = \Gamma_{S_c} : I \rightarrow [0, 1]$. It is important to note that $\Gamma_{S_c}(p)$ is zero when the pixel p is inside or outside of S_c , and that $\Gamma_{S_c}(p)$ possesses a high value when p is on the contour represented by S_c . Similarly, $\Gamma_{S_r}(p)$ is high when p belongs to the region, and $\Gamma_{S_r}(p)$ is low or zero when p does not belong to the region. With respect to the fusion operator, four situations may be identified considering the position of p (see Figure 5.76):




FIGURE 5.76

The four different situations treated by the fusion operator. Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, “Fuzzy fusion operators to combine results of complementary medical image segmentation techniques”, *Journal of Electronic Imaging*, 12(3): 379 – 389, 2003. © SPIE and IS&T.

1. p belongs to the intersection of S_r and S_c [that is, $\Gamma_{S_r}(p)$ is high and $\Gamma_{S_c}(p)$ is zero]. In this case the pixel p belongs to the final segmentation result with a high membership degree. The sources agree with respect to the inclusion of the pixel p in the final result. This is a case of certainty.
2. p does not belong to S_r or belongs to S_r with a low membership degree, and is inside S_c [that is, $\Gamma_{S_r}(p)$ is low or zero and $\Gamma_{S_c}(p)$ is zero]. In this case the sources disagree with respect to the inclusion of the pixel p in the final result. This is a case of uncertainty.
3. p belongs to the contour line of S_c [that is, $\Gamma_{S_c}(p)$ is high] and does not belong to S_r [that is, $\Gamma_{S_r}(p)$ is low or zero]. As in Item 2 above, this is an uncertainty situation. Note that although the inputs are different from those presented in Item 2 above, the result of the fusion operator is expected to represent uncertainty.
4. p belongs to S_r [that is, $\Gamma_{S_r}(p)$ is high] and is outside of S_c [that is, $\Gamma_{S_c}(p)$ is zero]. Here again we have an uncertainty case. Observe that although the inputs are similar to those in Item 1 [that is, $\Gamma_{S_r}(p)$ is high and $\Gamma_{S_c}(p)$ is zero], the result of the fusion operator is expected to be different.

We can conclude from the discussion above that a practically applicable fusion operator should be composed of a number of basic fusion operators, and that the spatial position of the pixel being analyzed is an important item of information that should be used in determining the basic fusion operator to be applied to the pixel. Based upon these observations, Guliato et al. [277] proposed a general fusion operator oriented by a finite automaton, where the finite set of states of the automaton is determined by the spatial position of

the pixel being analyzed, and where the transition function (to be defined later) depends on the strategy used to traverse the image.

An important question to be considered in fusion is the reliability of the sources (original segmentation results). The result of the fusion operator depends on how good the original segmentation results are. The evaluation of the individual segmentation results is not a component of the fusion procedure, although parameters are included in the definitions of the operators to represent the reliability of the sources; it is assumed that the parameters are determined using other methods.

General finite-automaton-oriented fusion operators: Formally, a *fusion operator oriented by a finite automaton* that aggregates n sources may be defined as an ordered pair $\langle H, M \rangle$, where [351]:

- $H = \{h_1, h_2, \dots, h_k\}$ is a finite set of *basic fusion operators*, where h_i are functions that map $[0, 1]^n \rightarrow [0, 1]$, $n \geq 2$.
- $M = (Q, \Sigma, \delta, q_0, F)$ is a finite automaton, where:
 - Q is a finite set of states,
 - Σ is a finite input alphabet,
 - δ is a transition function that maps $Q \times \Sigma \rightarrow Q$, where \times is the Cartesian product operator,
 - $q_0 \in Q$ is an initial state, and
 - $F \subset Q$ is the set of final states.

In the present case, the alphabet Σ is given by a finite collection of labels associated with the Cartesian product of finite partitions of the interval $[0, 1]$. For example, suppose that, coming from different motivations, we are dividing $[0, 1]$ into two finite partitions P_1 and P_2 , where P_1 divides the values between ‘low’ and ‘high’, and P_2 between ‘good’ and ‘bad’. Our alphabet may be composed of $\Sigma = \{0, 1, 2\}$ representing, for example, the combinations (low, good), (low, bad), and (high, good), respectively. Observe that we are not necessarily using the whole set of possibilities.

The interpretation of the transition function δ of a finite automaton is the following: $\delta(q_i, a) = q_j$ is a valid transition if, and only if, the automaton can go from the state q_i to q_j through the input a . Sometimes, q_i and q_j could be the same state. If there is a transition from the state q_i to q_j through the input a , then there is a directed arc from q_i to q_j with the label a in the graphical representation (transition diagram) of the specific automaton; see Figure 5.77.

Application of the fusion operator to image segmentation: The fusion operator proposed by Guliato et al. [277] is designed to combine the results obtained from two segmentation techniques that explore complementary characteristics of the image: one based on region growing, and the other




FIGURE 5.77

Graphical representation of the transition function given by $\delta(q_i, a) = q_j$ and $\delta(q_i, b) = q_i$. Reproduced with permission from D. Guliatò, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, “Fuzzy fusion operators to combine results of complementary medical image segmentation techniques”, *Journal of Electronic Imaging*, 12(3): 379 – 389, 2003. © SPIE and IS&T.

based on closed-contour detection [276]. The result of fusion obtained is a fuzzy set that represents the agreement or disagreement between the input sources.

Let S_r (based on region growing) and S_c (based on closed-contour detection) represent the two segmented images to be combined, as shown in Figure 5.75. The process starts with a seed pixel selected by the user. The seed pixel must belong to the intersection $S_r \cap S_c$. It is assumed that S_r and S_c are each endowed with a reliability measure, given by a number in the interval $[0, 1]$.

The fusion operator is represented by $\mathcal{O} = \langle H, M \rangle$, where $H = \{h_1, h_2, \dots, h_6\}$ is a collection of six basic fusion operators (that take into consideration the reliability measures of the sources, as explained below), and M is a finite automaton that governs the actions of the operator.

In the following description of the basic fusion operators, the parameters C_r and C_c range within the interval $[0, 1]$ and denote the reliability measures of the sources S_r and S_c , respectively. The parameters are used to indicate the influence that a given source should have on the final result of the fusion operation: the higher the value, the larger the influence of the source.

The result of each basic fusion operator should give information about the agreement among the sources being analyzed. The absence of conflict is represented by a membership degree equal to 1 or 0; that is, both the sources agree or do not agree with respect to the membership of the given pixel in the ROI. Maximal conflict is represented by membership degree equal to 0.5; in this case, the sources do not agree with respect to the membership of the given pixel. Intermediate membership degrees denote intermediate degrees of agreement.

Let p_{ij} be the j^{th} pixel of the segmented image S_i and $\Gamma_{S_i}(p_{ij})$ be the membership degree of the pixel p_{ij} , where $i \in \{r, c\}$, $j = 1, 2, \dots, m$, and m is the total number of pixels in the image I . (Note: In the present section, we are using only one index j to represent the position of a pixel in an image.) Then, the basic fusion operators are defined as follows:

1. $h_1 = \max\{C_r \Gamma_{S_r}(p_{rj}), C_c, 0.5\}$.

This is a disjunctive operator that associates with the pixels in $S_r \cap S_c$ new membership degrees taking into account the source with the greater reliability measure (see h_1 in Figure 5.78).

2. if $\max(C_r, C_c) \leq 0.5$

then $h_2 = 0.5$

else if $(C_r \leq 0.5)$

then $h_2 = C_c \Gamma_{S_c}(p_{cj})$

else if $(C_c < 0.5)$

then $h_2 = C_r \Gamma_{S_r}(p_{rj})$

else $h_2 = \frac{1}{(C_r + C_c)} [C_r \Gamma_{S_r}(p_{rj}) + C_c \Gamma_{S_c}(p_{cj})]$.

This is a compromise operator that acts on the pixels belonging to the transition region between the interior and the exterior of the result of contour detection (see h_2 in Figure 5.78).

3. if $\max(C_r, C_c) \leq 0.5$

then $h_3 = 0.5$

else if $(C_r \leq 0.5)$

then $h_3 = C_c$

else if $(C_c \leq 0.5)$

then $h_3 = C_r \Gamma_{S_r}(p_{rj})$

else $h_3 = \frac{1}{(C_r + C_c)} [C_r \Gamma_{S_r}(p_{rj}) + C_c]$.

This is a compromise operator that acts on the pixels lying outside the result of region growing and belonging to the interior of the result of contour detection (see h_3 in Figure 5.78).

4. if $\max(C_r, C_c) \leq 0.5$

then $h_4 = 0.5$

else if $(C_r \leq 0.5)$

then $h_4 = 0$

else if ($C_c \leq 0.5$)

then $h_4 = C_r \Gamma_{S_r}(p_{rj})$

else $h_4 = \frac{1}{(C_r + C_c)} \{C_r \Gamma_{S_r}(p_{rj}) + [1 - \sqrt{C_c}]^2\}$.

This is a compromise operator that acts on the pixels lying outside the result of contour detection and belonging to the interior of the result of region growing (see h_4 in Figure 5.78). Artifacts within the region-growing result (as indicated schematically by the line segment inside the circle in Figure 5.78) are rejected by this operator.

5. $h_5 = \max\{C_r \Gamma_{S_r}(p_{rj}), C_c \Gamma_{S_c}(p_{cj}), 0.5\}$.

This is a disjunctive operator that acts on the transition pixels lying in the intersection $S_r \cap S_c$ (see h_5 in Figure 5.78).

6. if $\max(C_r, C_c) \leq 0.5$

then $h_6 = 0.0$

else if ($C_r \leq 0.5$)

then $h_6 = 0.0$

else if ($C_c \leq 0.5$)

then $h_6 = C_r \Gamma_{S_r}(p_{rj})$

else $h_6 = \min\{C_r \Gamma_{S_r}(p_{rj}), [1 - C_c]\}$.

This is a conjunctive operator that acts on the exterior of $S_r \cup S_c$ and determines a limiting or stopping condition for the operator (see h_6 in Figure 5.78).

Description of the finite automaton: The finite automaton $M = (Q, \Sigma, \delta, q_0, F)$ in \mathcal{O} is defined by the following entities:

- A set of finite states $Q = \{a, b, c\}$, where
 - state a indicates that the pixel being analyzed belongs to the interior of the contour,
 - state b indicates that the pixel being analyzed belongs to the contour, and
 - state c indicates that the pixel being analyzed belongs to the exterior of the contour (see Figure 5.79).




FIGURE 5.78

The regions where the six basic fusion operators are applied are indicated by $\{h_1, h_2, h_3, h_4, h_5, h_6\}$. Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, “Fuzzy fusion operators to combine results of complementary medical image segmentation techniques”, *Journal of Electronic Imaging*, 12(3): 379 – 389, 2003. © SPIE and IS&T.

- A finite input alphabet $\Sigma = \{I_1, I_2, I_3, I_4\}$.

Let π_1 and π_2 be two finite partitions of $[0, 1]$, where $\pi_1 = \pi_2 = \{high, low\}$. We can choose the classes *high* and *low* as follows:

$$low = [0, 0.5),$$

$$high = [0.5, 1.0];$$

$p_{ij} \in high$ if $\Gamma_{S_i}(p_{ij}) \geq 0.5$ for $j = 1, 2, \dots, m$, and

$p_{ij} \in low$ if $\Gamma_{S_i}(p_{ij}) < 0.5$ for $j = 1, 2, \dots, m$,

where p_{ij} , $i \in \{r, c\}$, and $j = 1, 2, \dots, m$, identify the i^{th} source and the j^{th} pixel; and $\Gamma_{S_i}(p_{ij})$ is the membership degree of the pixel p_j in S_i .

The finite input alphabet Σ is produced by the function $\mu : \pi_1 \times \pi_2 \rightarrow \Sigma$, where:

- $\mu(high, low) = I_1$. The pixel being analyzed presents a high membership degree in the region-growing segmentation result and a low membership degree in the closed-contour result. This input represents a certainty or uncertainty situation depending on the spatial position of the pixel being analyzed; see Figure 5.80.
- $\mu(high, high) = I_2$. The pixel being analyzed presents a high membership degree in the region-growing segmentation result and a high

membership degree in the closed-contour result. This indicates an intersection case; see Figure 5.80.

- $\mu(\text{low}, \text{high}) = I_3$. The pixel being analyzed presents a low membership degree in the region-growing segmentation result and a high membership degree in the closed-contour result. This indicates an uncertainty case; see Figure 5.80.
 - $\mu(\text{low}, \text{low}) = I_4$. The pixel being analyzed presents a low membership degree in the region-growing segmentation result and a low membership degree in the closed-contour result. This indicates an uncertainty case if the pixel belongs to the interior of the contour; in the opposite case, this indicates a stopping or limiting condition of the fusion operator; see Figure 5.80.
- A transition diagram δ of M , as shown in Figure 5.81.

The transition diagram illustrates the situations when the basic fusion operator is executed. The analysis begins with a pixel that belongs to the intersection of the two segmentation results. The first input must be of type I_1 ; the initial state of the automaton is a , which corresponds to the fact that the pixel belongs to the interior of the contour. The analysis procedure is first applied to all of the pixels inside the contour. While the inputs are I_1 or I_4 , the operators h_1 or h_3 will be applied and the automaton remains in state a . When an input of type I_2 or I_3 arrives, the automaton goes to state b to inform the analysis process that the pixel being processed is on the boundary given by the contour-detection method. At this stage, all the pixels on the contour are processed. While the inputs are I_2 or I_3 and the operators h_5 or h_2 are applied, the automaton will remain in state b . If, while in state b , the input I_1 or I_4 occurs (and the operators h_4 or h_6 applied), the automaton goes to state c , indicating that the pixel being analyzed is outside the contour. All the pixels outside the contour are processed at this stage. Observe that, depending upon the state of the automaton, different fusion operators may be applied to the same inputs. As indicated by the transition diagram in Figure 5.81, all of the pixels in the interior of the contour are processed first; all of the pixels on the contour are processed next, followed by the pixels outside the contour.

- The initial state $q_0 \in Q$, with $q_0 = \{a\}$.
- The set of final states $F = \{c\}$, where $F \subseteq Q$. In the present case, F has only one element.

Behavior of the basic fusion operators: The fusion operator described above can combine several results of segmentation, two at a time. The result yielded by the fusion operator is a fuzzy set that identifies the certainty and uncertainty present in the inputs to the fusion process. It is expected that

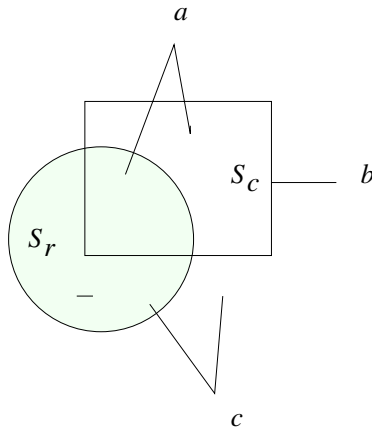


FIGURE 5.79

The three states of the automaton. Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, “Fuzzy fusion operators to combine results of complementary medical image segmentation techniques”, *Journal of Electronic Imaging*, 12(3): 379 – 389, 2003. © SPIE and IS&T.




FIGURE 5.80

The four possible input values $\{I_1, I_2, I_3, I_4\}$ for the fusion operator. The short line segments with the labels I_2 and I_3 represent artifacts in the segmentation result. Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, “Fuzzy fusion operators to combine results of complementary medical image segmentation techniques”, *Journal of Electronic Imaging*, 12(3): 379 – 389, 2003. © SPIE and IS&T.




FIGURE 5.81

Transition diagram that governs the actions of the fusion operator. Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, “Fuzzy fusion operators to combine results of complementary medical image segmentation techniques”, *Journal of Electronic Imaging*, 12(3): 379 – 389, 2003. © SPIE and IS&T.

maximal certainty will be represented by a membership degree equal to 1 or 0 (that is, the pixel being analyzed certainly belongs to or does not belong to the final segmentation result). When individual segmentation results disagree with respect to a pixel belonging or not belonging to the final result, or when both sources do not present sufficient reliability, the fusion operator yields a membership degree equal to 0.5 to represent a situation with maximal uncertainty. Other situations are represented by membership degrees ranging in the interval $(0, 0.5)$ and $(0.5, 1)$, depending on the evidence with respect to the membership of the analyzed pixel in the ROI and the reliability of the sources. Two illustrative studies on the behavior of the basic fusion operators h_1 and h_2 are presented below, taking into consideration a limited set of entries:

- State a of the automaton (inside the contour), entry 1 (*high, low*), basic fusion operator $h_1 = \max\{C_r \Gamma_{S_r}(p_{rj}), C_c, 0.5\}$.

This is the starting condition of the fusion operator; see Figure 5.81. The starting pixel must lie in the intersection of S_r and S_c (see Figures 5.78 and 5.79). In this case, $\Gamma_{S_r}(p_{rj}) \geq 0.5$ (that is, p belongs to the region-growing result with a high membership degree) and $\Gamma_{S_c}(p_{c_j}) < 0.5$ (that is, p is inside the contour). This situation represents the condition that both sources agree with respect to the pixel p belonging to the ROI. Table 5.2 provides explanatory comments describing the behavior of h_1 for several values of the reliability parameters and inputs from the two sources.

- State a of the automaton (inside the contour) or state b (on the contour), entry 3 (*low, high*), basic fusion operator h_2 .

The operator h_2 is applied when the automaton is in the state a or b and a transition occurs from a pixel inside the contour to a pixel on

the contour [that is, $\delta(a, I_3) \rightarrow b$] or from a pixel on the contour to a neighboring pixel on the contour [that is, $\delta(b, I_3) \rightarrow b$].

In this case, $\Gamma_{S_r}(p_{rj}) < 0.5$ (that is, p does not belong to the region-growing result or belongs with a low membership degree), and $\Gamma_{S_c}(p_{cj}) > 0.5$ (that is, p is on the contour). This situation represents the condition where the sources disagree with respect to the pixel p belonging to the ROI. The result of h_2 is a weighted average of the input membership values. Table 5.3 provides explanatory comments describing the behavior of h_2 for several values of the reliability parameters and inputs from the two sources.

TABLE 5.2
Behavior of the Basic Fusion Operator h_1 .

C_r	$\Gamma_{S_r}(p_{rj})$	C_c	$\Gamma_{S_c}(p_{cj})$	h_1	Comments
1.0	1.0	1.0	0.0	1.0	p belongs to the ROI with maximal certainty
1.0	1.0	0.0	0.0	1.0	Result depends on the source with the higher reliability
0.0	1.0	1.0	0.0	1.0	Result depends on the source with the higher reliability
0.0	1.0	0.0	0.0	0.5	Both sources do not present reliability
0.8	1.0	1.0	0.0	1.0	Source S_c presents the higher reliability
0.8	1.0	0.8	0.0	0.8	Result depends on the source with the higher reliability
0.9	1.0	0.3	0.0	0.9	Result depends on the source with the higher reliability

Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, "Fuzzy fusion operators to combine results of complementary medical image segmentation techniques", *Journal of Electronic Imaging*, 12(3): 379 – 389, 2003. © SPIE and IS&T.

Application of the fusion operator to the segmentation of breast tumors: Figure 5.82 shows the results of the fusion of the ROIs represented in

TABLE 5.3Behavior of the Basic Fusion Operator h_2 .

C_r	$\Gamma_{S_r}(p_{rj})$	C_c	$\Gamma_{S_c}(p_{cj})$	h_2	Comments
1.0	0.0	1.0	1.0	0.5	Weighted averaging
1.0	0.0	0.0	1.0	0.0	Result depends on the source with the higher reliability
0.0	0.0	1.0	1.0	1.0	Result depends on the source with the higher reliability
0.0	0.0	0.0	1.0	0.5	Both sources do not present reliability; maximal uncertainty
0.8	0.0	1.0	1.0	0.56	Weighted averaging

Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, "Fuzzy fusion operators to combine results of complementary medical image segmentation techniques", *Journal of Electronic Imaging*, 12(3): 379 – 389, 2003. © SPIE and IS&T.

Figures 5.30 and 5.32. The results have been superimposed with the contours drawn by an expert radiologist, for comparison. Figure 5.83 demonstrates the application of the methods for contour detection, fuzzy region growing, and fusion to a segment of a mammogram with a malignant tumor. Observe that the fusion results reduce the uncertainty present in the interior of the regions, but also reduce the certainty of the boundaries. The features of the results of the individual segmentation procedures contribute to the fusion results, allowing the postponement of a crisp decision (if necessary) on the ROI or its boundary to a higher level of the image analysis system.

Evaluation of the results of fusion using a measure of fuzziness:

In order to evaluate the results of the fusion operator, Guliato et al. [277] compared the degree of agreement between the reference contour given by an expert radiologist and each segmentation result: contour segmentation, region-growing segmentation, and the result of fusion. The reference contour and a segmentation result were aggregated using the fusion operator. The fusion operator yields a fuzzy set that represents the certainty and uncertainty identified during the aggregation procedure. The maximal certainty occurs when $\Gamma(p) = 0$ or $\Gamma(p) = 1$, where Γ is the membership degree of the pixel p . The maximal uncertainty occurs when $\Gamma(p) = 0.5$. In the former case, the information sources agree completely with respect to the pixel p ; in the latter, the information sources present maximal conflict with respect to the pixel p . Intermediate values of the membership degree represent intermediate degrees of agreement among the information sources. If the uncertainty presented



Figure 5.82 (a)

by the fusion result can be quantified, the result could be used to evaluate the degree of agreement among two different information sources. In order to quantify the uncertainty, Guliato et al. [277] proposed a measure of fuzziness.


In general, a *measure of fuzziness* is a function

$$f : \mathcal{F}(X) \rightarrow \mathcal{R}^+, \quad (5.89)$$

where $\mathcal{F}(X)$ denotes the set of all fuzzy subsets of X . For each fuzzy set A of X , this function assigns a nonnegative real number $f(A)$ that characterizes the degree of fuzziness of A . The function f must satisfy the following three requirements:

- $f(A) = 0$ if, and only if, A is a crisp set;
- $f(A)$ assumes its maximal value if, and only if, A is maximally fuzzy, that is, all of the elements of A are equal to 0.5; and
- if set A is undoubtedly sharper than set B , then $f(A) \leq f(B)$.

There are different ways of measuring fuzziness that satisfy all of the three essential requirements [351]. Guliato et al. [277] chose to measure fuzziness in terms of the distinctions between a set and its complement, observing that it is the lack of distinction between a set and its complement that distinguishes



(b)

FIGURE 5.82

Result of the fusion of the contour and region (a) in Figures 5.30 (c) and (d) for the case with a malignant tumor; and (b) in Figures 5.32 (c) and (d) for the case with a benign mass, with $C_r = 1.0$, $C_c = 1.0$. The contours drawn by the radiologist are superimposed for comparison. Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, “Fuzzy fusion operators to combine results of complementary medical image segmentation techniques”, *Journal of Electronic Imaging*, 12(3): 379 – 389, 2003. © SPIE and IS&T.




Figure 5.83 (a)




Figure 5.83 (b)




Figure 5.83 (c)

a fuzzy set from a crisp set. The implementation of this concept depends on the definition of the fuzzy complement; the standard complement is defined as $\bar{A}(x) = 1 - A(x)$, for all $x \in X$. Choosing the Hamming distance, the local distinction between a given set A and its complement \bar{A} is measured by

$$|A(x) - \{1 - A(x)\}| = |2A(x) - 1|, \quad (5.90)$$

and the lack of local distinction is given by


$$1 - |2A(x) - 1|. \quad (5.91)$$

The measure of fuzziness, $f(A)$, is then obtained by adding the local measurements:

$$f(A) = \sum_{x \in X} [1 - |2A(x) - 1|]. \quad (5.92)$$

The range of the function f is $[0, |X|]$: $f(A) = 0$ if, and only if, A is a crisp set; $f(A) = |X|$ when $A(x) = 0.5$ for all $x \in X$.

In the work reported by Guliato et al. [277], for each mammogram the reference contour drawn by the expert radiologist was combined, using the fusion operator, with each of the results obtained by contour detection, fuzzy region growing, and fusion, denoted by RS_c , RS_r , and RF_r , respectively. The fusion operator was applied with both the reliability measures equal to unity,



(d)

FIGURE 5.83

(a) A 700×700 -pixel portion of a mammogram with a spiculated malignant tumor. Pixel size = $62.5 \mu\text{m}$. (b) Contour extracted (white line) by fuzzy-set-based preprocessing and region growing. The black line represents the boundary drawn by a radiologist (shown for comparison). (c) Result of fuzzy region growing. The contour drawn by the radiologist is superimposed for comparison. (d) Result of the fusion of the contour in (b) and the region in (c) with $C_r = 1.0$, $C_c = 0.9$. The contour drawn by the radiologist is superimposed for comparison. Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, "Fuzzy fusion operators to combine results of complementary medical image segmentation techniques", *Journal of Electronic Imaging*, 12(3): 379 – 389, 2003. © SPIE and IS&T.

that is, $C_r = C_c = 1.0$, for the two information sources being combined in each case. When the result of contour detection was combined with the contour drawn by the radiologist, the former was converted into a region because the fusion method is designed to accept a contour and a region as the inputs.

Considering the results shown in Figure 5.83, the measures of fuzziness obtained were $f(RS_c) = 14,774$, $f(RS_r) = 14,245$, and $f(RF_r) = 9,710$, respectively. The aggregation or fusion of the two segmentation results presents lower uncertainty than either, yielding a better result as expected.

The methods were tested with 14 mammographic images of biopsy-proven cases; the values of the measure of fuzziness for the cases are shown in Table 5.4. The values of C_c and C_r used to obtain the result of fusion for the 14 mammograms are also listed in the table. Both C_c and C_r were maintained equal to unity when computing the measure of fuzziness with respect to the contour drawn by the radiologist for all the cases. In 11 cases, the fusion operator yielded improvement over the original results. There was no improvement by fusion in three of the cases: in one of these cases both segmentation results were not accurate, and in the other two, the fuzzy region segmentation was much better than the result of contour segmentation (based upon visual comparison with the reference contour drawn by the radiologist). The results provide good evidence that the fusion operator obtains regions with a higher degree of certainty than the results of the individual segmentation methods.

The measure of fuzziness may be normalized by division by $|X|$. However, in the context of the work of Guliato et al., this would lead to very small values because the number of boundary pixels is far less than the number of pixels inside a mass. The measure of fuzziness without normalization is adequate in the assessment of the results of fusion because the comparison is made using the measure for each mammogram separately.

5.12 Remarks

We have explored several methods to detect the edges of objects or to segment ROIs. We have also studied methods to detect objects of known characteristics, and methods to improve initial estimates of edges, contours, or regions. The class of filters based upon mathematical morphology [8, 192, 220, 221, 222] has not been dealt with in this chapter.

After ROIs have been detected and extracted from a given image, they may be analyzed further in terms of representation, feature extraction, pattern classification, and image understanding. Some of the measures and approaches that could be used for these purposes are listed below. It should be recognized that the accuracy of the measures derived will depend upon the accuracy of the results of detection or segmentation [400].

TABLE 5.4

Measures of Fuzziness for the Results of Segmentation and Fusion for 14 Mammograms.

Mammogram	C_r, C_c	$f(RS_c)$	$f(RS_r)$	$f(RF_r)$	Is the result of fusion better?
spic-s-1	1.0, 1.0	14,774	14,245	9,711	Yes
circ-fb-010	1.0, 1.0	8,096	9,223	7,905	Yes
spx111m	1.0, 1.0	5,130	9,204	4,680	Yes
spic-fh0	1.0, 0.6	28,938	23,489	21,612	Yes
circ-x-1	1.0, 0.8	6,877	2,990	3,862	No
spic-fh2	0.8, 1.0	45,581	38,634	34,969	Yes
circ-fb-005	1.0, 1.0	26,176	34,296	25,084	Yes
circ-fb-012	1.0, 0.9	16,170	15,477	12,693	Yes
spic-db-145	1.0, 0.9	8,306	7,938	7,658	Yes
circ-fb-025	1.0, 0.6	56,060	44,277	49,093	No
spic-fb-195	1.0, 1.0	11,423	12,511	10,458	Yes
spic-s-112	1.0, 0.6	31,413	17,784	12,838	Yes
spic-s-401	1.0, 0.6	13,225	11,117	11,195	No
circ-fb-069	1.0, 1.0	46,835	53,321	38,832	Yes

Reproduced with permission from D. Guliato, R.M. Rangayyan, W.A. Carnielli, J.A. Zuffo, and J.E.L. Desautels, "Fuzzy fusion operators to combine results of complementary medical image segmentation techniques", *Journal of Electronic Imaging*, 12(3): 379 – 389, 2003. © SPIE and IS&T.

- External characteristics:
 - boundary or contour morphology,
 - boundary roughness,
 - boundary complexity.

It is desirable that boundary descriptors are invariant to translation, scaling, and rotation. Methods for the analysis of contours and shape complexity are described in Chapter 6.

- Internal characteristics:
 - gray level,
 - color,
 - texture,
 - statistics of pixel population.

Methods for the analysis of texture are presented in Chapters 7 and 8.

- Description of (dis)similarity:
 - distance measures,
 - correlation coefficient.

Chapter 12 contains the descriptions of several methods based upon measures of similarity and distance; however, the methods are described in the context of pattern classification using vectors of features. Some of the methods may be extended to compare sets of pixels representing segmented regions.

- Relational description:
 - placement rules,
 - string, tree, and web grammar,
 - structural description,
 - syntactic analysis.

See Gonzalez and Thomason [401] and Duda et al. [402] for details on syntactic pattern recognition and image understanding.

5.13 Study Questions and Problems

Selected data files related to some of the problems and exercises are available at the site

www.enel.ucalgary.ca/People/Ranga/enel697

1. Give the definition of the 3×3 Sobel masks, and explain how they may be used to detect edges of any orientation in an image.
What are the limitations of this approach to edge detection?
What type of further processing steps could help in improving edge representation?
2. Prepare a 5×5 image with the value 10 in the central 3×3 region and the value zero in the remainder of the image.
Calculate the results of application of the 3×3 Laplacian operator and the masks given in Section 5.2 for the detection of isolated points and lines.
Evaluate the results in terms of the detection of edges, lines, points, and corners.

5.14 Laboratory Exercises and Projects

1. Create a test image with objects of various shapes and sizes. Process the image with LoG functions of various scales. Analyze the results in terms of the detection of features of varying size and shape.
2. Prepare a test image with a few straight lines of various slopes and positions. Apply the Hough transform for the detection of straight lines. Study the effect of varying the number of bins in the Hough space. Analyze the spreading of values in the Hough space and develop strategies for the detection of the straight lines of interest in the image.
What are the causes of artifacts with this method?

Analysis of Shape

Several human organs and biological structures possess readily identifiable shapes. The shapes of the human heart, brain, kidneys, and several bones are well known, and, in normal cases, do not deviate much from an “average” shape. However, disease processes can affect the structure of organs, and cause deviation from their expected or average shapes. Even abnormal entities, such as masses and calcifications in the breast, tend to demonstrate differences in shape between benign and malignant conditions. For example, most benign masses in the breast appear as well-circumscribed areas on mammograms, with smooth boundaries that are circular or oval; some benign masses may be macrolobulated. On the other hand, malignant masses (cancerous tumors) are typically ill-defined on mammograms, and possess a rough or stellate (star-like) shape with strands or spicules appearing to radiate from a central mass; some malignant masses may be microlobulated [54, 345, 403]. Shape is a key feature in discriminating between normal and abnormal cells in Pap-smear tests [272, 273]. However, biological entities demonstrate wide ranges of manifestation, with significant overlap between their characteristics for various categories. Furthermore, it should be borne in mind that the imaging geometry, 3D-to-2D projection, and the superimposition of multiple objects commonly affect the shapes of objects as perceived on biomedical images.

Several techniques have been proposed to characterize shape [404, 405, 406]. We shall study a selection of shape analysis techniques in this chapter. A few applications will be described to demonstrate the usefulness of shape characteristics in the analysis of biomedical images.

6.1 Representation of Shapes and Contours

The most general form of representation of a contour in discretized space is in terms of the (x, y) coordinates of the digitized points (pixels) along the contour. A contour with N points could be represented by the series of coordinates $\{x(n), y(n)\}$, $n = 0, 1, 2, \dots, N - 1$. Observe that there is no gray level associated with the pixels along a contour. A contour may be depicted as a binary or bilevel image.

6.1.1 Signatures of contours

The dimensionality of representation of a contour may be reduced from two to one by converting from a coordinate-based representation to distances from each contour point to a reference point. A convenient reference is the centroid or center of mass of the contour, whose coordinates are given by

$$\bar{x} = \frac{1}{N} \sum_{n=0}^{N-1} x(n), \quad \text{and} \quad \bar{y} = \frac{1}{N} \sum_{n=0}^{N-1} y(n). \quad (6.1)$$

The signature of the contour is then defined as

$$d(n) = \sqrt{[x(n) - \bar{x}]^2 + [y(n) - \bar{y}]^2}, \quad (6.2)$$

$n = 0, 1, 2, \dots, N - 1$; see Figure 6.1. It should be noted that the centroids of regions that are concave or have holes could lie outside the regions.

A radial-distance signature may also be derived by computing the distance from the centroid to the contour point(s) intersected for angles of the radial line spanning the range ($0^\circ, 360^\circ$). However, for irregular contours, such a signature may be multivalued for some angles; that is, a radial line may intersect the contour more than once (see, for example, Pohlman et al. [407]).

It is obvious that going around a contour more than once generates the same signature; hence, the signature signal is periodic with the period equal to N , the number of pixels on the contour. The signature of a contour provides general information on the nature of the contour, such as its smoothness or roughness.

Examples: Figures 6.2 (a) and 6.3 (a) show the contours of a benign breast mass and a malignant tumor, respectively, as observed on mammograms [345]. The ‘*’ marks within the contours represent their centroids. Figures 6.2 (b) and 6.3 (b) show the signatures of the contours as defined in Equation 6.2. It is evident that the smooth contour of the benign mass possesses a smooth signature, whereas the spiculated malignant tumor has a rough signature with several significant rapid variations over its period.

6.1.2 Chain coding

An efficient representation of a contour may be achieved by specifying the (x, y) coordinates of an arbitrary starting point on the contour, the direction of traversal (clockwise or counter-clockwise), and a code to indicate the manner of movement to reach the next contour point on a discrete grid. A coarse representation may be achieved by using only four possible movements: to the point at the left of, right of, above, or below the current point, as indicated in Figure 6.4 (a). A finer representation may be achieved by using eight possible movements, including diagonal movements, as indicated in Figure 6.4 (b). The sequence of codes required to traverse through all the points along the

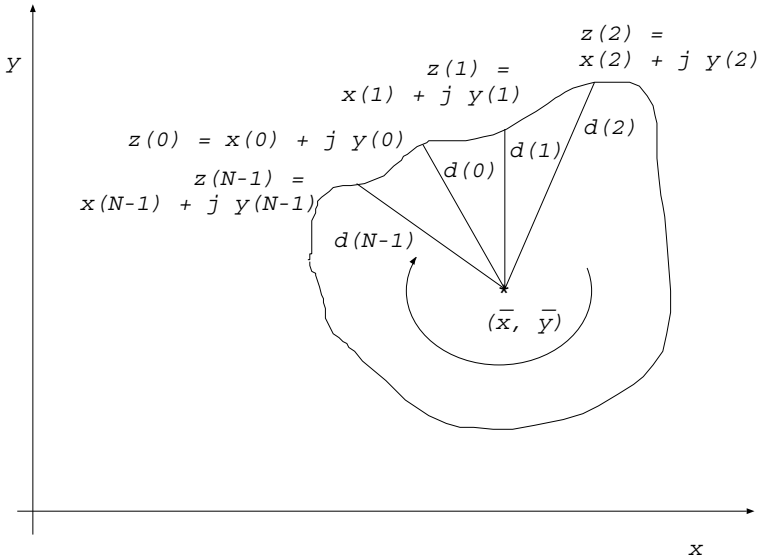


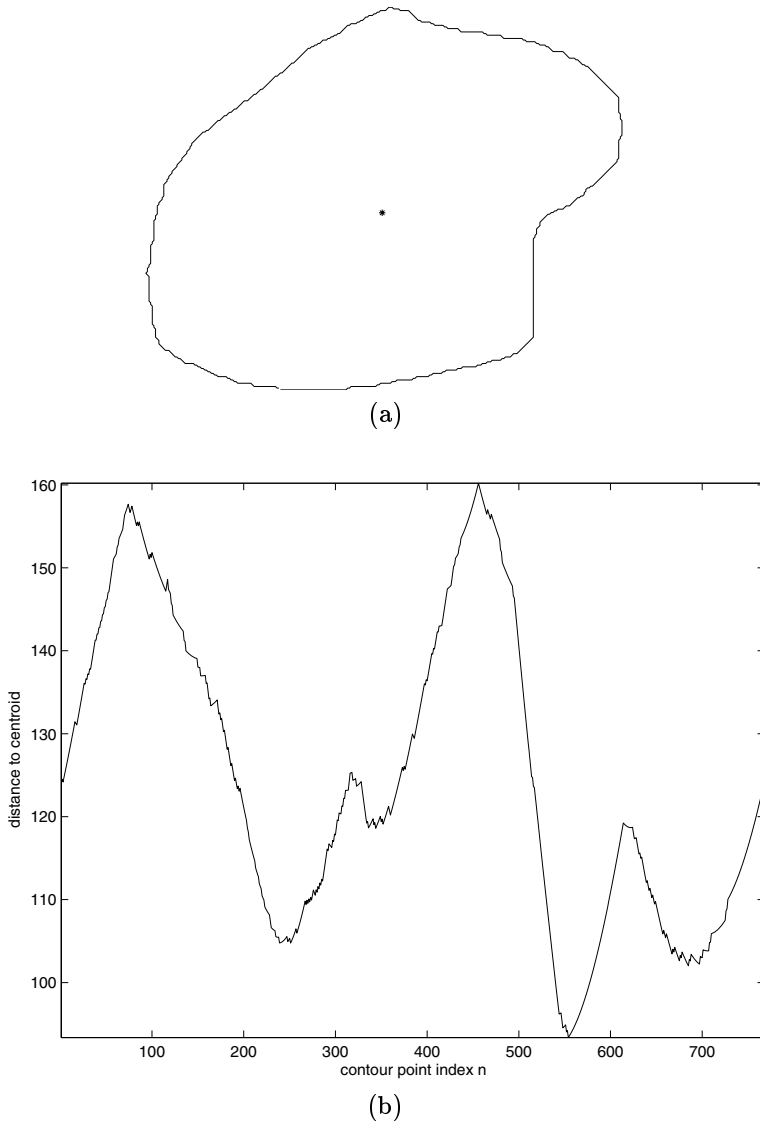
FIGURE 6.1

A contour represented by its boundary points $z(n)$ and distances $d(n)$ to its centroid.

contour is known as the chain code [8, 245]. The technique was proposed by Freeman [408], and is also known as the Freeman chain code.

The chain code facilitates more compact representation of a contour than the direct specification of the (x, y) coordinates of all of its points. Except the initial point, the representation of each point on the contour requires only two or three bits, depending upon the type of code used. Furthermore, chain coding provides the following advantages:

- The code is invariant to shift or translation because the starting point is kept out of the code.
- To a certain extent, the chain code is invariant to size (scaling). Contours of different sizes may be generated from the same code by using different sampling grids (step sizes). A contour may also be enlarged by a factor of n by repeating each code element n times and maintaining the same sampling grid [408]. A contour may be shrunk to half of the original size by reducing pairs of code elements to single numbers, with approximation of unequal pairs by their averages reduced to integers.
- The chain code may be normalized for rotation by taking the first difference of the code (and adding 4 or 8 to negative differences, depending upon the code used).

**FIGURE 6.2**

(a) Contour of a benign breast mass; $N = 768$. The '*' mark represents the centroid of the contour. (b) Signature $d(n)$ as defined in Equation 6.2.





FIGURE 6.3

(a) Contour of a malignant breast tumor; $N = 3,281$. The ‘*’ mark represents the centroid of the contour. (b) Signature $d(n)$ as defined in Equation 6.2.

- With reference to the 8-symbol code, the rotation of a given contour by $n \times 90^\circ$ in the counter-clockwise direction may be achieved by adding a value of $2n$ to each code element, followed by integer division by 8. The addition of an odd number rotates the contour by the corresponding multiple of 45° ; however, the rotation of a contour by angles other than integral multiples of 90° on a discrete grid is subject to approximation.
- In the case of the 8-symbol code, the length of a contour is given by the number of even codes plus $\sqrt{2}$ times the number of odd codes, multiplied by the grid sampling interval.
- The chain code may also be used to achieve reduction, check for closure, check for multiple loops, and determine the area of a closed loop [408].

Examples: Figure 6.5 shows a contour represented using the chain codes with four and eight symbols. The use of a discrete grid with large spacings leads to the loss of fine detail in the contour. However, this feature may be used advantageously to filter out minor irregularities due to noise, artifacts due to drawing by hand, etc.





FIGURE 6.4

Chain code with (a) four directional codes and (b) eight directional codes.

6.1.3 Segmentation of contours

The segmentation of a contour into a set of piecewise-continuous curves is a useful step before analysis and modeling. Segmentation may be performed by locating the points of inflection on the contour.

Consider a function $f(x)$. Let $f'(x)$, $f''(x)$, and $f'''(x)$ represent the first, second, and third derivatives of $f(x)$. A point of inflection of the function or



Chain code: [0 1 0 3 3 0 0 3 2 3 2 3 3 2 1 2 2 3 2 1 1 1 2 1 0 1 1 0 0 3]

Figure 6.5 (a)

curve $f(x)$ is defined as a point where $f''(x)$ changes its sign. Note that the derivation of $f''(x)$ requires $f(x)$ and $f'(x)$ to be continuous and differentiable. It follows that the following conditions apply at a point of inflection:


$$\begin{aligned}
 f''(x) &= 0, \\
 f'(x) &\neq 0, \\
 f'(x) f''(x) &= 0, \text{ and} \\
 f'(x) f'''(x) &\neq 0.
 \end{aligned}
 \tag{6.3}$$

Let $\mathbf{C} = \{(x(n), y(n))\}$, $n = 0, 1, 2, \dots, N - 1$, represent in vector form the (x, y) coordinates of the N points on the given contour. The points of inflection on the contour are obtained by solving

$$\begin{aligned}
 \mathbf{C}' \times \mathbf{C}'' &= 0, \\
 \mathbf{C}' \times \mathbf{C}''' &\neq 0,
 \end{aligned}
 \tag{6.4}$$

where \mathbf{C}' , \mathbf{C}'' , and \mathbf{C}''' are the first, second, and third derivatives of \mathbf{C} , respectively, and \times represents the vector cross product. Solving Equation 6.4 is equivalent to solving the system of equations given by

$$x''(n) y'(n) - x'(n) y''(n) = 0,$$



Chain code: [0 1 6 6 0 7 4 5 6 6 3 4 4 5 2 2 2 3 1 1 7]

(b)

FIGURE 6.5

A closed contour represented using the chain code (a) using four directional codes as in Figure 6.4 (a), and (b) with eight directional codes as in Figure 6.4 (b). The 'o' mark represents the starting point of the contour, which is traversed in the clockwise direction to derive the code.

$$x'(n) y'''(n) - x'''(n) y'(n) \neq 0, \quad (6.5)$$

where $x'(n)$, $y'(n)$, $x''(n)$, $y''(n)$, $x'''(n)$, and $y'''(n)$ are the first, second, and third derivatives of $x(n)$ and $y(n)$, respectively.

Segments of contours of breast masses between successive points of inflection were modeled as parabolas by Menut et al. [354]. Difficulty lies in segmentation because the contours of masses are, in general, not smooth. False or irrelevant points of inflection could appear on relatively straight parts of a contour when $x''(n)$ and $y''(n)$ are not far from zero. In order to address this problem, smoothed derivatives at each contour point could be estimated by considering the cumulative sum of weighted differences of a certain number of pairs of points on either side of the point $x(n)$ under consideration as

$$x'(n) = \sum_{i=1}^m \frac{[x(n+i) - x(n-i)]}{i}, \quad (6.6)$$

where m represents the number of pairs of points used to compute the derivative $x'(n)$; the same procedure applies to the computation of $y'(n)$.


In the works reported by Menut et al. [354] and Rangayyan et al. [345], the value of m was varied from 3 to 60 to compute derivatives that resulted in varying numbers of inflection points for a given contour. The number of inflection points detected as a function of the number of differences used was analyzed to determine the optimal number of differences that would provide the most appropriate inflection points: the value of m at the first straight segment on the function was selected.

Examples: Figure 6.6 shows the contour of a spiculated malignant tumor. The points of inflection detected are marked with '*'. The number of inflection points detected is plotted in Figure 6.7 as a function of the number of differences used (m in Equation 6.6); the horizontal and vertical lines indicate the optimal number of differences used to compute the derivative at each contour point and the corresponding number of points of inflection that were located on the contour.

The contour in Figure 6.6 is shown in Figure 6.8, overlaid on the corresponding part of the original mammogram. Segments of the contours are shown in black or white, indicating if they are concave or convex, respectively. Figure 6.9 provides a similar illustration for a circumscribed benign mass. Analysis of concavity of contours is described in Section 6.4.

6.1.4 Polygonal modeling of contours

Pavlidis and Horowitz [361] and Pavlidis and Ali [409] proposed methods for segmentation and approximation of curves and shapes by polygons for computer recognition of handwritten numerals, cell outlines, and ECG signals. Ventura and Chen [410] presented an algorithm for segmenting and polygonal modeling of 2D curves in which the number of segments is to be prespecified

**FIGURE 6.6**

Contour of a spiculated malignant tumor with the points of inflection indicated by '*'. Number of points of inflection = 58. See also Figure 6.8.




FIGURE 6.7

Number of inflection points detected as a function of the number of differences used to estimate the derivative for the contour in Figure 6.6. The horizontal and vertical lines indicate the optimal number of differences used to compute the derivative at each contour point and the corresponding number of points of inflection that were located on the contour.




FIGURE 6.8

Concave and convex parts of the contour of a spiculated malignant tumor, separated by the points of inflection. See also Figure 6.6. The concave parts are shown in black and the convex parts in white. The image size is 770×600 pixels or 37.2×47.7 mm with a pixel size of $62 \mu\text{m}$. Shape factors $f_{cc} = 0.47$, $SI = 0.62$, $cf = 0.94$. Reproduced with permission from R.M. Rangayyan, N.R. Mudigonda, and J.E.L. Desautels, "Boundary modeling and shape analysis methods for classification of mammographic masses", *Medical and Biological Engineering and Computing*, 38: 487 – 496, 2000. © IFMBE.




FIGURE 6.9

Concave and convex parts of the contour of a circumscribed benign mass, separated by the points of inflection. The concave parts are shown in black and the convex parts in white. The image size is 730×630 pixels or 31.5×36.5 mm with a pixel size of $50 \mu\text{m}$. Shape factors $f_{cc} = 0.16$, $SI = 0.22$, $cf = 0.30$. Reproduced with permission from R.M. Rangayyan, N.R. Mudigonda, and J.E.L. Desautels, "Boundary modeling and shape analysis methods for classification of mammographic masses", *Medical and Biological Engineering and Computing*, 38: 487 – 496, 2000. © IFMBE.

for initiating the process, in relation to the complexity of the shape. This is not a desirable step when dealing with complex or spiculated shapes of breast tumors [163]. In a modified approach proposed by Rangayyan et al. [345], the polygon formed by the points of inflection detected on the original contour was used as the initial input to the polygonal modeling procedure. This step helps in automating the polygonalization algorithm: the method does not require any interaction from the user in terms of the initial number of segments.

Given an irregular contour \mathbf{C} as specified by the set of its (x, y) coordinates, the polygonal modeling algorithm starts by dividing the contour into a set of piecewise-continuous curved parts by locating the points of inflection on the contour as explained in Section 6.1.3. Each segmented curved part is represented by a pair of linear segments based on its arc-to-chord deviation. The procedure is iterated subject to predefined boundary conditions so as to minimize the error between the true length of the contour and the cumulative length computed from the polygonal segments.

Let $\mathbf{C} = \{x(n), y(n)\}$, $n = 0, 1, 2, \dots, N - 1$, represent the given contour. Let SC_{mk} , $SC_{mk} \in \mathbf{C}$, $m = 1, 2, \dots, M$, be M curved parts, each containing a set of contour points, at the start of the k^{th} iteration, such that $SC_{1k} \cup SC_{2k} \cup \dots \cup SC_{Mk} \equiv \mathbf{C}$. The iterative procedure proposed by Rangayyan et al. [345] is as follows:

1. In each curved part represented by SC_{mk} , the arc-to-chord distance is computed for all the points, and the point on the curve with the maximum arc-to-chord deviation (d_{\max}) is located.
2. If $d_{\max} \geq 0.25 \text{ mm}$ (5 pixels in the images with a pixel size of $50 \mu\text{m}$ used in the work of Rangayyan et al. [345]), the curved part is segmented at the point of maximum deviation to approximate the same with a pair of linear segments, irrespective of the length of the resulting linear segments. If $0.1 \text{ mm} \leq d_{\max} < 0.25 \text{ mm}$, the curved part is segmented at the point of maximum deviation subject to the condition that the resulting linear segments satisfy a minimum-length criterion, which was specified as 1 mm in the work of Rangayyan et al. [345]. If $d_{\max} < 0.1 \text{ mm}$, the curved part SC_{mk} is considered to be almost linear and is not segmented any further.
3. After performing Steps 1 and 2 on all the curved parts of the contour available in the current k^{th} iteration, the resulting vector of the polygon's vertices is updated.
4. If the number of polygonal segments following the k^{th} iteration equals that of the previous iteration, the algorithm is considered to have converged and the polygonalization process is terminated. Otherwise, the procedure (Steps 1 to 3) is repeated until the algorithm converges.

The criterion for choosing the threshold for arc-to-chord deviation was based on the assumption that any segment possessing a smaller deviation is insignificant in the analysis of contours of breast masses.

Examples: Figure 6.10 (a) shows the points of inflection (denoted by “*”) and the initial stage of polygonal modeling (straight-line segments) of the contour of a spiculated malignant tumor (see also Figure 6.8). Figure 6.10 (b) shows the final result of polygonal modeling of the same contour. The algorithm converged after four iterations, as shown by the convergence plot in Figure 6.11. The result of the application of the polygonal modeling algorithm to the contour of a circumscribed benign mass is shown in Figure 6.12.

The number of linear segments required for the approximation of a contour increases with its shape complexity; polygons with the number of sides in the range 20 – 400 were used in the work of Rangayyan et al. [345]) to model contours of breast masses and tumors. The number of iterations required for the convergence of the algorithm did not vary much for different mass contour shapes, remaining within the range 3 – 5. This is due to the fact that the relative complexity of the contour to be segmented is taken into consideration during the initial preprocessing step of locating the points of inflection; hence, the subsequent polygonalization process is robust and computationally efficient. The algorithm performed well and delivered satisfactory results on various irregular shapes of spiculated cases of benign and malignant masses.

6.1.5 Parabolic modeling of contours

Menut et al. [354] proposed the modeling of segments of contours of breast masses between successive points of inflection as parabolas. An inspection of the segments of the contours illustrated in Figures 6.6 and 6.12 (a) (see also Figures 6.8 and 6.9) indicates that most of the curved portions between successive points inflection lend themselves well to modeling as parabolas. Some of the segments are relatively straight; however, such segments may not contribute much to the task of discrimination between benign masses and malignant tumors.

Let us consider a segment of a contour represented in the continuous 2D space by the points $[x(s), y(s)]$ over the interval $S_1 \leq s \leq S_2$, where s indicates distance along the contour and S_1 and S_2 are the end-points of the segment. Let us now consider the approximation of the curve by a parabola. Regardless of the position and orientation of the given curve, let us consider the simplest representation of a parabola as $Y = A X^2$ in the coordinate space (X, Y) . The parameter A controls the narrowness of the parabola: the larger the value of A , the narrower is the parabola. Allowing for a rotation of θ and a shift of (c, d) between the (x, y) and (X, Y) spaces, we have

$$\begin{aligned} x(s) &= X(s) \cos \theta - Y(s) \sin \theta + c, \\ y(s) &= X(s) \sin \theta + Y(s) \cos \theta + d. \end{aligned} \tag{6.7}$$




FIGURE 6.10

Polygonal modeling of the contour of a spiculated malignant tumor. (a) Points of inflection (indicated by ‘*’) and the initial polygonal approximation (straight-line segments); number of sides = 58. (b) Final model after four iterations; number of sides = 146. See also Figure 6.8. Reproduced with permission from R.M. Rangayyan, N.R. Mudigonda, and J.E.L. Desautels, “Boundary modeling and shape analysis methods for classification of mammographic masses”, *Medical and Biological Engineering and Computing*, 38: 487 – 496, 2000. © IFMBE.




FIGURE 6.11

Convergence plot of the iterative polygonal modeling procedure for the contour of the spiculated malignant tumor in Figure 6.10. Reproduced with permission from R.M. Rangayyan, N.R. Mudigonda, and J.E.L. Desautels, "Boundary modeling and shape analysis methods for classification of mammographic masses", *Medical and Biological Engineering and Computing*, 38: 487 – 496, 2000. © IFMBE.




FIGURE 6.12

Polygonal modeling of the contour of a circumscribed benign mass. (a) Points of inflection (indicated by ‘*’) and the initial polygonal approximation (straight-line segments); initial number of sides = 14. (b) Final model; number of sides = 36, number of iterations = 4. See also Figure 6.9. Reproduced with permission from R.M. Rangayyan, N.R. Mudigonda, and J.E.L. Desautels, “Boundary modeling and shape analysis methods for classification of mammographic masses”, *Medical and Biological Engineering and Computing*, 38: 487 – 496, 2000. © IFMBE.

We also have the following relationships:

$$\begin{aligned} X(s) &= [x(s) - c] \cos \theta + [y(s) - d] \sin \theta, \\ Y(s) &= -[x(s) - c] \sin \theta + [y(s) - d] \cos \theta; \end{aligned} \quad (6.8)$$

$$\begin{aligned} X(s) &= s, \\ Y(s) &= A s^2. \end{aligned} \quad (6.9)$$

Taking the derivatives of Equation 6.9 with respect to s , we get the following:

$$\begin{aligned} X'(s) &= 1, \\ Y'(s) &= 2As; \end{aligned} \quad (6.10)$$

$$\begin{aligned} X''(s) &= 0, \\ Y''(s) &= 2A. \end{aligned} \quad (6.11)$$

Similarly, taking the derivatives of Equation 6.8 with respect to s , we get the following:

$$\begin{aligned} X''(s) &= x''(s) \cos \theta + y''(s) \sin \theta, \\ Y''(s) &= -x''(s) \sin \theta + y''(s) \cos \theta. \end{aligned} \quad (6.12)$$

Combining Equations 6.11 and 6.12, we get

$$X''(s) = 0 = x''(s) \cos \theta + y''(s) \sin \theta, \quad (6.13)$$

which, upon multiplication with $\sin \theta$, yields

$$x''(s) \sin \theta \cos \theta + y''(s) \sin \theta \sin \theta = 0. \quad (6.14)$$

Similarly, we also get

$$Y''(s) = 2A = -x''(s) \sin \theta + y''(s) \cos \theta, \quad (6.15)$$

which, upon multiplication with $\cos \theta$, yields

$$2A \cos \theta = -x''(s) \sin \theta \cos \theta + y''(s) \cos \theta \cos \theta. \quad (6.16)$$

Combining Equations 6.14 and 6.16 we get

$$2A \cos \theta = y''(s). \quad (6.17)$$

The equations above indicate that $y''(s)$ and $x''(s)$ are constants with values related to A and θ . The values of the two derivatives may be computed from the given curve over all available points, and averaged to obtain the corresponding (constant) values. Equations 6.14 and 6.17 may then be solved

simultaneously to obtain θ and A . Thus, the parameter of the parabolic model is obtained from the given contour segment.

Menut et al. hypothesized that malignant tumors, due to the presence of narrow spicules or microlobulations, would have several parabolic segments with large values of A ; on the other hand, benign masses, due to their characteristics of being oval or macrolobulated, would have a small number of parabolic segments with small values of A . The same reasons were also expected to lead to a larger standard deviation of A for malignant tumors than for benign masses. In addition to the parameter A , Menut et al. proposed to use the width of the projection of each parabola on to the X axis, with the expectation that its values would be smaller for malignant tumors than for benign masses. A classification accuracy of 76% was obtained with a set of 54 contours.

6.1.6 Thinning and skeletonization

Objects that are linear or oblong, or structures that have branching (anastomotic) patterns may be effectively characterized by their skeletons. The skeleton of an object or region is obtained by its medial-axis transform or via a thinning algorithm [8, 245, 411].

The medial-axis transformation proposed by Blum [412] is as follows. First, the given image needs to be binarized so as to include only the patterns of interest. Let the set of pixels in the binary pattern be denoted as B , let C be the set of contour pixels of B , and let c_i be an arbitrary contour point in C . For each point b in B , a point c_i is found such that the distance between the point b and c_i , represented as $d(b, c_i)$, is at its minimum. If a second point c_k is found in C such that $d(b, c_k) = d(b, c_i)$, then b is a part of the skeleton of B ; otherwise, b is not a part of the skeleton.

A simple algorithm for thinning is as follows [8, 411, 413]. Assume that the image has been binarized, with the pixels inside the ROIs being labeled as 1 and the background pixels as 0. A contour point is defined as any pixel having the value 1 and at least one 8-connected neighbor valued 0. Let the 8-connected neighboring pixels of the pixel p_1 being processed be indexed as

$$\begin{bmatrix} p_9 & p_2 & p_3 \\ p_8 & p_1 & p_4 \\ p_7 & p_6 & p_5 \end{bmatrix}. \quad (6.18)$$

1. Flag a contour point p_1 for deletion if the following are true:

- (a) $2 \leq N(p_1) \leq 6$;
- (b) $S(p_1) = 1$;
- (c) $p_2 \times p_4 \times p_6 = 0$;
- (d) $p_4 \times p_6 \times p_8 = 0$;

where $N(p_1)$ is the number of nonzero neighbors of p_1 , and $S(p_1)$ is the number of 0 – 1 transitions in the sequence $p_2, p_3, \dots, p_9, p_2$.

2. Delete all flagged pixels.
3. Do the same as Step 1 above replacing the conditions (c) and (d) with
(c') $p_2 \times p_4 \times p_8 = 0$;
(d') $p_2 \times p_6 \times p_8 = 0$.
4. Delete all flagged pixels.
5. Iterate Steps 1 – 4 until no further pixels are deleted.

The algorithm described above has the properties that it does not remove end points, does not break connectivity, and does not cause excessive erosion of the region [8].

Example: Figure 6.13 (a) shows a pattern of blood vessels in a section of a ligament [414, 415]. The vessels were perfused with black ink prior to extraction of the tissue for study. Figure 6.13 (b) shows the skeleton of the image in part (a) of the figure. It is seen that the skeleton represents the general orientational pattern and overall shape of the blood vessels in the original image. However, information regarding the variation in the thickness (diameter) of the blood vessels is lost in skeletonization. Eng et al. [414] studied the effect of injury and healing on the microvascular structure of ligaments by analyzing the statistics of the volume and directional distribution of blood vessels as illustrated in Figure 6.13; see Section 8.7.2 for details.

6.2 Shape Factors

Although contours may be effectively characterized by representations and models such as the chain code and the polygonal model described in the preceding section, it is often desirable to encode the nature or form of a contour using a small number of measures, commonly referred to as shape factors. The nature of the contour to be encapsulated in the measures may vary from one application to another. Regardless of the application, a few basic properties are essential for efficient representation, of which the most important are:

- invariance to shift in spatial position,
- invariance to rotation, and
- invariance to scaling (enlargement or reduction).

Invariance to reflection may also be desirable in some applications. Shape factors that meet the criteria listed above can effectively and efficiently represent contours for pattern classification.




FIGURE 6.13

(a) Binarized image of blood vessels in a ligament perfused with black ink. Image courtesy of R.C. Bray and M.R. Doschak, University of Calgary. (b) Skeleton of the image in (a) after 15 iterations of the algorithm described in Section 6.1.6.

A basic method that is commonly used to represent shape is to fit an ellipse or a rectangle to the given (closed) contour. The ratio of the major axis of the ellipse to its minor axis (or, equivalently, the ratio of the larger side to the smaller side of the bounding rectangle) is known as its *eccentricity*, and represents its deviation from a circle (for which the ratio will be equal to unity). Such a measure, however, represents only the elongation of the object, and may have, on its own, limited application in practice. Several shape factors of increasing complexity and specificity of application are described in the following sections.

6.2.1 Compactness

Compactness is a simple and popular measure of the efficiency of a contour to contain a given area, and is commonly defined as

$$Co = \frac{P^2}{A}, \quad (6.19)$$

where P and A are the contour perimeter and area enclosed, respectively. The smaller the area contained by a contour of a given length, the larger will be the value of compactness. Compactness, as defined in Equation 6.19, has a lower bound of 4π for a circle (except for the trivial case of zero for $P = 0$), but no upper bound. It is evident that compactness is invariant to shift, scaling, rotation, and reflection of a contour.

In order to restrict and normalize the range of the parameter to $[0, 1]$, as well as to obtain increasing values with increase in complexity of the shape, the definition of compactness may be modified as [274, 163]


$$cf = 1 - \frac{4\pi A}{P^2}. \quad (6.20)$$

With this expression, cf has a lower bound of zero for a circle, and increases with the complexity of the contour to a maximum value of unity.

Examples: Figure 6.14 illustrates a few simple geometric shapes along with their values of compactness. Elongated contours with large values of the perimeter and small enclosed areas possess high values of compactness.

Figure 6.15 illustrates a few objects with simple geometric shapes including scaling and rotation; the values of compactness Co and cf for the contours of the objects are listed in Table 6.1 [274, 320, 334, 416]. It is evident that both definitions of compactness provide the desired invariance to scaling and rotation (within the limitations due to the use of a discrete grid).

Figure 6.16 illustrates a few objects of varying shape complexity, prepared by cutting construction paper [215, 320, 334, 416]. The values of compactness cf for the contours of the objects are listed in Table 6.2. It is seen that compactness increases with shape roughness and/or complexity.

**FIGURE 6.14**

Examples of contours with their values of compactness Co and cf , as defined in Equations 6.19 and 6.20. (a) Circle. (b) Square. (c) Rectangle with sides equal to 1.0 and 0.5 units. (d) Rectangle with sides equal to 1.0 and 0.25 units. (e) Right-angled triangle of height 1.0 and base 0.25 units.




FIGURE 6.15

A set of simple geometric shapes, including scaling and rotation, created on a discrete grid, to test shape factors. Reproduced with permission from L. Shen, R.M. Rangayyan, and J.E.L. Desautels, "Application of shape analysis to mammographic calcifications", *IEEE Transactions on Medical Imaging*, 13(2): 263 - 274, 1994. © IEEE.

TABLE 6.1

Shape Factors for the Shapes in Figure 6.15 [274, 320, 334, 416].

Shape	Co	cf	$F_1 = F'_1$	F_2	F'_2	F_3	F'_3	mf	ff
Large circle (a)	14.08	0.1078	0.0056	0.4105	0.0042	2.0271	0.0067	0.0011	0.0358
Medium circle (b)	14.13	0.1105	0.0066	0.1731	0.0037	1.9285	0.0078	0.0012	0.0380
Small circle (c)	14.29	0.1205	0.0085	0.1771	0.0048	1.9334	0.0100	0.0015	0.0432
Large square (d)	15.77	0.2034	0.1083	0.5183	0.0870	1.9987	0.1288	0.0205	0.1416
Medium square (e)	15.70	0.1997	0.1081	0.5122	0.0865	2.0126	0.1287	0.0207	0.1389
Rotated square (f)	16.00	0.2146	0.1101	0.5326	0.0893	1.9987	0.1309	0.0208	0.1434
Small square (g)	15.56	0.1926	0.1078	0.4943	0.0853	2.0495	0.1290	0.0212	0.1362
Large rectangle (i)	17.60	0.2858	0.2491	-0.3313	-0.1724	1.5385	0.2775	0.0283	0.1494
Medium rectangle (j)	17.47	0.2807	0.2483	-0.3267	-0.1710	1.5429	0.2767	0.0284	0.1483
Rotated rectangle (k)	17.47	0.2807	0.2483	-0.3267	-0.1710	1.5429	0.2767	0.0284	0.1483
Small rectangle (l)	17.23	0.2707	0.2468	-0.3165	-0.1682	1.5583	0.2758	0.0290	0.1420
Large isosceles triangle (m)	22.41	0.4392	0.3119	0.0737	0.1308	2.3108	0.3846	0.0727	0.2248
Medium isosceles triangle (n)	22.13	0.4322	0.3051	0.2027	0.1792	2.2647	0.3743	0.0692	0.2233
Rotated isosceles triangle (o)	22.13	0.4322	0.3051	0.2027	0.1792	2.2647	0.3743	0.0692	0.2238
Small isosceles triangle (p)	21.61	0.4185	0.3014	0.1518	0.1608	2.2880	0.3707	0.0693	0.2198
Large right-angled triangle (q)	27.68	0.5459	0.3739	0.0475	0.1355	1.9292	0.4407	0.0668	0.2217
Medium right-angled triangle (r)	27.18	0.5377	0.3707	0.0534	0.1396	1.9433	0.4377	0.0670	0.2221
Rotated right-angled triangle (s)	26.99	0.5345	0.3752	0.0022	0.0487	1.8033	0.4347	0.0596	0.2216
Small right-angled triangle (t)	26.26	0.5215	0.3644	0.0646	0.1462	1.9750	0.4319	0.0676	0.2180

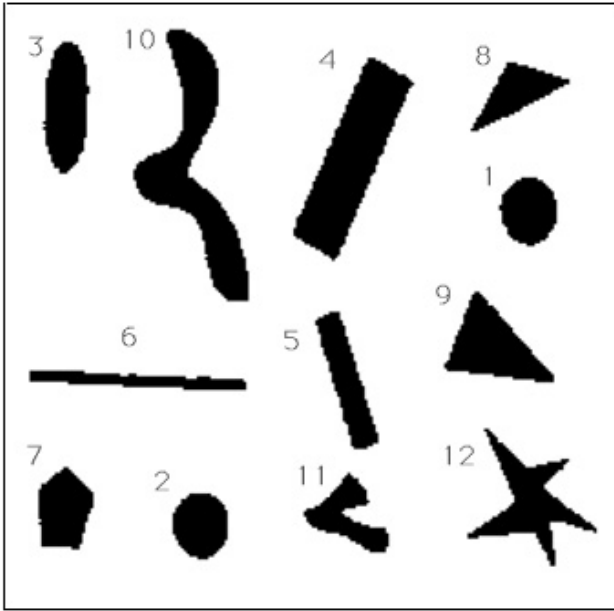


FIGURE 6.16

A set of objects of varying shape complexity. The objects were prepared by cutting construction paper. The contours of the objects include imperfections and artifacts. Reproduced with permission from L. Shen, R.M. Rangayyan, and J.E.L. Desautels, “Application of shape analysis to mammographic calcifications”, *IEEE Transactions on Medical Imaging*, 13(2): 263 – 274, 1994. © IEEE.

Shen et al. [274, 334] applied compactness to shape analysis of mammographic calcifications. The details of this application are presented in Sections 6.6 and 12.7. The use of compactness in benign-versus-malignant classification of breast masses is discussed in Sections 6.7, 12.11, and 12.12.

6.2.2 Moments

Statistical moments of PDFs and other data distributions have been utilized as pattern features in a number of applications; the same concepts have been extended to the analysis of images and contours [8, 417, 418, 419, 420]. Given a 2D continuous image $f(x, y)$, the regular moments m_{pq} of order $(p + q)$ are

TABLE 6.2
 Shape Factors for the Objects in
 Figure 6.16 Arranged in Increasing Order
 of ff [334, 274, 416, 320].

Shape	cf	mf	ff	Type
1	0.13	0.022	0.14	Circle
2	0.11	0.019	0.14	Circle
3	0.35	0.047	0.14	Ellipse
4	0.55	0.047	0.17	Rectangle
5	0.62	0.060	0.18	Rectangle
6	0.83	0.084	0.18	Rectangle
7	0.22	0.038	0.19	Pentagon
8	0.50	0.063	0.24	Triangle
9	0.44	0.063	0.25	Triangle
10	0.75	0.090	0.30	Other
11	0.63	0.106	0.36	Other
12	0.81	0.077	0.42	Other

defined as [420, 8]:

$$m_{pq} = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} x^p y^q f(x, y) dx dy, \quad (6.21)$$

for $p, q = 0, 1, 2, \dots$. A uniqueness theorem [128] states that if $f(x, y)$ is piecewise continuous and has nonzero values only in a finite part of the (x, y) plane, then moments of all orders exist, and the moment sequence m_{pq} , $p, q = 0, 1, 2, \dots$, is uniquely determined by $f(x, y)$. Conversely, the sequence m_{pq} uniquely determines $f(x, y)$ [8].

The central moments are defined with respect to the centroid of the image as

$$\mu_{pq} = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} (x - \bar{x})^p (y - \bar{y})^q f(x, y) dx dy, \quad (6.22)$$

where

$$\bar{x} = \frac{m_{10}}{m_{00}}, \quad \bar{y} = \frac{m_{01}}{m_{00}}. \quad (6.23)$$

Observe that the gray levels of the pixels provide weights for the moments as defined above. If moments are to be computed for a contour, only the contour pixels would be used with weights equal to unity; the internal pixels would have weights of zero, and effectively do not participate in the computation of the moments.

For an $M \times N$ digital image, the integrals are replaced by summations; for example, Equation 6.22 becomes

$$\mu_{pq} = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} (m - \bar{x})^p (n - \bar{y})^q f(m, n). \quad (6.24)$$

The central moments have the following relationships [8]:

$$\mu_{00} = m_{00} = \mu, \tag{6.25}$$

$$\mu_{10} = \mu_{01} = 0, \tag{6.26}$$

$$\mu_{20} = m_{20} - \mu\bar{x}^2, \tag{6.27}$$

$$\mu_{11} = m_{11} - \mu\bar{x}\bar{y}, \tag{6.28}$$

$$\mu_{02} = m_{02} - \mu\bar{y}^2, \tag{6.29}$$

$$\mu_{30} = m_{30} - 3m_{20}\bar{x} + 2\mu\bar{x}^3, \tag{6.30}$$

$$\mu_{21} = m_{21} - m_{20}\bar{y} - 2m_{11}\bar{x} + 2\mu\bar{x}^2\bar{y}, \tag{6.31}$$

$$\mu_{12} = m_{12} - m_{02}\bar{x} - 2m_{11}\bar{y} + 2\mu\bar{x}\bar{y}^2, \tag{6.32}$$

$$\mu_{03} = m_{03} - 3m_{02}\bar{y} + 2\mu\bar{y}^3. \tag{6.33}$$

Normalization with respect to size is achieved by dividing each of the moments by μ_{00}^γ , where $\gamma = \frac{p+q}{2} + 1$, to obtain the normalized moments as [8]

$$\nu_{pq} = \frac{\mu_{pq}}{\mu_{00}^\gamma}. \tag{6.34}$$

Hu [420] (see also Gonzalez and Woods [8]) defined a set of seven shape factors that are functions of the second-order and third-order central moments as follows:

$$M_1 = \nu_{20} + \nu_{02}, \tag{6.35}$$

$$M_2 = (\nu_{20} - \nu_{02})^2 + 4\nu_{11}^2, \tag{6.36}$$

$$M_3 = (\nu_{30} - 3\nu_{12})^2 + (3\nu_{21} - \nu_{03})^2, \tag{6.37}$$

$$M_4 = (\nu_{30} + \nu_{12})^2 + (\nu_{21} + \nu_{03})^2, \tag{6.38}$$

$$M_5 = (\nu_{30} - 3\nu_{12})(\nu_{30} + \nu_{12})[(\nu_{30} + \nu_{12})^2 - 3(\nu_{21} + \nu_{03})^2] + (3\nu_{21} - \nu_{03})(\nu_{21} + \nu_{03})[3(\nu_{30} + \nu_{12})^2 - (\nu_{21} + \nu_{03})^2], \tag{6.39}$$

$$M_6 = (\nu_{20} - \nu_{02})[(\nu_{30} + \nu_{12})^2 - (\nu_{21} + \nu_{03})^2] + 4\nu_{11}(\nu_{30} + \nu_{12})(\nu_{21} + \nu_{03}), \tag{6.40}$$

$$M_7 = (3\nu_{21} - \nu_{03})(\nu_{30} + \nu_{12})[(\nu_{30} + \nu_{12})^2 - 3(\nu_{21} + \nu_{03})^2] - (\nu_{30} - 3\nu_{12})(\nu_{21} + \nu_{03})[3(\nu_{30} + \nu_{12})^2 - (\nu_{21} + \nu_{03})^2]. \tag{6.41}$$

The shape factors M_1 through M_7 are invariant to shift, scaling, and rotation (within limits imposed by representation on a discrete grid), and have

been found to be useful for pattern analysis. Rangayyan et al. [163] computed several versions of the factors M_1 through M_7 for 54 breast masses and tumors, using the mass ROIs with and without their gray levels, as well as the contours of the masses with and without their gray levels. The features provided benign-versus-malignant classification accuracies in the range 56 – 75%; see Section 6.7.

Moments of distances to the centroid: When an ROI is represented using only its contour, an alternative definition of moments is based upon a sequence that represents the Euclidean distances between the centroid of the region and all of the points or pixels along the contour, shown as $d(n)$ in Figure 6.1. The distances from the center of a circle to its contour points are all equal to the radius of the circle; the variance of the values is zero. On the other hand, for rough shapes, the distances will vary considerably; see Figures 6.2 and 6.3 for examples. The variance and higher-order moments of the distance values could be expected to provide indicators of shape complexity. The p^{th} moment of the sequence $d(n)$ is defined as [417]

$$m_p = \frac{1}{N} \sum_{n=0}^{N-1} [d(n)]^p, \quad (6.42)$$

and the p^{th} central moment is defined as

$$M_p = \frac{1}{N} \sum_{n=0}^{N-1} [d(n) - m_1]^p. \quad (6.43)$$

The corresponding normalized moments are defined as

$$\overline{m}_p = \frac{m_p}{(M_2)^{p/2}} = \frac{\frac{1}{N} \sum_{n=0}^{N-1} [d(n)]^p}{\left\{ \frac{1}{N} \sum_{n=0}^{N-1} [d(n) - m_1]^2 \right\}^{p/2}}, \quad (6.44)$$

$$\overline{M}_p = \frac{M_p}{(M_2)^{p/2}} = \frac{\frac{1}{N} \sum_{n=0}^{N-1} [d(n) - m_1]^p}{\left\{ \frac{1}{N} \sum_{n=0}^{N-1} [d(n) - m_1]^2 \right\}^{p/2}}. \quad (6.45)$$

Gupta and Srinath [417] showed that the normalized moments \overline{m}_p and \overline{M}_p are invariant to translation, rotation, and scaling. This set of moments (in an infinite series) reversibly represents the shape of a contour.

Although moments of any arbitrarily large order can be derived from a contour and used as features for shape classification, high-order moments are

sensitive to noise, and hence, the resulting classifier will be less tolerant to noise. Therefore, Gupta and Srinath [417] selected four normalized low-order moments to form a set of shape features as follows:

$$F_1 = \frac{(M_2)^{1/2}}{m_1} = \frac{\left\{ \frac{1}{N} \sum_{n=0}^{N-1} [d(n) - m_1]^2 \right\}^{1/2}}{m_1}, \quad (6.46)$$

$$F_2 = \frac{M_3}{(M_2)^{3/2}} = \frac{\frac{1}{N} \sum_{n=0}^{N-1} [d(n) - m_1]^3}{\left\{ \frac{1}{N} \sum_{n=0}^{N-1} [d(n) - m_1]^2 \right\}^{3/2}}, \quad (6.47)$$

$$F_3 = \frac{M_4}{(M_2)^2} = \frac{\frac{1}{N} \sum_{n=0}^{N-1} [d(n) - m_1]^4}{\left\{ \frac{1}{N} \sum_{n=0}^{N-1} [d(n) - m_1]^2 \right\}^2}. \quad (6.48)$$

A study by Shen et al. [334, 416] showed that the variations in F_2 and F_3 for differing shape complexity are small and do not show a simple progression. Furthermore, F_2 was observed to vary significantly for the same (geometric) shape with scaling and rotation; see Figure 6.15 and Table 6.1. In order to overcome these limitations, F_2 and F_3 were modified by Shen et al. [274, 334, 416] as follows:

$$F'_2 = \frac{M_3^{1/3}}{m_1} = \frac{\left\{ \frac{1}{N} \sum_{n=0}^{N-1} [d(n) - m_1]^3 \right\}^{1/3}}{m_1}, \quad (6.49)$$

$$F'_3 = \frac{M_4^{1/4}}{m_1} = \frac{\left\{ \frac{1}{N} \sum_{n=0}^{N-1} [d(n) - m_1]^4 \right\}^{1/4}}{m_1}. \quad (6.50)$$

Compared with the feature set proposed by Gupta and Srinath [417], the set $\{F_1, F'_2, F'_3\}$ has the following properties:

- All of the three features are directly comparable.
- F'_3 describes the roughness of a contour better than F_3 . In general, the larger the value of F'_3 , the rougher is the contour.

Although F'_2 was observed by Shen et al. [334] to have better invariance with respect to size and rotation for a given geometric shape, it showed no better variation than F_2 across the shape categories tested. However, it was shown that the combination $mf = F'_3 - F'_1$ is a good indicator of shape roughness because the fourth-order term in F'_3 will be much larger than the second-order term in F'_1 as the contour becomes rougher. Also, mf provides the desired invariance for a given contour type as well as the desired variation across the various shape categories; see Figure 6.15 and Table 6.1. Note that the definition of the features F'_1 and F'_3 makes it possible to perform the subtraction directly, and that mf is limited to the range $[0, 1]$.

The values of mf for the contours of the objects in Figure 6.16 are listed in Table 6.2. Observe that the values of mf do not demonstrate the same trends as those of the other shape factors listed in Tables 6.1 and 6.2 for the same contours: the shape factors characterize different notions of shape complexity; see Section 6.6.

6.2.3 Chord-length statistics

Methods to characterize 2D closed contours using their chord-length distributions were proposed by You and Jain [421]. A chord-length measure L_k is defined as the length of the line segment that links a pair of contour points, normalized by the length of the longest chord. The complete set of chords for a given object consists of all possible chords drawn from every boundary pixel to every other boundary pixel; see Figure 6.17. You and Jain considered the $K = N(N - 1)/2$ unique chords of the N boundary points of an object as a sample distribution set, and computed the Kolmogorov-Smirnov (K-S) statistics of the chord-length distribution for use as shape factors as follows:

$$M_{c1} = \frac{1}{K} \sum_{k=1}^K L_k, \quad (6.51)$$

$$M_{c2}^2 = \frac{1}{K} \sum_{k=1}^K (L_k - M_{c1})^2, \quad (6.52)$$

$$M_{c3} = \frac{1}{M_{c2}^3} \frac{1}{K} \sum_{k=1}^K (L_k - M_{c1})^3, \quad (6.53)$$

$$M_{c4} = \frac{1}{M_{c2}^4} \frac{1}{K} \sum_{k=1}^K (L_k - M_{c1})^4. \quad (6.54)$$

The measures listed above, in order, represent the mean, variance, skewness, and kurtosis of the chord-length distributions.

The chord-length statistics are invariant to translation, scaling, and rotation, and are robust in the presence of noise and distortion in the shape




FIGURE 6.17

The set of all possible chords for a contour with $N = 6$ boundary points. There exist $K = N(N - 1)/2 = 15$ unique chords (including the sides of the polygonal contour) in the example. The contour points (0 - 5) are shown in regular font; the chord numbers (1 - 15) are shown in italics.

boundary. The method has a major disadvantage: it is possible for contours of different shapes to have the same chord-length distribution.

The technique was applied by You and Jain to the boundary maps of seven countries and six machine parts with different levels of resolution, and the results indicated good discrimination between the shapes. Rangayyan et al. [163] applied chord-length statistics to the analysis of contours of breast masses, but obtained accuracies of no more than 68% in discriminating between benign masses and malignant tumors; see Section 6.7.

6.3 Fourier Descriptors

Given a contour with N points having the coordinates $\{x(n), y(n)\}$, $n = 0, 1, 2, \dots, N - 1$, we could form a complex sequence $z(n) = x(n) + jy(n)$, $n = 0, 1, 2, \dots, N - 1$; see Figure 6.1. Traversing the contour, it is evident that $z(n)$ is a periodic signal with a period of N samples. The sequence $|z(n)|$ may be used as a signature of the contour.

Periodic signals lend themselves to analysis via the Fourier series. Given a discrete-space sequence $z(n)$, we could derive its Fourier series as one period of its DFT $Z(k)$, defined as

$$Z(k) = \frac{1}{N} \sum_{n=0}^{N-1} z(n) \exp \left[-j \frac{2\pi}{N} nk \right], \quad (6.55)$$

$k = -\frac{N}{2}, \dots, -1, 0, 1, 2, \dots, \frac{N}{2} - 1$. The frequency index k could be interpreted as the index of the harmonics of a fundamental frequency. The contour sample sequence $z(n)$ is given by the inverse DFT as

$$z(n) = \frac{1}{N} \sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} Z(k) \exp \left[j \frac{2\pi}{N} nk \right], \quad (6.56)$$

$n = 0, 1, 2, \dots, N - 1$. The coefficients $Z(k)$ are known as the *Fourier descriptors* of the contour $z(n)$ [8, 422]. (*Note:* Other definitions of Fourier descriptors exist [422, 423, 424, 425, 426].)

Observe that $Z(k)$ represents a two-sided complex spectrum: with folding or shifting of the DFT or FFT array, the frequency index would run as $k = -\frac{N}{2}, \dots, -1, 0, 1, 2, \dots, \frac{N}{2} - 1$; without folding, as usually provided by common DFT or FFT algorithms, the coefficients are provided in the order $0, 1, 2, \dots, \frac{N}{2} - 1, -\frac{N}{2}, \dots, -2, -1$.

A few important properties and characteristics of Fourier descriptors are as follows [8]:

- The frequency index k represents the harmonic number or order in a Fourier series representation of the periodic signal $z(n)$. The fundamental frequency ($k = 1$) represents a sinusoid in each of the coordinates x and y that exhibits one period while traversing once around the closed contour $z(n)$.
- Differing from the usual spectra of real signals, $Z(k)$ does not possess conjugate symmetry due to the complex nature of $z(n)$.
- The zero-frequency (DC) coefficient $Z(0)$ represents the centroid or center of mass (\bar{x}, \bar{y}) , as

$$Z(0) = \frac{1}{N} \sum_{n=0}^{N-1} z(n) = (\bar{x}, \bar{y}). \quad (6.57)$$

- Each of the fundamental frequency coefficients $Z(1)$ and $Z(-1)$ represents a circle. The set of coefficients $\{Z(-1), Z(1)\}$ represents an ellipse.
- High-order Fourier descriptors represent fine details or rapid excursions of the contour.
- The rotation of a contour by an angle θ may be expressed as $z_1(n) = z(n) \exp(j\theta)$, where $z_1(n)$ represents the rotated contour. Rotation leads to an additional phase component as $Z_1(k) = Z(k) \exp(j\theta)$.
- If $z(n)$ represents the points along a contour obtained by traversing the contour in the clockwise direction, and $z_1(n)$ represents the points obtained by traversing in the counter-clockwise direction, we have $z_1(n) = z(-n) = z(N - n)$, and $Z_1(k) = Z(-k)$.
- Shifting or translating $z(n)$ by (x_o, y_o) to obtain $z_1(n) = z(n) + (x_o + j y_o)$ leads to an additional DC component as $Z_1(k) = Z(k) + (x_o + j y_o) \delta(k)$.
- Scaling a contour as $z_1(n) = \alpha z(n)$ leads to a similar scaling of the Fourier descriptors as $Z_1(k) = \alpha Z(k)$.
- Shifting the starting point by n_o samples, expressed as $z_1(n) = z(n - n_o)$, leads to an additional linear-phase component, with $Z_1(k) = Z(k) \exp[-j \frac{2\pi}{N} n_o k]$.

Fourier descriptors may be filtered in a manner similar to the filtering of signals and images in the frequency domain. The full set or a subset of the coefficients may also be used to represent the contour, and to derive shape factors.

Examples: Figure 6.18 shows the normalized Fourier descriptors (up to $k = \pm 15$ only) for the squares and right-angled triangles in Figure 6.15. It is seen that the magnitude of the normalized Fourier descriptors is invariant to scaling and rotation.





FIGURE 6.18

Normalized Fourier descriptors (NFD, up to $k = \pm 15$) for (a) the squares and (b) the right-angled triangles in Figure 6.15. Each figure shows the NFD for four objects; however, due to invariance with respect to scaling and rotation, the functions overlap completely. Reproduced with permission from L. Shen, R.M. Rangayyan, and J.E.L. Desautels, "Application of shape analysis to mammographic calcifications", *IEEE Transactions on Medical Imaging*, 13(2): 263 – 274, 1994. © IEEE.

Figures 6.19 and 6.20 show the $|z(n)|$ signatures and Fourier-descriptor sequences for the benign-mass and malignant-tumor contours shown in Figures 6.2 (a) and 6.3 (a). It is evident that the signatures reflect the smoothness or roughness of the contours: the Fourier descriptors of the spiculated contour indicate the presence of more high-frequency energy than those of the nearly oval contour of the benign mass.

Figure 6.21 shows the results of filtering the benign-mass contour in Figure 6.2 (a) using Fourier descriptors. The coefficients $Z(1)$ and $Z(-1)$ provide two circles, with one of them fitting the contour better than the other (depending upon the direction of traversal of the contour). The combined use of $Z(1)$ and $Z(-1)$ has provided an ellipse that fits the original contour well. The use of additional Fourier descriptors has provided contours that fit the original contour better.

Figure 6.22 shows the results of filtering the malignant-tumor contour in Figure 6.3 (a) using Fourier descriptors. The inclusion of more high-order coefficients has led to contours that approximate the original contour to better levels of fit. The filtered contours illustrate clearly the role played by high-order Fourier descriptors in representing the finer details of the given contour.


Persoon and Fu [423] showed that Fourier descriptors may be used to characterize the skeletons of objects, with applications in character recognition and machine-part recognition. Lin and Chellappa [427] showed that the classification of 2D shapes based on Fourier descriptors is accurate even when 20 – 30% of the data are missing.

Shape factor based upon Fourier descriptors: In a procedure proposed by Shen et al. [274, 334] to derive a single shape factor, the Fourier descriptors are normalized as follows: $Z(0)$ is set equal to zero in order to make the descriptors independent of position, and each coefficient is divided by the magnitude of $Z(1)$ in order to normalize for size. After these steps, the magnitudes of the Fourier descriptors are independent of position, size, orientation, and starting point of the contour; note that the orientation and starting point affect only the phase of the Fourier descriptors. The normalized Fourier descriptors $Z_o(k)$ are defined as


$$Z_o(k) = \begin{cases} 0, & k = 0; \\ \frac{Z(k)}{|Z(1)|}, & \text{otherwise.} \end{cases}$$

[*Note:* For normalization as above, the points of the contour must be indexed from 0 to $(N - 1)$ in counter-clockwise order; in the opposite case, $|Z(-1)|$ should be used.]

Contours with sharp excursions possess more high-frequency energy than smooth contours. However, applying a weighting factor that increases with frequency leads to unbounded values that are also sensitive to noise. A shape factor ff based upon the normalized Fourier descriptors was defined by Shen



(a)



(b)

FIGURE 6.19

- (a) Signature with $|z(n)|$ of the benign-mass contour in Figure 6.2 (a).
(b) Magnitude of the Fourier descriptors, shown only for $k = [-50, 50]$.





FIGURE 6.20

(a) Signature with $|z(n)|$ of the malignant-tumor contour in Figure 6.3 (a).
(b) Magnitude of the Fourier descriptors, shown only for $k = [-50, 50]$.




FIGURE 6.21

Filtering of the benign-mass contour in Figure 6.2 (a) using Fourier descriptors. (a) Using coefficients for $k = 1$ (smaller circle in dashed line), $k = -1$ (larger circle in dashed line), and $k = \{-1, 0, 1\}$ (ellipse in solid line). (b) Using coefficients for $k = [-2, 2]$ (dashed line) and $k = [-3, 3]$ (solid line). The original contour is indicated with a dotted line for reference.




FIGURE 6.22

Filtering of the malignant-tumor contour in Figure 6.3 (a) using Fourier descriptors. (a) Using coefficients for $k = 1$ (smaller circle in dashed line), $k = -1$ (larger circle in dashed line), and $k = \{-1, 0, 1\}$ (ellipse in solid line). (b) Using coefficients for $k = [-10, 10]$ (dashed line) and $k = [-20, 20]$ (solid line). The original contour is indicated with a dotted line for reference.

et al. [274] as

$$ff = 1 - \frac{\sum_{k=-N/2+1}^{N/2} |Z_o(k)|/|k|}{\sum_{k=-N/2+1}^{N/2} |Z_o(k)|}. \quad (6.58)$$

The advantage of this measure is that it is limited to the range $[0, 1]$, and is not sensitive to noise, which would not be the case if weights increasing with frequency were used. ff is invariant to translation, rotation, starting point, and contour size, and increases in value as the object shape becomes more complex and rough.

Other forms of weighting could be used in Equation 6.58 to derive several variants or different shape factors based upon Fourier descriptors. For example, the normalized frequency given by $\frac{|k|}{N/2}$ could be used to provide weights increasing with frequency, and the computation limited to frequencies up to a fraction of the highest available frequency (such as 0.2) in order to limit the effect of noise and high-frequency artifacts. High-order moments could also be computed by using powers of the normalized frequency. Subtraction from unity as in Equation 6.58 could then be removed so as to obtain shape factors that increase with roughness.

The values of ff for the contours of the objects in Figures 6.15 and 6.16 are listed in Tables 6.1 and 6.2. The values of ff do not demonstrate the same trends as those of the other shape factors listed in the tables for the same contours. Several shape factors that characterize different notions of shape complexity may be required for efficient pattern classification of contours in some applications; see Sections 6.6 and 6.7 for illustrations.

Malignant calcifications that have elongated and rough contours lead to larger ff values than benign calcifications that are mostly smooth, round, or oval in shape. Furthermore, tumors with microlobulations and jagged boundaries are expected to have larger ff values than masses with smooth or macrolobulated boundaries. Shen et al. [274, 334] applied ff to shape analysis of mammographic calcifications. The details of this application are presented in Section 6.6. Rangayyan et al. [163] used ff to discriminate between benign breast masses and malignant tumors, and obtained an accuracy of 76%; see Section 6.7. Sahiner et al. [428] tested the classification performance of several shape factors and texture measures with a dataset of 122 benign breast masses and 127 malignant tumors; ff was found to give the best individual performance with an accuracy of 0.82.

6.4 Fractional Concavity

Most benign mass contours are smooth, oval, or have major portions of convex macrolobulations. Some benign masses may have minor concavities and

spicules. On the other hand, malignant tumors typically possess both concave and convex segments as well as microlobulations and prominent spicules. Rangayyan et al. [345] proposed a measure of fractional concavity (f_{cc}) of contours to characterize and quantify these properties.

In order to compute f_{cc} , after performing segmentation of the contour as explained in Section 6.1.3, the individual segments between successive inflection points are labeled as concave or convex parts. A convex part is defined as a segment of the contour that encloses a portion of the mass (inside of the contour), whereas a concave part is one formed by the presence of a background region within the segment. Figure 6.9 shows a section of a mammogram with a circumscribed benign mass, overlaid with the contour drawn by a radiologist specialized in mammography; the black and white portions represent the concave and convex parts, respectively. Figure 6.8 shows a similar result of the analysis of the contour of a spiculated malignant tumor.

The contours used in the work of Rangayyan et al. [345] were manually drawn, and included artifacts and minor modulations that could lead to inefficient representation for pattern classification. The polygonal modeling procedure described in Section 6.1.4 was applied in order to reduce the effect of the artifacts. The cumulative length of the concave segments was computed using the polygonal model, and normalized by the total length of the contour to obtain f_{cc} . It is obvious that f_{cc} is limited to the range $[0, 1]$, and is independent of rotation, shift, and the size (scaling) of the contour. The performance of f_{cc} in discriminating between benign masses and malignant tumors is illustrated in Sections 6.7, 12.11, and 12.12.

Lee et al. [429] proposed an irregularity index for the classification of cutaneous melanocytic lesions based upon their contours. The index was derived via an analysis of the curvature of the contour and the detection of local indentations (concavities) and protrusions (convexities). The irregularity index was observed to have a higher correlation with clinical assessment of the lesions than other shape factors based upon compactness (see Section 6.2.1) and fractal analysis (see Section 7.5).

6.5 Analysis of Spicularity

It is known that invasive carcinomas, due to their nature of infiltration into surrounding tissues, form narrow, stellate distortions or spicules at their boundaries. Based upon this observation, Rangayyan et al. [345] proposed a spiculation index (SI) to represent the degree of spiculation of a mass contour. In order to emphasize narrow spicules and microlobulations, a weighting factor was included to enhance the contributions of narrow spicules in the computation of SI .

For each curved part of a mass contour or the corresponding polygonal model segment, obtained as described in Sections 6.1.3 and 6.1.4, the ratio of its length to the base width can represent the degree of narrowness or spiculation. A nonlinear weighting function was proposed by Rangayyan et al. [345], based upon the segment's length S and angle of spiculation θ , to deliver progressively increasing weighting with increase in the narrowness of spiculation of each segment. Spicule candidates were identified as portions of the contour delimited by pairs of successive points of inflection. The polygonal model, obtained as described in Section 6.1.4, was used to compute the parameters S and θ for each spicule candidate.

If a spicule includes M polygonal segments, then there exist $M - 1$ angles at the points of intersection of the successive segments. Let s_m , $m = 1, 2, \dots, M$, be the polygonal segments, and Θ_n , $n = 1, 2, \dots, M - 1$, be the angles subtended. Then, the segment length (S) and the angle of narrowness (θ) of the spicule under consideration are computed as follows:

1. If $M = 1$, the portion of the contour that has been delimited by successive points of inflection is relatively straight; see Figure 6.10. Such parts are merged into the spicules that include them, thus enhancing the lengths of the corresponding spicules without affecting their angles of spiculation. The merging process discards the redundant points of inflection lying on relatively straight parts of the contour. This may be verified by comparing the initial points of inflection present on the contour in Figure 6.10 with the points of inflection that are retained to compute SI in the corresponding contour shown in Figure 6.23, specifically in the spicule with the angle of spiculation labeled as 116° .
2. If $M = 2$, then the length of spicule is $S = s_1 + s_2$, and the angle subtended by the linear segments at the point of intersection represents the angle of narrowness (θ) of the spicule.
3. If $M > 2$, then the length of the spicule is $S = \sum_{m=1}^M s_m$. In order to estimate the angle of narrowness, an adaptive threshold is applied by using the mean of the set of angles Θ_n , $n = 1, 2, \dots, M - 1$, as the threshold (Θ_{th}) for rejecting insignificant angles (that is, angles that are close to 180°). The mean of the angles that are less than or equal to Θ_{th} is taken as an estimate of the angle of narrowness of the spicule.

Figure 6.24 illustrates the computation of S and θ using the procedure given above for two different examples of spicules with $M = 2$ and $M = 5$, respectively.

Figure 6.23 shows the spicule candidates used in the computation of SI for the contour of the spiculated malignant tumor in Figure 6.8; the corresponding polygonal model is shown in Figure 6.10. The angles of spiculation computed are indicated in Figure 6.23 for all spicule candidates; some of the candidates

may not be considered to be spicules by a radiologist. (*Note:* Visual assessment of the angles of spicules may not agree well with the computed values due to the thresholding and averaging process.) Observe that most of the angles computed for narrow spicules are acute; on the other hand, the angles computed are obtuse for large lobulations and relatively straight segments. The procedure described above adapts to the complexity of each spicule and delivers reliable estimates of the lengths and angles of narrowness of spicules required for computing SI , following polygonal modeling. The computation of SI for a given mass contour is described next.

Let S_n and θ_n , $n = 1, 2, \dots, N$, be the length and angle of N sets of polygonal model segments corresponding to the N spicule candidates of a mass contour. Then, SI is computed as

$$SI = \frac{\sum_{n=1}^N (1 + \cos \theta_n) S_n}{\sum_{n=1}^N S_n}. \quad (6.59)$$

The factor $(1 + \cos \theta_n)$ modulates the length of each segment (possible spicule) according to its narrowness. Spicules with narrow angles between 0° and 30° get high weighting, as compared to macrolobulations that usually form obtuse angles, and hence get low weighting.

The majority of the angles of spicules of the masses and tumors in the MIAS database [376], computed by using the procedure described above, were found to be in the range of 30° to 150° [345]. The function $(1 + \cos \theta_n)$ in Equation 6.59 is progressively decreasing within this range, giving lower weighting to segments with larger angles. Relatively flat segments having angles ranging between 150° and 180° receive low weighting, and hence are treated as insignificant segments.

The denominator in Equation 6.59 serves as a normalization factor to take into account the effect of the size of the contour; it ensures that SI represents only the severity of the spiculated nature of the contour, which in turn may be linked to the invasive properties of the tumor under consideration. The value of SI as in Equation 6.59 is limited to the range $[0, 2]$, and may be normalized to the range $[0, 1]$ by dividing by 2. Circumscribed masses with smooth contours could be expected to have low SI values, whereas sharp, stellate contours with acute spicules should have high SI values. The performance of SI in discriminating between benign masses and malignant tumors is illustrated in Sections 6.7, 12.11, and 12.12.





FIGURE 6.23

The polygonal model used in the procedure to compute SI for the spiculated malignant tumor shown in Figure 6.8 (with the corresponding polygonal model in Figure 6.10). The '⊕' marks correspond to the points of inflection retained to represent the starting and the ending points of spicule candidates, and the '⊙' marks indicate the points of intersection of linear segments within the spicules in the corresponding complete polygonal model. The numbers inside or beside each spicule candidate are the angles in degrees computed for the derivation of SI . Reproduced with permission from R.M. Rangayyan, N.R. Mudigonda, and J.E.L. Desautels, "Boundary modeling and shape analysis methods for classification of mammographic masses", *Medical and Biological Engineering and Computing*, 38: 487 – 496, 2000. © IFMBE.

**FIGURE 6.24**


Computation of segment length S and angle of spiculation θ for two examples of spicule candidates with the number of segments $M = 2$ and $M = 5$, respectively. Θ_{th} is the threshold computed to reject insignificant angles (that is, angles that are close to 180°). Reproduced with permission from R.M. Rangayyan, N.R. Mudigonda, and J.E.L. Desautels, "Boundary modeling and shape analysis methods for classification of mammographic masses", *Medical and Biological Engineering and Computing*, 38: 487 – 496, 2000. © IFMBE.

6.6 Application: Shape Analysis of Calcifications

Because of the higher attenuation coefficient of calcium as compared with normal breast tissues, the main characteristic of calcifications in mammograms is that they are relatively bright. This makes calcifications readily distinguishable on properly acquired mammograms. However, calcifications that appear against a background of dense breast tissue may be difficult to detect; see Sections 5.4.9 and 5.4.10 for illustrations.

Malignant calcifications tend to be numerous, clustered, small, varying in size and shape, angular, irregularly shaped, and branching in orientation [430, 431]. On the other hand, calcifications associated with benign conditions are generally larger, more rounded, smaller in number, more diffusely distributed, and more homogeneous in size and shape. One of the key differences between benign and malignant calcifications lies in the roughness of their shapes.

Shen et al. [274, 334] applied shape analysis to the classification of mammographic calcifications as benign or malignant. Eighteen mammograms of biopsy-proven cases from the Radiology Teaching Library of the Foothills Hospital (Calgary, Alberta, Canada) were digitized with high resolution of up to 2560×4096 pixels with 12 bits per pixel using the Eikonix 1412 scanner. Sixty-four benign calcifications from 11 mammograms and 79 malignant calcifications from seven mammograms were manually selected for shape analysis. Multitolerance region growing (see Section 5.4.9) was performed, and the shape factors (mf , ff , cf) based upon moments (Section 6.2.2), Fourier descriptors (Section 6.3), and compactness (Section 6.2.1) were computed from their boundaries. Figures 5.26 and 5.27 illustrate parts of two mammograms, one with benign calcifications and the other with malignant calcifications, along with the contours of the calcifications that were detected. A plot of the shape factors (mf , ff , cf) for the 143 calcifications in the study of Shen et al. is shown in Figure 6.25. It is evident that most of the malignant calcifications have large values, whereas most of the benign calcifications have low values for the three shape factors. The bar graph in Figure 6.26 indicates that the means of the three features possess good levels of differences between the benign and malignant categories with respect to the corresponding standard deviation values. The three measures represent shape complexity from different perspectives, and hence could be combined for improved discrimination between benign and malignant calcifications. The three features permitted classification of the 143 calcifications with 100% accuracy using the nearest-neighbor method (see Section 12.2.3) as well as neural networks (see Section 12.7).

**FIGURE 6.25**

Plot of the shape factors (mf, ff, cf) of 143 calcifications. The + symbols represent 79 malignant calcifications, and the o symbols represent 64 benign calcifications.




FIGURE 6.26

Means of the shape factors (mf, ff, cf) of 64 benign calcifications ('b') and 79 malignant calcifications ('m'). The error bars indicate the range of mean plus or minus one standard deviation.

6.7 Application: Shape Analysis of Breast Masses and Tumors

Rangayyan et al. [163, 345] applied several shape factors to the analysis of a set of contours of 28 benign breast masses and 26 malignant tumors. The dataset included 39 mammograms from the MIAS database [376] and 15 images from Screen Test: Alberta Program for the Early Detection of Breast Cancer [61]. The contours were drawn on digitized mammograms by an expert radiologist. Figure 6.27 shows the 54 contours arranged in order of increasing shape complexity as characterized by the magnitude of the feature vector (cf, f_{cc}, SI) ; Figure 6.28 shows a scatter plot of the three features. Each of the three features has, in general, the distinction of reflecting low values for circumscribed benign masses and high values for spiculated malignant tumors.

In benign-versus-malignant pattern classification experiments using linear discriminant analysis [163], ff , cf , and mf provided accuracies of 76%, 72%, and 67%, respectively; the moment-based shape factors provided classification accuracy of up to 75%; chord-length statistics provided accuracy up to 68% only. The use of the parameters obtained via parabolic models of segments of the contours separated by their points of inflection led to a classification accuracy of 76% [354]. In a different study [345], the shape factors f_{cc} and SI provided classification accuracies of 74% and 80%, respectively; the set (cf, f_{cc}, SI) provided the highest accuracy of 82%. The MIAS database was observed to include an unusually high proportion of benign masses with spiculated contours, which led to reduced accuracy of benign-versus-malignant classification via shape analysis.

In pattern classification experiments to discriminate between circumscribed and spiculated masses, several combinations of the shape factors mentioned above provided accuracies of up to 91% [163, 345]. However, the classification of a contour as circumscribed or spiculated is a subjective decision of a radiologist; on the other hand, benign-versus-malignant classification via pathology is objective. Furthermore, circumscribed-versus-spiculated classification is of academic interest, with the discrimination between benign disease and malignancy being of clinical relevance and importance. For these reasons, circumscribed-versus-spiculated classification is not important. Sections 7.9, 8.8, 12.11, and 12.12 provide further details on pattern classification of breast masses and tumors.





FIGURE 6.27

Contours of 54 breast masses. 'b': benign masses (28). 'm': malignant tumors (26). The contours are arranged in order of increasing magnitude of the feature vector (cf, f_{cc}, SI). Note that the masses and their contours are of widely differing size, but have been scaled to the same size in the illustration.

**FIGURE 6.28**

Feature-space plot of cf , f_{cc} , and SI : \circ for benign masses (28) and $*$ for malignant tumors (26). SI : spiculation index, f_{cc} : fractional concavity, and cf : modified compactness. See Figure 6.27 for an illustration of the contours.

6.8 Remarks

In this chapter, we have explored several methods to model, characterize, and parameterize contours. Closed contours were considered in most of the discussion and illustrations, although some of the techniques described may be extended to open contours or contours with missing parts.

Regardless of the success of some of the methods and applications illustrated, it should be noted that obtaining contours with good accuracy could be difficult in many applications. It is not common clinical practice to draw the contours of tumors or organs. Malignant tumors typically exhibit poor definitions of their margins due to their invasive and metastatic nature: this makes the identification and drawing of their contours difficult, if not impossible, either manually or by computer methods. Hand-drawn and computer-detected contours may contain imperfections and artifacts that could corrupt shape factors; furthermore, there could be significant variations between the contours drawn by different individuals for the same objects. It should be recognized that the contour of a 3D entity (such as a tumor) as it appears on a 2D image (for example, a mammogram) depends upon the imaging and projection geometry. Above all, contours of biological entities often present significant overlap in their characteristics between various categories, such as for benign and malignant diseases. The inclusion of measures representing other image characteristics, such as texture and gradient, could complement shape factors, and assist in improved analysis of biomedical images. For example, Sahiner et al. [428] showed that the combined use of shape and texture features could improve the accuracy in discriminating between benign breast masses and malignant tumors. Methods for the characterization of texture and gradient information are described in Chapters 7 and 8. The use of the fractal dimension as a measure of roughness is described in Section 7.5. See Chapter 12 for several examples of pattern classification via shape analysis.

6.9 Study Questions and Problems

Selected data files related to some of the problems and exercises are available at the site

www.enel.ualgary.ca/People/Ranga/enel697

1. Prove that the zeroth-order Fourier descriptor represents the centroid of the given contour.
2. Prove that the first-order Fourier descriptors ($k = 1$ or $k = -1$) represent circles.

3. A robotic inspection system is required to discriminate between flat (planar) objects arriving on a conveyor belt. The objects may arrive at any orientation. The set of possible objects includes squares, circles, and triangles of variable size.

Propose an image analysis procedure to detect each object and recognize it as being one of the three types mentioned above. Describe each step of the algorithm briefly. Provide equations for the measures that you may propose.

6.10 Laboratory Exercises and Projects

1. Using black or dark-colored paper, cut out at least 20 pieces of widely varying shapes. Include a few variations of the same geometric shape (square, triangle, etc.) with varying size and orientation.

Lay out the objects on a flat surface and capture an image. Develop a program to detect the objects and derive their contours. Verify that the contours are closed, are one-pixel thick, and do not include knots.

Derive several shape factors for each contour, including compactness, Fourier descriptors, moments, and fractional concavity. Rank-order the objects by each shape factor individually, and by all of the factors combined into a single vector. Study the characterization of various notions of shape complexity by the different shape factors.

Request a number of your friends and colleagues to assign a measure of roughness to each object on a scale of 0 – 100. Normalize the values by dividing by 100 and average the scores over all the observers. Analyze the correlation between the subjective ranking and the objective measures of roughness.

2. Synthesize a digital image with rectangles, triangles, and circles of various sizes. Compute several of the shape factors described in this chapter for each object in the image.

Study the variation in the shape factors from one category of shapes to another in your test image. Is the variation adequate to facilitate pattern classification?

Study the variation in the shape factors within each category of shapes in your test image. Explain the cause of the variation.

Analysis of Texture

Texture is one of the important characteristics of images, and texture analysis is encountered in several areas [432, 433, 434, 435, 436, 437, 438, 439, 440, 441, 442]. We find around us several examples of texture: on wooden furniture, cloth, brick walls, floors, and so on. We may group texture into two general categories: (quasi-) periodic and random. If there is a repetition of a texture element at almost regular or (quasi-) periodic intervals, we may classify the texture as being (quasi-) periodic or ordered; the elements of such a texture are called *textons* [438] or *textels*. Brick walls and floors with tiles are examples of periodic texture. On the other hand, if no *texton* can be identified, such as in clouds and cement-wall surfaces, we can say that the texture is random. Rao [432] gives a more detailed classification, including weakly ordered or oriented texture that takes into account hair, wood grain, and brush strokes in paintings. Texture may also be related to visual and/or tactile sensations such as fineness, coarseness, smoothness, granularity, periodicity, patchiness, being mottled, or having a preferred orientation [441].

A significant amount of work has been done in texture characterization [441, 442, 439, 432, 438] and synthesis [443, 438]; see Haralick [441] and Haralick and Shapiro [440] (Chapter 9) for detailed reviews. According to Haralick et al. [442], texture relates to information about the spatial distribution of gray-level variation; however, this is a general observation. It is important to recognize that, due to the existence of a wide variety of texture, no single method of analysis would be applicable to several different situations. Statistical measures such as gray-level co-occurrence matrices and entropy [442] characterize texture in a stochastic sense; however, they do not convey a physical or perceptual sense of the texture. Although periodic texture may be modeled as repetitions of *textons*, not many methods have been developed for the structural analysis of texture [444].

In this chapter, we shall explore the nature of texture found in biomedical images, study methods to characterize and analyze such texture, and investigate approaches for the classification of biomedical images based upon texture. We shall concentrate on random texture in this chapter; due to the extensive occurrence of oriented patterns and texture in biomedical images, we shall treat this topic on its own, in Chapter 8.

7.1 Texture in Biomedical Images

A wide variety of texture is encountered in biomedical images. Oriented texture is common in medical images due to the fibrous nature of muscles and ligaments, as well as the extensive presence of networks of blood vessels, veins, ducts, and nerves. A preferred or dominant orientation is associated with the functional integrity and strength of such structures. Although truly periodic texture is not commonly encountered in biomedical images, ordered texture is often found in images of the skins of reptiles, the retina, the cornea, the compound eyes of insects, and honeycombs.


Organs such as the liver are made up of clusters of parenchyma that are of the order of $1 - 2 \text{ mm}$ in size. The pixels in CT images have a typical resolution of $1 \times 1 \text{ mm}$, which is comparable to the size of the parenchymal units. With ultrasonic imaging, the wavelength of the probing radiation is of the order of $1 - 2 \text{ mm}$, which is also comparable to the size of parenchymal clusters. Under these conditions, the liver appears to have a speckled random texture.

Several samples of biomedical images with various types of texture are shown in Figures 7.1, 7.2, and 7.3; see also Figures 1.5, 1.8, 9.18, and 9.20. It is evident from these illustrations that no single approach can succeed in characterizing all types of texture.


Several approaches have been proposed for the analysis of texture in medical images for various diagnostic applications. For example, texture measures have been derived from X-ray images for automatic identification of pulmonary diseases [433], for the analysis of MR images [445], and processing of mammograms [165, 275, 446]. In this chapter, we shall investigate the nature of texture in a few biomedical images, and study some of the commonly used methods for texture analysis.

7.2 Models for the Generation of Texture


Martins et al. [447], in their work on the auditory display of texture in images (see Section 7.8), outlined the following similarities between speech and texture generation. The sounds produced by the human vocal system may be grouped as voiced, unvoiced, and plosive sounds [31, 176]. The first two types of speech signals may be modeled as the convolution of an input excitation signal with a filter function. The excitation signal is quasi-periodic when we use the vocal cords to create voiced sounds, or random in the case of unvoiced sounds. Figure 7.4 (a) illustrates the basic model for speech generation.

**FIGURE 7.1**

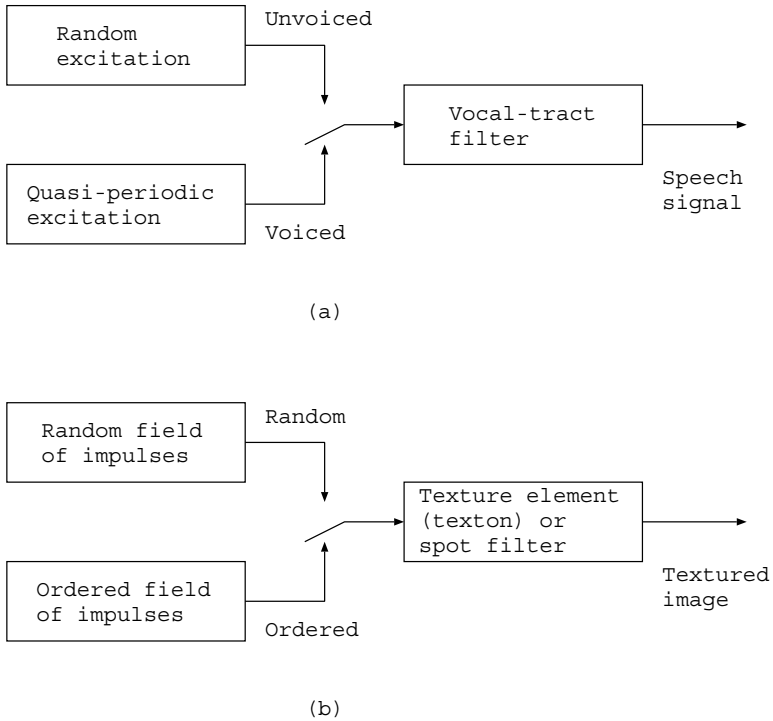
Examples of texture in CT images: (a) Liver. (b) Kidney. (c) Spine. (d) Lung. The true size of each image is 55×55 mm. The images represent widely differing ranges of tissue density, and have been enhanced to display the inherent texture. Image data courtesy of Alberta Children's Hospital.

**FIGURE 7.2**

Examples of texture in mammograms (from the MIAS database [376]): (a) – (c) oriented texture; true image size 60×60 mm; (d) random texture; true image size 40×40 mm. For more examples of oriented texture, see Figures 9.20 and 1.8, as well as Chapter 8.

**FIGURE 7.3**

Examples of ordered texture: (a) Endothelial cells in the cornea. Image courtesy of J. Jaroszewski. (b) Part of a fly's eye. Reproduced with permission from D. Suzuki, "Behavior in *drosophila melanogaster*: A geneticist's view", *Canadian Journal of Genetics and Cytology*, XVI(4): 713 – 735, 1974. © Genetics Society of Canada. (c) Skin on the belly of a cobra snake. Image courtesy of Implora, Colonial Heights, VA. <http://www.implora.com>. See also Figure 1.5.

**FIGURE 7.4**

(a) Model for speech signal generation. (b) Model for texture synthesis. Reproduced with permission from A.C.G. Martins, R.M. Rangayyan, and R.A. Ruschioni, "Audification and sonification of texture in images", *Journal of Electronic Imaging*, 10(3): 690 – 705, 2001. © SPIE and IS&T.

Texture may also be modeled as the convolution of an input impulse field with a spot or a texton that would act as a filter. The “spot noise” model of van Wijk [443] for synthesizing random texture uses this model, in which the Fourier spectrum of the spot acts as a filter that modifies the spectrum of a 2D random-noise field. Ordered texture may be generated by specifying the basic pattern or texton to be used, and a placement rule. The placement rule may be expressed as a field of impulses. Texture is then given by the convolution of the impulse field with the texton, which could also be represented as a filter. A one-to-one correspondence may thus be established between speech signals and texture in images. Figure 7.4 (b) illustrates the model for texture synthesis: the correspondence between the speech and image generation models in Figure 7.4 is straightforward.


7.2.1 Random texture

According to the model in Figure 7.4, random texture may be modeled as a filtered version of a field of white noise, where the filter is represented by a spot of a certain shape and size (usually of small spatial extent compared to the size of the image). The 2D spectrum of the noise field, which is essentially a constant, is shaped by the 2D spectrum of the spot. Figure 7.5 illustrates a random-noise field of size 256×256 pixels and its Fourier spectrum. Parts (a) – (d) of Figure 7.6 show two circular spots of diameter 12 and 20 pixels and their spectra; parts (e) – (h) of the figure show the random texture generated by convolving the noise field in Figure 7.5 (a) with the circular spots, and their Fourier spectra. It is readily seen that the spots have filtered the noise, and that the spectra of the textured images are essentially those of the corresponding spots.

Figures 7.7 and 7.8 illustrate a square spot and a hash-shaped spot, as well as the corresponding random texture generated by the spot-noise model and the corresponding spectra; the anisotropic nature of the images is clearly seen in their spectra.

7.2.2 Ordered texture

Ordered texture may be modeled as the placement of a basic pattern or texton (which is of a much smaller size than the total image) at positions determined by a 2D field of (quasi-) periodic impulses. The separations between the impulses in the x and y directions determine the periodicity or “pitch” in the two directions. This process may also be modeled as the convolution of the impulse field with the texton; in this sense, the only difference between ordered and random texture lies in the structure of the impulse field: the former uses a (quasi-) periodic field of impulses, whereas the latter uses a random-noise field. Once again, the spectral characteristics of the texton could be seen as a filter that modifies the spectrum of the impulse field (which is essentially a 2D field of impulses as well).


**FIGURE 7.5**

(a) Image of a random-noise field (256×256 pixels). (b) Spectrum of the image in (a). Reproduced with permission from A.C.G. Martins, R.M. Rangayyan, and R.A. Ruschioni, “Audification and sonification of texture in images”, *Journal of Electronic Imaging*, 10(3): 690 – 705, 2001. © SPIE and IS&T.


Figure 7.9 (a) illustrates a 256×256 field of impulses with horizontal periodicity $p_x = 40$ pixels and vertical periodicity $p_y = 40$ pixels. Figure 7.9 (b) shows the corresponding periodic texture with a circle of diameter 20 pixels as the spot or texton. Figure 7.9 (c) shows a periodic texture with the texton being a square of side 20 pixels, $p_x = 40$ pixels, and $p_y = 40$ pixels. Figure 7.9 (d) depicts a periodic-textured image with an isosceles triangle of sides 12, 16, and 23 pixels as the spot, and periodicity $p_x = 40$ pixels and $p_y = 40$ pixels. See Section 7.6 for illustrations of the Fourier spectra of images with ordered texture.

7.2.3 Oriented texture


Images with oriented texture may be generated using the spot-noise model by providing line segments or oriented motifs as the spot. Figure 7.10 shows a spot with a line segment oriented at 135° and the result of convolution of the spot with a random-noise field; the log-magnitude Fourier spectra of the spot and the textured image are also shown. The preferred orientation of the texture and the directional concentration of the energy in the Fourier domain are clearly seen in the figure. See Figure 7.2 for examples of oriented texture in mammograms. See Chapter 8 for detailed discussions on the analysis of oriented texture and several illustrations of oriented patterns.




(a)




(b)




(c)



(d)



(e)



(f)






Figure 7.6 (g)




(h)

FIGURE 7.6

(a) Circle of diameter 12 pixels. (b) Circle of diameter 20 pixels. (c) Fourier spectrum of the image in (a). (d) Fourier spectrum of the image in (b). (e) Random texture with the circle of diameter 12 pixels as the spot. (f) Random texture with the circle of diameter 20 pixels as the spot. (g) Fourier spectrum of the image in (e). (h) Fourier spectrum of the image in (f). The size of each image is 256×256 pixels. Reproduced with permission from A.C.G. Martins, R.M. Rangayyan, and R.A. Ruschioni, "Audification and sonification of texture in images", *Journal of Electronic Imaging*, 10(3): 690 – 705, 2001. © SPIE and IS&T.

**FIGURE 7.7**

(a) Square of side 20 pixels. (b) Random texture with the square of side 20 pixels as the spot. (c) Spectrum of the image in (a). (d) Spectrum of the image in (b). The size of each image is 256×256 pixels. Reproduced with permission from A.C.G. Martins, R.M. Rangayyan, and R.A. Ruschioni, "Audification and sonification of texture in images", *Journal of Electronic Imaging*, 10(3): 690 – 705, 2001. © SPIE and IS&T.

**FIGURE 7.8**

(a) Hash of side 20 pixels. (b) Random texture with the hash of side 20 pixels as the spot. (c) Spectrum of the image in (a). (d) Spectrum of the image in (b). The size of each image is 256×256 pixels. Reproduced with permission from A.C.G. Martins, R.M. Rangayyan, and R.A. Ruschioni, "Audification and sonification of texture in images", *Journal of Electronic Imaging*, 10(3): 690 – 705, 2001. © SPIE and IS&T.

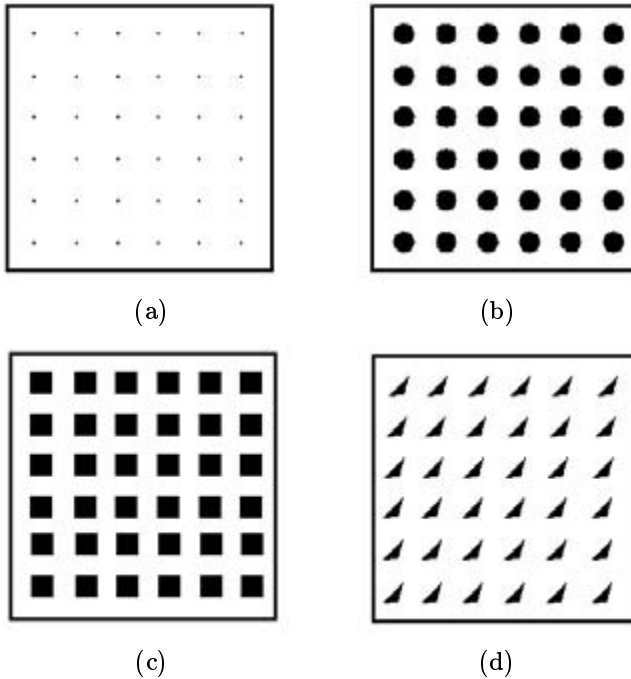



FIGURE 7.9

(a) Periodic field of impulses with $p_x = 40$ pixels and $p_y = 40$ pixels. (b) Ordered texture with a circle of diameter 20 pixels, $p_x = 40$ pixels, and $p_y = 40$ pixels as the spot. (c) Ordered texture with a square of side 20 pixels, $p_x = 40$ pixels, and $p_y = 40$ pixels as the spot. (d) Ordered texture with a triangle of sides 12, 16, and 23 pixels as the spot; $p_x = 40$ pixels; and $p_y = 40$ pixels. The size of each image is 256×256 pixels. Reproduced with permission from A.C.G. Martins, R.M. Rangayyan, and R.A. Ruschioni, "Audification and sonification of texture in images", *Journal of Electronic Imaging*, 10(3): 690 – 705, 2001. © SPIE and IS&T.

**FIGURE 7.10**

Example of oriented texture generated using the spot-noise model in Figure 7.4: (a) Spot with a line segment oriented at 135° . (b) Oriented texture generated by convolving the spot in (a) with a random-noise field. (c) and (d) Log-magnitude Fourier spectra of the spot and the textured image, respectively. The size of each image is 256×256 pixels.

7.3 Statistical Analysis of Texture

Simple measures of texture may be derived based upon the moments of the gray-level PDF (or normalized histogram) of the given image. The k^{th} central moment of the PDF $p(l)$ is defined as

$$m_k = \sum_{l=0}^{L-1} (l - \mu_f)^k p(l), \quad (7.1)$$

where $l = 0, 1, 2, \dots, L - 1$ are the gray levels in the image f , and μ_f is the mean gray level of the image given by

$$\mu_f = \sum_{l=0}^{L-1} l p(l). \quad (7.2)$$

The second central moment, which is the variance of the gray levels and is given by

$$\sigma_f^2 = m_2 = \sum_{l=0}^{L-1} (l - \mu_f)^2 p(l), \quad (7.3)$$

can serve as a measure of inhomogeneity. The normalized third and fourth moments, known as the skewness and kurtosis, respectively, and defined as

$$\text{skewness} = \frac{m_3}{m_2^{3/2}}, \quad (7.4)$$

and

$$\text{kurtosis} = \frac{m_4}{m_2^2}, \quad (7.5)$$

indicate the asymmetry and uniformity (or lack thereof) of the PDF. High-order moments are affected significantly by noise or error in the PDF, and may not be reliable features. The moments of the PDF can only serve as basic representatives of gray-level variation.

Byng et al. [448] computed the skewness of the histograms of 24×24 (3.12×3.12 mm) sections of mammograms. An average skewness measure was computed for each image by averaging over all the section-based skewness measures of the image. Mammograms of breasts with increased fibroglandular density were observed to have histograms skewed toward higher density, resulting in negative skewness. On the other hand, mammograms of fatty breasts tended to have positive skewness. The skewness measure was found to be useful in predicting the risk of development of breast cancer.

7.3.1 The gray-level co-occurrence matrix


Given the general description of texture as a pattern of the occurrence of gray levels in space, the most commonly used measures of texture, in particular of random texture, are the statistical measures proposed by Haralick et al. [441, 442]. Haralick's measures are based upon the moments of a joint PDF that is estimated as the joint occurrence or co-occurrence of gray levels, known as the gray-level co-occurrence matrix (GCM). GCMs are also known as spatial gray-level dependence (SGLD) matrices, and may be computed for various orientations and distances.

The GCM $P_{(d, \theta)}(l_1, l_2)$ represents the probability of occurrence of the pair of gray levels (l_1, l_2) separated by a given distance d at angle θ . GCMs are constructed by mapping the gray-level co-occurrence counts or probabilities based on the spatial relations of pixels at different angular directions (specified by θ) while scanning the image from left-to-right and top-to-bottom.

Table 7.1 shows the GCM for the image in Figure 7.11 with eight gray levels (3 $b/pixel$) by considering pairs of pixels with the second pixel immediately below the first. For example, the pair of gray levels $\begin{bmatrix} 1 \\ 2 \end{bmatrix}$ occurs 10 times in the image. Observe that the table of counts of occurrence of pairs of pixels shown in Table 11.2 and used to compute the first-order entropy also represents a GCM, with the second pixel appearing immediately after the first in the same row. Due to the fact that neighboring pixels in natural images tend to have nearly the same values, GCMs tend to have large values along and around the main diagonal, and low values away from the diagonal.

Observe that, for an image with B $b/pixel$, there will be $L = 2^B$ gray levels; the GCM is then of size $L \times L$. Thus, for an image quantized to 8 $b/pixel$, there will be 256 gray levels, and the GCM will be of size 256×256 . Fine quantization to large numbers of gray levels, such as $2^{12} = 4,096$ levels in high-resolution mammograms, will increase the size of the GCM to unmanageable levels, and also reduce the values of the entries in the GCM. It may be advantageous to reduce the number of gray levels to a relatively small number before computing GCMs. A reduction in the number of gray levels with smoothing can also reduce the effect of noise on the statistics computed from GCMs.

GCMs are commonly formed for unit pixel distances and the four angles of 0° , 45° , 90° , and 135° . (Strictly speaking, the distances to the diagonally connected neighboring pixels at 45° and 135° would be $\sqrt{2}$ times the pixel size.) For an $M \times N$ image, the number of pairs of pixels that can be formed will be less than MN due to the fact that it may not be possible to pair the pixels in a few rows or columns at the borders of the image with another pixel according to the chosen parameters (d, θ) .



1	1	1	1	1	1	1	1	1	1	2	3	2	2	1	2
0	1	1	1	1	1	1	1	1	1	1	2	2	3	4	5
1	0	0	0	1	1	1	1	1	1	1	1	2	2	4	6
2	2	3	5	4	3	1	0	1	1	1	1	1	2	3	5
4	6	5	4	3	1	1	2	2	1	1	1	1	1	2	4
5	5	2	1	2	3	2	2	2	3	3	4	3	2	1	3
4	3	1	2	1	1	1	2	2	2	1	2	2	2	3	5
2	0	2	0	1	3	1	3	5	3	3	2	2	3	3	6
1	1	2	2	1	2	1	2	3	3	3	4	4	6	5	6
1	1	2	4	1	0	0	1	3	4	5	5	5	4	4	6
1	1	1	4	2	1	2	3	5	5	5	4	4	3	4	6
1	1	1	4	4	4	5	6	6	5	4	3	2	3	5	6
1	1	2	5	5	4	5	5	4	3	3	2	3	4	5	6
2	1	4	5	5	5	5	4	3	1	1	1	4	6	5	6
2	2	5	5	5	4	3	2	2	1	1	4	6	6	6	7
4	4	4	4	3	2	2	1	0	1	5	6	6	6	6	7

FIGURE 7.11

A 16×16 part of the image in Figure 2.1 (a) quantized to $3 b/\text{pixel}$, shown as an image and as a 2D array of pixel values.

TABLE 7.1

Gray-level Co-occurrence Matrix for the Image in Figure 7.11, with the Second Pixel Immediately Below the First.

Current Pixel	Next Pixel Below							
	0	1	2	3	4	5	6	7
0	0	3	4	1	0	1	0	0
1	6	44	10	9	5	1	0	0
2	3	13	13	5	8	3	1	0
3	1	5	11	5	3	5	2	0
4	0	1	5	7	5	9	3	0
5	0	0	1	5	11	10	4	0
6	0	0	0	0	2	3	10	1
7	0	0	0	0	0	0	0	1

Pixels in the last row were not processed. The GCM has not been normalized. See also Table 11.2.

7.3.2 Haralick's measures of texture

Based upon normalized GCMs, Haralick et al. [441, 442] proposed several quantities as measures of texture. In order to define these measures, let us normalize the GCM as

$$p(l_1, l_2) = \frac{P(l_1, l_2)}{\sum_{l_1=0}^{L-1} \sum_{l_2=0}^{L-1} P(l_1, l_2)}. \quad (7.6)$$

A few other entities used in the derivation of Haralick's texture measures are as follows:

$$p_x(l_1) = \sum_{l_2=0}^{L-1} p(l_1, l_2), \quad (7.7)$$

$$p_y(l_2) = \sum_{l_1=0}^{L-1} p(l_1, l_2), \quad (7.8)$$

$$p_{x+y}(k) = \underbrace{\sum_{l_1=0}^{L-1} \sum_{l_2=0}^{L-1}}_{l_1+l_2=k} p(l_1, l_2), \quad (7.9)$$

where $k = 0, 1, 2, \dots, 2(L-1)$, and

$$p_{x-y}(k) = \underbrace{\sum_{l_1=0}^{L-1} \sum_{l_2=0}^{L-1}}_{|l_1-l_2|=k} p(l_1, l_2), \quad (7.10)$$

where $k = 0, 1, 2, \dots, L-1$.

The texture measures are then defined as follows.

The energy feature F_1 , which is a measure of homogeneity, is defined as

$$F_1 = \sum_{l_1=0}^{L-1} \sum_{l_2=0}^{L-1} p^2(l_1, l_2). \quad (7.11)$$

A homogeneous image has a small number of entries along the diagonal of the GCM with large values, which will lead to a large value of F_1 . On the other hand, an inhomogeneous image will have small values spread over a larger number of GCM entries, which will result in a low value for F_1 .

The contrast feature F_2 is defined as

$$F_2 = \sum_{k=0}^{L-1} k^2 \underbrace{\sum_{l_1=0}^{L-1} \sum_{l_2=0}^{L-1}}_{|l_1-l_2|=k} p(l_1, l_2). \quad (7.12)$$

The correlation measure F_3 , which represents linear dependencies of gray levels, is defined as

$$F_3 = \frac{1}{\sigma_x \sigma_y} \left[\sum_{l_1=0}^{L-1} \sum_{l_2=0}^{L-1} l_1 l_2 p(l_1, l_2) - \mu_x \mu_y \right], \quad (7.13)$$

where μ_x and μ_y are the means, and σ_x and σ_y are the standard deviations of p_x and p_y , respectively.

The sum of squares feature is given by

$$F_4 = \sum_{l_1=0}^{L-1} \sum_{l_2=0}^{L-1} (l_1 - \mu_f)^2 p(l_1, l_2), \quad (7.14)$$

where μ_f is the mean gray level of the image.

The inverse difference moment, a measure of local homogeneity, is defined as

$$F_5 = \sum_{l_1=0}^{L-1} \sum_{l_2=0}^{L-1} \frac{1}{1 + (l_1 - l_2)^2} p(l_1, l_2). \quad (7.15)$$

The sum average feature F_6 is given by

$$F_6 = \sum_{k=0}^{2(L-1)} k p_{x+y}(k), \quad (7.16)$$

and the sum variance feature F_7 is defined as

$$F_7 = \sum_{k=0}^{2(L-1)} (k - F_6)^2 p_{x+y}(k). \quad (7.17)$$

The sum entropy feature F_8 is given by

$$F_8 = - \sum_{k=0}^{2(L-1)} p_{x+y}(k) \log_2 [p_{x+y}(k)]. \quad (7.18)$$

Entropy, a measure of nonuniformity in the image or the complexity of the texture, is defined as

$$F_9 = - \sum_{l_1=0}^{L-1} \sum_{l_2=0}^{L-1} p(l_1, l_2) \log_2 [p(l_1, l_2)]. \quad (7.19)$$

The difference variance measure F_{10} is defined as the variance of p_{x-y} , in a manner similar to that given by Equations 7.16 and 7.17 for its sum counterpart.

The difference entropy measure is defined as

$$F_{11} = - \sum_{k=0}^{L-1} p_{x-y}(k) \log_2 [p_{x-y}(k)]. \quad (7.20)$$

Two information-theoretic measures of correlation are defined as

$$F_{12} = \frac{H_{xy} - H_{xy1}}{\max\{H_x, H_y\}}, \quad (7.21)$$

and

$$F_{13} = \{1 - \exp[-2 (H_{xy2} - H_{xy})]\}^2, \quad (7.22)$$

where $H_{xy} = F_9$; H_x and H_y are the entropies of p_x and p_y , respectively;

$$H_{xy1} = - \sum_{l_1=0}^{L-1} \sum_{l_2=0}^{L-1} p(l_1, l_2) \log_2 [p_x(l_1) p_y(l_2)], \quad (7.23)$$

and

$$H_{xy2} = - \sum_{l_1=0}^{L-1} \sum_{l_2=0}^{L-1} p_x(l_1) p_y(l_2) \log_2 [p_x(l_1) p_y(l_2)]. \quad (7.24)$$

The maximal correlation coefficient feature F_{14} is defined as the square root of the second largest eigenvalue of \mathbf{Q} , where

$$\mathbf{Q}(l_1, l_2) = \sum_{k=0}^{L-1} \frac{p(l_1, k) p(l_2, k)}{p_x(k) p_y(k)}. \quad (7.25)$$

The subscripts d and θ in the representation of the GCM $P_{(\theta, d)}(l_1, l_2)$ have been removed in the definitions above for the sake of notational simplicity. However, it should be noted that each of the measures defined above may be derived for each value of d and θ of interest. If the dependence of texture upon angle is not of interest, GCMs over all angles may be averaged into a single GCM. The distance d should be chosen taking into account the sampling interval (pixel size) and the size of the texture units of interest. More details on the derivation and significance of the features defined above are provided by Haralick et al. [441, 442].

Some of the features defined above have values much greater than unity, whereas some of the features have values far less than unity. Normalization to a predefined range, such as $[0, 1]$, over the dataset to be analyzed, may be beneficial.

Parkkinen et al. [449] studied the problem of detecting periodicity in texture using statistical measures of association and agreement computed from GCMs. If the displacement and orientation (d, θ) of a GCM match the same parameters of the texture, the GCM will have large values for the elements

along the diagonal corresponding to the gray levels present in the texture elements. A measure of association is the χ^2 statistic, which may be expressed using the notation above as

$$\chi^2 = \sum_{l_1=0}^{L-1} \sum_{l_2=0}^{L-1} \frac{[p(l_1, l_2) - p_x(l_1) p_y(l_2)]^2}{p_x(l_1) p_y(l_2)}. \quad (7.26)$$

The measure may be normalized by dividing by L , and expected to possess a high value for an image with periodic texture under the condition described above.

Parkkinen et al. [449] discussed some limitations of the χ^2 statistic in the analysis of periodic texture, and proposed a measure of agreement given by

$$\kappa = \frac{P_o - P_c}{1 - P_c}, \quad (7.27)$$

where

$$P_o = \sum_{l=0}^{L-1} p(l, l), \quad (7.28)$$

and

$$P_c = \sum_{l=0}^{L-1} p_x(l) p_y(l). \quad (7.29)$$

The measure κ has its maximal value of unity when the GCM is a diagonal matrix, which indicates perfect agreement or periodic texture.

Haralick's measures have been applied for the analysis of texture in several types of images, including medical images. Chan et al. [450] found the three features of correlation, difference entropy, and entropy to perform better than other combinations of one to eight features selected in a specific sequence. Sahiner et al. [428, 451] defined a "rubber-band straightening transform" (RBST) to map ribbons around breast masses in mammograms into rectangular arrays (see Figure 7.26), and then computed Haralick's measures of texture. Mudigonda et al. [165, 275] computed Haralick's measures using adaptive ribbons of pixels extracted around mammographic masses, and used the features to distinguish malignant tumors from benign masses; details of this work are provided in Sections 7.9 and 8.8. See Section 12.12 for a discussion on the application of texture measures for content-based retrieval and classification of mammographic masses.

7.4 Laws' Measures of Texture Energy

Laws [452] proposed a method for classifying each pixel in an image based upon measures of local "texture energy". The texture energy features rep-

resent the amounts of variation within a sliding window applied to several filtered versions of the given image. The filters are specified as separable 1D arrays for convolution with the image being processed.

The basic operators in Laws' method are the following:

$$\begin{aligned} L3 &= [1 \quad 2 \quad 1], \\ E3 &= [-1 \quad 0 \quad 1], \\ S3 &= [-1 \quad 2 \quad -1]. \end{aligned} \tag{7.30}$$

The operators $L3$, $E3$, and $S3$ perform center-weighted averaging, symmetric first differencing (edge detection), and second differencing (spot detection), respectively [453]. Nine 3×3 masks may be generated by multiplying the transposes of the three operators (represented as vectors) with their direct versions. The result of $L3^T E3$ gives one of the 3×3 Sobel masks.

Operators of length five pixels may be generated by convolving the $L3$, $E3$, and $S3$ operators in various combinations. Of the several filters designed by Laws, the following five were said to provide good performance [452, 453]:

$$\begin{aligned} L5 &= L3 * L3 = [1 \quad 4 \quad 6 \quad 4 \quad 1], \\ E5 &= L3 * E3 = [-1 \quad -2 \quad 0 \quad 2 \quad 1], \\ S5 &= -E3 * E3 = [-1 \quad 0 \quad 2 \quad 0 \quad -1], \\ R5 &= S3 * S3 = [1 \quad -4 \quad 6 \quad -4 \quad 1], \\ W5 &= -E3 * S3 = [-1 \quad 2 \quad 0 \quad -2 \quad 1], \end{aligned} \tag{7.31}$$

where $*$ represents 1D convolution.

The operators listed above perform the detection of the following types of features: $L5$ – local average; $E5$ – edges; $S5$ – spots; $R5$ – ripples; and $W5$ – waves [453]. In the analysis of texture in 2D images, the 1D convolution operators given above are used in pairs to achieve various 2D convolution operators (for example, $L5L5 = L5^T L5$ and $L5E5 = L5^T E5$), each of which may be represented as a 5×5 array or matrix. Following the application of the selected filters, texture energy measures are derived from each filtered image by computing the sum of the absolute values in a 15×15 sliding window.

All of the filters listed above, except $L5$, have zero mean, and hence the texture energy measures derived from the filtered images represent measures of local deviation or variation. The result of the $L5$ filter may be used for normalization with respect to luminance and contrast.

The use of a large sliding window to smooth the filtered images could lead to the loss of boundaries across regions with different texture. Hsiao and

Sawchuk [454] applied a modified LLMMSE filter so as to derive Laws' texture energy measures while preserving the edges of regions, and applied the results for pattern classification.

Example: The results of the application of the operators $L5L5$, $E5E5$, and $W5W5$ to the 128×128 Lenna image in Figure 10.5 (a) are shown in Figure 7.12 (a) – (c). Also shown in parts (e) – (f) of the figure are the sums of the absolute values of the filtered images using a 9×9 moving window. It is evident that the $L5L5$ filter results in a measure of local brightness. Careful inspection of the results of the $E5E5$ and $W5W5$ filters shows that they have high values for different regions of the original image possessing different types of texture (edges and waves, respectively). Feature vectors composed of the values of various Laws' operators for each pixel may be used for classifying the image into texture categories on a pixel-by-pixel basis. The results may be used for texture segmentation and recognition.

In an example provided by Laws [452] (see also Pietikäinen et al. [453]), the texture energy measures have been shown to be useful in the segmentation of an image composed of patches with different texture. Miller and Astley [372, 455] used features of mammograms based upon the $R5R5$ operator, and obtained an accuracy of 80.3% in the segmentation of the nonfat (glandular) regions in mammograms. See Section 8.8 for a discussion on the application of Laws' and other methods of texture analysis for the detection of breast masses in mammograms.

7.5 Fractal Analysis

Fractals are defined in several different ways, the most common of which is that of a pattern composed of repeated occurrences of a basic unit at multiple scales of detail in a certain order of generation; this definition includes the notion of “self-similarity” or nested recurrence of the same motif at smaller and smaller scales (see Section 11.9 for a discussion on self-similar, space-filling curves). The relationship to texture is evident in the property of repeated occurrence of a motif. Fractal patterns occur abundantly in nature as well as in biological and physiological systems [456, 457, 458, 459, 460, 461, 462]: the self-replicating patterns of the complex leaf structures of ferns (see Figure 7.13), the ramifications of the bronchial tree in the lung (see Figure 7.1), and the branching and spreading (anastomotic) patterns of the arteries in the heart (see Figure 9.20), to name a few. Fractals and the notion of chaos are related to the area of nonlinear dynamic systems [456, 463], and have found several applications in biomedical signal and image analysis.




FIGURE 7.12

Results of convolution of the Lenna test image of size 128×128 pixels [see Figure 10.5 (a)] using the following 5×5 Laws' operators: (a) $L5L5$, (b) $E5E5$, and (c) $W5W5$. (d) – (f) were obtained by summing the absolute values of the results in (a) – (c), respectively, in a 9×9 moving window, and represent three measures of texture energy. The image in (c) was obtained by mapping the range $[-200, 200]$ out of the full range of $[-1338, 1184]$ to $[0, 255]$.



FIGURE 7.13
The leaf of a fern with a fractal pattern.

7.5.1 Fractal dimension

Whereas the self-similar aspect of fractals is apparent in the examples mentioned above, it is not so obvious in other patterns such as clouds, coastlines, and mammograms, which are also said to have fractal-like characteristics. In such cases, the “fractal nature” perceived is more easily related to the notion of complexity in the dimensionality of the object, leading to the concept of the *fractal dimension*. If one were to use a large ruler to measure the length of a coastline, the minor details present in the border having small-scale variations would be skipped, and a certain length would be derived. If a smaller ruler were to be used, smaller details would get measured, and the total length that is measured would increase (between the same end points as before). This relationship may be expressed as [457]

$$l(\eta) = l_0 \eta^{1-d_f}, \quad (7.32)$$

where $l(\eta)$ is the length measured with η as the measuring unit (the size of the ruler), d_f is the fractal dimension, and l_0 is a constant. Fractal patterns exhibit a linear relationship between the log of the measured length and the log of the measuring unit:

$$\log[l(\eta)] = \log[l_0] + (1 - d_f) \log[\eta]; \quad (7.33)$$

the slope of this relationship is related to the fractal dimension d_f of the pattern. This method is known as the caliper method to estimate the fractal dimension of a curve. It is obvious that $d_f = 1$ for a straight line.

Fractal dimension is a measure that quantifies how the given pattern fills space. The fractal dimension of a straight line is unity, that of a circle or a 2D perfectly planar (sheet-like) object is two, and that of a sphere is three. As the irregularity or complexity of a pattern increases, its fractal dimension increases up to its own Euclidean dimension d_E plus one. The fractal dimension of a jagged, rugged, convoluted, kinky, or crinkly curve will be greater than unity, and reaches the value of two as its complexity increases. The fractal dimension of a rough 2D surface will be greater than two, and approaches three as the surface roughness increases. In this sense, fractal dimension may be used as a measure of the roughness of texture in images.

Several methods have been proposed to estimate the fractal dimension of patterns [464, 465, 466, 467, 464, 468, 469]. Among the methods described by Schepers et al. [467] for the estimation of the fractal dimension of 1D signals is that of computing the relative dispersion $RD(\eta)$, defined as the ratio of the standard deviation to the mean, using varying bin size or number of samples of the signal η . For a fractal signal, the expected variation of $RD(\eta)$ is

$$RD(\eta) = RD(\eta_0) \left[\frac{\eta}{\eta_0} \right]^{H-1}, \quad (7.34)$$

where η_0 is a reference value for the bin size, and H is the Hurst coefficient that is related to the fractal dimension as

$$d_f = d_E + 1 - H. \tag{7.35}$$

(Note: $d_E = 1$ for 1D signals, 2 for 2D images, etc.) The value of H , and hence d_f , may be estimated by measuring the slope of the straight-line approximation to the relationship between $\log[RD(\eta)]$ and $\log(\eta)$.

7.5.2 Fractional Brownian motion model

Fractal signals may be modeled in terms of fractional Brownian motion [466, 467, 470]. The expectation of the differences between the values of such a signal at a position η and another at $\eta + \Delta\eta$ follow the relationship

$$E[|f(\eta + \Delta\eta) - f(\eta)|] \propto |\Delta\eta|^H. \tag{7.36}$$

The slope of a plot of the averaged difference as above versus $\Delta\eta$ (on a log – log scale) may be used to estimate H and the fractal dimension.

Chen et al. [470] applied fractal analysis for the enhancement and classification of ultrasonographic images of the liver. Burdett et al. [471] derived the fractal dimension of 2D ROIs of mammograms with masses by using the expression in Equation 7.36. Benign masses, due to their smooth and homogeneous texture, were found to have low fractal dimensions of about 2.38, whereas malignant tumors, due to their rough and heterogeneous texture, had higher fractal dimensions of about 2.56.

The PSD of a fractional Brownian motion signal $\Phi(\omega)$ is expected to follow the so-called power law as

$$\Phi(\omega) \propto \frac{1}{|\omega|^{(2H+1)}}. \tag{7.37}$$

The derivative of a signal generated by a fractional Brownian motion model is known as a fractional Gaussian noise signal; the exponent in the power-law relationship for such a signal is changed to $(2H - 1)$.

7.5.3 Fractal analysis of texture

Based upon a fractional Brownian motion model, Wu et al. [472] defined an averaged intensity-difference measure $id(k)$ for various values of the displacement or distance parameter k as

$$id(k) = \frac{1}{2N(N-k-1)} \left[\sum_{m=0}^{N-1} \sum_{n=0}^{N-k-1} |f(m, n) - f(m, n+k)| + \sum_{m=0}^{N-k-1} \sum_{n=0}^{N-1} |f(m, n) - f(m+k, n)| \right]. \tag{7.38}$$

The slope of a plot of $\log[id(k)]$ versus $\log[k]$ was used to estimate H and the fractal dimension. Wu et al. applied multiresolution fractal analysis as well GCM features, Fourier spectral features, gray-level difference statistics, and Laws' texture energy measures for the classification of ultrasonographic images of the liver as normal, hepatoma, or cirrhosis; classification accuracies of 88.9%, 83.3%, 80.0%, 74.4%, and 71.1%, respectively, were obtained. In a related study, Lee et al. [473] derived features based upon fractal analysis including the application of multiresolution wavelet transforms. Classification accuracies of 96.7% in distinguishing between normal and abnormal liver images, and 93.6% in discriminating between cirrhosis and hepatoma were obtained.

Byng et al. [448] (see also Peleg et al. [464], Yaffe et al. [474], and Caldwell et al. [475]) describe a surface-area measure to represent the complexity of texture in an image by interpreting the gray level as the height of a function of space; see Figure 7.14. In a perfectly uniform image of size $N \times N$ pixels, with each pixel being of size $\eta \times \eta$ units of area, the surface area would be equal to $(N\eta)^2$. When adjacent pixels are of unequal value, more surface area of the blocks representing the pixels will be exposed, as shown in Figure 7.14. The total surface area for the image may be calculated as

$$A(\eta) = \sum_{m=0}^{N-2} \sum_{n=0}^{N-2} \{ \eta^2 + \eta [|f_\eta(m, n) - f_\eta(m, n+1)| + |f_\eta(m, n) - f_\eta(m+1, n)|] \}, \quad (7.39)$$

where $f_\eta(m, n)$ is the 2D image expressed as a function of the pixel size η . The method is analogous to the popular box-counting method [460, 476, 469]. In order to estimate the fractal dimension of the image, we could derive several smoothed and downsampled versions of the given image (representing various scales η), and estimate the slope of the plot of $\log[A(\eta)]$ versus $\log[\eta]$; the fractal dimension is given as two minus the slope. Smoothing and downsampling may be achieved simply by averaging pixels in blocks of 2×2 , 3×3 , 4×4 , etc., and replacing the blocks by a single pixel with the corresponding average. A perfectly uniform image would demonstrate no change in its area, and have a fractal dimension of two; images with rough texture would have increasing values of the fractal dimension, approaching three. Yaffe et al. [474] obtained fractal dimension values in the range of [2.23, 2.54] with 60 mammograms. Byng et al. [448] demonstrated the usefulness of the fractal dimension as a measure of increased fibroglandular density in the breast, and related it to the risk of development of breast cancer. Fractal dimension was found to complement histogram skewness (see Section 7.3) as an indicator of breast cancer risk.




FIGURE 7.14

Computation of the exposed surface area for a pixel $f(m, n)$ with respect to its neighboring pixels $f(m, n + 1)$ and $f(m + 1, n)$. A pixel at (m, n) is viewed as a box (or building) with base area $\eta \times \eta$ and height equal to the gray level $f(m, n)$. The total exposed surface area for the pixel $f(m, n)$, with respect to its neighboring pixels at $(m, n + 1)$ and $(m + 1, n)$, is the sum of the areas of the rectangles $ABCD$, $CBEF$, and $DCGH$ [448].

7.5.4 Applications of fractal analysis

Chaudhuri and Sarkar [477] proposed a modified box-counting method to estimate fractal measures of texture from a given image, its horizontally and vertically smoothed versions, as well as high- and low-gray-valued versions derived by thresholding operations. The features were applied for the segmentation of multitextured images.

Zheng and Chan [478] used the fractal dimension of sections of mammograms to select areas with rough texture for further processing toward the detection of tumors. Pohlman et al. [407] derived the fractal dimension of 1D signatures of radial distance versus angle of boundaries of mammographic masses; the measure provided an average accuracy of 81% in discriminating between benign masses and malignant tumors. It should be noted that the function of radial distance versus angle could be multivalued for spiculated and irregular contours, due to the fact that a radial line may cross the tumor contour more than once (see Section 6.1.1). A measure related to this characteristic was found to give an accuracy of 93% in discriminating between benign masses and malignant tumors [407].

Iftekharuddin et al. [476] proposed a modified box-counting method to estimate the fractal dimension of images, and applied the method to brain MR images. Their results indicated the potential of the methods in the detection of brain tumors.

Lundahl et al. [466] estimated the values of H from scan lines of X-ray images of the calcaneus (heel) bone, and showed that the value was decreased by injury and osteoporosis, indicating reduced complexity of structure (increased

gaps) as compared to normal bone. Saporin et al. [479], using symbol dynamics and measures of complexity, found that the complexity of the trabecular structure in bone declines more rapidly than bone density during the loss of bone in osteoporosis. Jennane et al. [480] applied fractal analysis to X-ray μ CT images of trabecular bone specimens extracted from the radius. It was found that the H value decreased with trabecular bone loss and osteoporosis. Samarabandhu et al. [481] proposed a morphological filtering approach to derive the fractal dimension, and indicated that the features they derived could serve as robust measures of the trabecular texture in bone. (For a discussion on the application of fractal analysis to bone images, see Geraets and van der Stelt [469].)

Sedivy et al. [482] showed that the fractal dimension of atypical nuclei in dysplastic lesions of the cervix uteri increased as the degree of dysplasia increased. They indicated that fractal dimension could quantify the irregularity and complexity of the outlines of nuclei, and facilitate objective nuclear grading. Esgiar et al. [483] found the fractal dimension to complement the GCM texture features of entropy and correlation in the classification of tissue samples from the colon: the inclusion of fractal dimension increased the sensitivity from 90% to 95%, and the specificity from 86% to 93%. Penn and Loew [484] discussed the limitations of the box-counting and PSD-based methods in fractal analysis, and proposed fractal interpolation function models to estimate the fractal dimension. The method was shown to provide improved results in the separation of normal and sickle-cell red blood cells.

Lee et al. [429] applied shape analysis for the classification of cutaneous melanocytic lesions based upon their contours. An irregularity index related to local protrusions and indentations was observed to have a higher correlation with clinical assessment of the lesions than compactness (see Section 6.2.1) and fractal dimension.

7.6 Fourier-domain Analysis of Texture

As is evident from the illustrations in Figure 7.6, the Fourier spectrum of an image with random texture contains the spectral characteristics of the spot involved in its generation (according to the spot-noise model shown in Figure 7.4). The effects of multiplication with the spectrum of the random-noise field (which is essentially, and on the average, a constant) may be removed by smoothing operations. Thus, the important characteristics of the texture are readily available in the Fourier spectrum.

On the other hand, the Fourier spectrum of an image with periodic texture includes not only the spectral characteristics of the spot, but also the effects of multiplication with the spectrum of the impulse field involved in its generation.

The Fourier spectrum of a train of impulses in 1D is a discrete spectrum with a constant value at the fundamental frequency (the inverse of the period) and its harmonics [1, 2]. Correspondingly, in 2D, the Fourier spectrum of a periodic field of impulses is a field of impulses, with high values only at the fundamental frequency of repetition of the impulses in the image domain and integral multiples thereof. Multiplication of the Fourier spectrum of the spot with the spectrum of the impulse field will cause significant modulation of the intensities in the former, leading to bright regions at regular intervals. With real-life images, the effects of windowing or finite data, as well as of quasi-periodicity, will lead to smearing of the impulses in the spectrum of the impulse field involved; regardless, the spectrum of the image may be expected to demonstrate a field of bright regions at regular intervals. The information related to the spectra of the spot and the impulse field components may be derived from the spectrum of the textured image by averaging in the polar-coordinate axes as follows [485].

Let $F(r, t)$ be the polar-coordinate representation of the Fourier spectrum of the given image; in terms of the Cartesian frequency coordinates (u, v) , we have $r = \sqrt{u^2 + v^2}$, and $t = \text{atan}(v/u)$. Derive the projection functions in r and t by integrating $F(r, t)$ in the other coordinate as

$$F(r) = \int_{t=0}^{\pi} F(r, t) dt, \quad (7.40)$$

and


$$F(t) = \int_{r=0}^{r^{\max}} F(r, t) dr. \quad (7.41)$$

The averaging effect of the integration above (or summation in the discrete case) leads to improved visualization of the spectral characteristics of periodic texture. Quantitative features may be derived from $F(r)$ and $F(t)$ [or directly from $F(r, t)$] for pattern classification purposes.

Example: Figure 7.15 shows the components involved in the generation of an image with periodic placement of a circular spot, and the related Fourier spectra. The spectra clearly demonstrate the effects of periodicity in the impulse field and the texture. It is evident that the spectrum of the texture also includes information related to the spectrum of the spot.

The spectrum of the texture in Figure 7.15 (f) is shown in polar coordinates in Figure 7.16. The projection functions derived by summing the spectrum in the radial and angular dimensions are shown in Figure 7.17. The projection functions demonstrate the effect of periodicity in the texture.

The spectrum of the image of a fly's eye in Figure 7.3 (b) is shown in Figure 7.18 in Cartesian and polar coordinates; the corresponding projection functions are shown in Figure 7.19. A similar set of results is shown in Figures 7.20 and 7.21 for the snake-skin image in Figure 7.3 (c). The spectra show regions of high intensity at quasi-periodic intervals in spite of the fact that the original images are only approximately periodic, and the texture elements vary considerably in size and orientation over the scope of the images.

**FIGURE 7.15**

Fourier spectral characteristics of periodic texture generated using the spot-noise model in Figure 7.4. (a) Periodic impulse field. (b) Circular spot. (c) Periodic texture generated by convolving the spot in (b) with the impulse field in (a). (d) – (f) Log-magnitude Fourier spectra of the images in (a) – (c), respectively.





FIGURE 7.16


The spectrum in Figure 7.15 (f) converted to polar coordinates; only the upper half of the spectrum was mapped to polar coordinates.

Jernigan and D'Astous [486] used the normalized PSD values within selected frequency bands as PDFs, and computed entropy values. It was expected that structured texture would lead to low entropy (due to spectral bands with concentrated energy) and random texture would lead to high entropy values (due to a uniform distribution of spectral energy). Their results indicated that the entropy values could provide discrimination between texture categories that was comparable to that provided by spectral energy and GCM-based measures. In addition to entropy, the locations and values of the spectral peaks may also be used as features. Liu and Jernigan [487] defined 28 measures in the Fourier spectral domain, including measures related to the frequency coordinates and relative orientation of the first and second spectral peaks; the percentages of energy and the moments of inertia of the normalized spectrum in the first and second quadrants; the Laplacian of the magnitude and phase at the first and second spectral peaks; and measures of isotropy and circularity of the spectrum. Their results indicated that the spectral measures were effective in discriminating between various types of texture, and also insensitive to additive noise.

Laine and Fan [488] proposed the use of wavelet packet frames or tree-structured filter banks for the extraction of features from textured images in the frequency domain. The features were used for the segmentation of multitextured images.



(a)



(b)

FIGURE 7.17

Projection functions in (a) the radial coordinate r , and (b) the angle coordinate t obtained by integrating (summing) the spectrum in Figure 7.16 in the other coordinate.






FIGURE 7.18

Fourier spectral characteristics of the quasi-periodic texture of the fly's eye image in Figure 7.3 (b): (a) The Fourier spectrum in Cartesian coordinates (u, v). (b) The upper half of the spectrum in (a) mapped to polar coordinates.




(a)




(b)

FIGURE 7.19

Projection functions in (a) the radial coordinate r , and (b) the angle coordinate t obtained by integrating (summing) the spectrum in Figure 7.18 (b) in the other coordinate.




(a)




(b)

FIGURE 7.20

Fourier spectral characteristics of the ordered texture of the snake-skin image in Figure 7.3 (c): (a) The Fourier spectrum in Cartesian coordinates (u, v) . (b) The upper half of the spectrum in (a) mapped to polar coordinates.



(a)



(b)

FIGURE 7.21

Projection functions in (a) the radial coordinate r , and (b) the angle coordinate t obtained by integrating (summing) the spectrum in Figure 7.20 (b) in the other coordinate.

McLean [489] applied vector quantization in the transform domain, and treated the method as a generalized template matching scheme, for the coding and classification of texture. The method yielded better texture classification accuracy than GCM-based features.

Bovik [490] discussed multichannel narrow-band filtering and modeling of texture. Highly granular and oriented texture may be expected to present spatio-spectral regions of concentrated energy. Gabor filters may then be used to filter, segment, and analyze such patterns. See Sections 5.10.2, 7.7, 8.4, 8.9, and 8.10 for further discussion on related topics.

7.7 Segmentation and Structural Analysis of Texture

Many methods have been reported in the literature for the analysis of texture, which may be broadly classified as statistical or structural methods [441, 439]. Most of the commonly used methods for texture analysis are based upon statistical characterization, such as GCMs (see Section 7.3.1) and ACFs; Fourier spectrum analysis, described in Section 7.6, may be considered to be equivalent to analysis based upon the ACF. Statistical methods are suitable for the analysis of random or fine texture with no large-scale motifs; for other types of texture and for multitextured images, structural methods could be more appropriate. Structural analysis of textured images requires some type of segmentation of the given image into its distinct or basic components.

Texture elements (or textons, as called by Julesz and Bergen [438]) play an important role in preattentive vision and texture perception. Ordered texture may be modeled as being composed of repeated placement of a basic motif or texton over the image field in accordance with a placement rule; see Section 7.2. The placement rule may be expressed as a field of impulses indicating the locations of the repeated textons; consequently, the textured image is given by the convolution of the ordered or (quasi-) periodic impulse field with the texton. Although this model does not directly permit scale and orientation differences between the various occurrences of the texton, such “jitter” could be introduced separately to synthesize realistic textured images.

Vilnrotter et al. [491] proposed a system to describe natural textures in terms of individual texture elements or primitives and their spatial relationships or arrangement. The main steps of the system include the generation of 1D descriptors of texture elements from edge repetition data, the extraction of elements that correspond to the preceding description, the generation of 2D descriptors of each texture primitive type, and the computation of spatial arrangements or placement rules (when the texture is homogeneous and regular). The method was used to classify several types of texture including floor

grating, raffia, brick, straw, and wool. The method does not extract a single version of a texture element or primitive; instead, all possible repeated structures are extracted. The analysis of a raffia pattern, for example, resulted in the extraction of three primitives.

He and Wang [492] defined “texture units” in terms of the 8-connected neighbors of each pixel. The values in each 3×3 unit were reduced to the range $\{0, 1, 2\}$, with the value of unity indicating that the value of the neighboring pixel was within a predefined range about the central pixel value; the values of 0 and 2 were used to indicate that the pixel value was lower or higher than the specified range, respectively. The “texture spectrum” was defined as the histogram or spectrum of the frequency of occurrence of all possible texture units in the image; it should be noted that the texture spectrum as above is not based upon a linear orthogonal transform, such as the Fourier transform. Methods were proposed to characterize as well as filter images based upon the texture spectrum.

Wang et al. [493] proposed a thresholding scheme for the extraction of texture primitives. It was assumed that the primitives would appear as regions of connected pixels demonstrating good contrast with their background. The primitives were characterized in terms of the statistics of their GCMs and shape attributes; the textured image could then be described in terms of its primitives and placement rules. Tomita et al. [494] proposed a similar approach based upon the extraction of texture elements, assumed to be regions of homogeneous gray levels, via segmentation. The centroids of the texture elements were used to define detailed placement rules.

The problem of segmentation of complex images containing regions of different types of texture has been addressed by several researchers. Gabor functions have been used by Turner [495] and Bovik et al. [496] for texture analysis and segmentation. Gabor functions may be used to design filters with tunable orientation, radial frequency bandwidth, and center frequencies that can achieve jointly optimal resolution in the space and frequency domains. Gabor filters are efficient in detecting discontinuities in texture phase, and are useful in texture segmentation. Porat and Zeevi [497] developed a method to describe texture primitives in terms of Gabor elementary functions. See Sections 5.10.2, 8.4, 8.9, and 8.10 for further discussion on Gabor filters.

Reed and Wechsler [498] described approaches to texture analysis and segmentation via the use of joint spatial and frequency-domain representations in the form of spectrograms, that is, functions of (x, y, u, v) obtained by the application of the Fourier transform in a moving window of the image, a bank of Gabor filters, DoG functions, and Wigner distributions. Reed et al. [499] described a texture segmentation method using the pseudo-Wigner distribution and a diffusion region-growing method. Jain and Farrokhnia [500] presented a method for texture analysis and segmentation based upon the application of multichannel Gabor filters. The results of the filter bank were processed in such a manner as to detect “blobs” in the given image; texture discrimination was performed by analyzing the attributes of the blobs detected in

different regions of the image. Other related methods for texture segmentation include wavelet frames for the characterization of texture properties at multiple scales [501], and circular Mellin features for rotation-invariant and scale-invariant texture analysis [502].

Tardif and Zaccarin [503] proposed a multiscale autoregressive (AR) model to analyze multifeatured images. The prediction error was used to segment a given image into different textured parts. Unser and Eden [504] proposed a multiresolution feature extraction method for texture segmentation. The method includes the use of a local linear transformation that is equivalent to processing the given image with a bank of FIR filters. (See Section 8.9.4 for further discussion on related topics.)

If the texton and the placement rule (impulse field) can be obtained from a given image with ordered texture, the most important characteristics of the image will have been determined. In particular, if a single texton or motif is extracted from the image, further analysis of its shape, morphology, spectral content, and internal details becomes possible. Martins and Rangayyan [444, 505] proposed cepstral filtering in the Radon domain of the image (see Section 10.3) to obtain the texton; their methods and results are described in Section 7.7.1.

7.7.1 Homomorphic deconvolution of periodic patterns

We have seen in Section 7.2 that an image with periodic texture may be modeled as the convolution of a texton or motif with an impulse field. Linear filters may be applied to the complex cepstrum for homomorphic deconvolution of signals that contain convolved components; see Section 10.3. A basic assumption in homomorphic deconvolution is that the complex cepstra of the components do not overlap. This assumption is usually met in 1D signal processing applications, such as in the case of voiced speech signals, where the basic wavelet is a relatively smooth signal [31, 176]. Whereas it would be questionable to make the assumption that the 2D cepstra of an arbitrary texton and an impulse field do not overlap, it would be acceptable to make the same assumption in the case of 1D projections (Radon transforms) of the same images. Then, the homomorphic deconvolution procedures described in Section 10.3 may be applied to recover the projections of a single texton [444, 505]. The texton may then be obtained via a procedure for image reconstruction from projections (see Chapter 9).

The distinction between the application of homomorphic deconvolution for the removal of visual echoes as described in Section 10.3 and for the extraction of a texton is minor. An image with visual echoes may contain only one copy or a few repetitions of a basic image, with possible overlap, and with possibly unequal spacing of the echoes; the basic image may be large in spatial extent [505]. On the other hand, an image with ordered texture typically contains several nonoverlapping repetitions of a relatively small texton or motif at regular or quasi-periodic spacing.

Although homomorphic deconvolution has been shown to successfully extract the basic wavelets or motifs in periodic signals, the extraction of the impulse train or field is made difficult by the presence of noise and artifacts related to the deconvolution procedure [31].

Example: An image of a part of a building with ordered arrangement of windows is shown in Figure 7.22 (a). A single window section of the image extracted by homomorphic deconvolution is shown in part (b) of the figure.





FIGURE 7.22

(a) An image of a part of a building with a periodic arrangement of windows. (b) A single window structure extracted by homomorphic deconvolution. Reproduced with permission from A.C.G. Martins and R.M. Rangayyan, “Texture element extraction via cepstral filtering in the Radon domain”, *IETE Journal of Research* (India), 48(3,4): 143 – 150, 2002. © IETE.

An image with a periodic arrangement of a textile motif is shown in Figure 7.23 (a). The result of the homomorphic deconvolution procedure of Martins and Rangayyan [444, 505] to extract the texton is shown in part (b) of the same figure. It is evident that a single motif has been extracted, albeit with some blurring and loss of detail. The procedure, however, was not successful with biomedical images due to the effects of quasi-periodicity as well as significant size and scale variations among the repeated versions of the basic pattern. More research is desirable in this area.

**FIGURE 7.23**

(a) An image with a periodic arrangement of a textile motif. (b) A single motif or texton extracted by homomorphic deconvolution. Reproduced with permission from A.C.G. Martins and R.M. Rangayyan, “Texture element extraction via cepstral filtering in the Radon domain”, *IETE Journal of Research (India)*, 48(3,4): 143 – 150, 2002. © IETE.

7.8 Audification and Sonification of Texture in Images

The use of sound in scientific data analysis is rather rare, and analysis and presentation of data are done almost exclusively by visual means. Even when the data are the result of vibrations or sounds, such as the heart sound signals or phonocardiograms, a Doppler ultrasound exam, or sonar, they are often mapped to a graphical display or an image and visual analysis is performed.

The auditory system has not been used much for image analysis in spite of the fact that it has several advantages over the visual system. Whereas many interesting methods have been proposed for the auditory display of scientific laboratory data and computer graphics representations of multidimensional data, not much work has been reported for deriving sounds from visual images. Chambers et al. [506] published a report on auditory data presentation in the early 1970s. The first international conference on auditory display of scientific data was held in 1992 [507], with specific interest in the use of sound for the presentation and analysis of information.

Meijer [508, 509] proposed a sonification procedure to present image data to the blind. In this method, the frequency of an oscillator is associated with the position of each pixel in the image, and the amplitude is made proportional

to the pixel intensity. The image is scanned one column at a time and the outputs of the associated oscillators are all presented as a sum, followed by a click before the presentation of the next column. In essence, the image is treated as a spectrogram or a time-frequency distribution [31, 176]. The sound produced by this method with simple images such as a line crossing the plane of an image can be easily analyzed; however, the sound patterns related to complex images could be complicated and confusing.

Texture analysis is often confounded by other neighboring or surrounding features. Martins et al. [447] explored the potential of auditory display procedures, including audification and sonification, for aural presentation and analysis of texture in images. An analogy was drawn between random texture and unvoiced speech, and between periodic texture and voiced speech, in terms of generation based on the filtering of an excitation function as shown in Figure 7.4. An audification procedure that played in sequence the projections (Radon transforms) of the given image at several angles was proposed for the auditory analysis of random texture. A linear-prediction model [510, 176, 31] was used to generate the sound signal from the projection data. Martins et al. also proposed a sonification procedure to convert periodic texture to sound, with the emphasis on displaying the essential features of the texture element and periodicity in the horizontal and vertical directions. Projections of the texton were used to compose sound signals including pitch like voiced speech as well as a rhythmic aspect, with the pitch period and rhythm related to the periodicities in the horizontal and vertical directions in the image. Data-mapping functions were designed to relate image characteristics to sound parameters in such a way that the sounds provided information in microstructure (timbre, individual pitch) and macrostructure (rhythm, melody, pitch organization) that were related to the objective or quantitative measures of texture.

In order to verify the potential of the proposed methods for aural analysis of texture, a set of pilot experiments was designed and presented to 10 subjects [447]. The results indicated that the methods could facilitate qualitative and comparative analysis of texture. In particular, it was observed that the methods could lead to the possibility of defining a sequence or order in the case of images with random texture, and that sound-to-image association could be achieved in terms of the size and shape of the spot used to synthesize the texture. Furthermore, the proposed mapping of the attributes of periodic texture to sound attributes could permit the analysis of features such as texton size and shape, as well as periodicity in qualitative and comparative manners. The methods could lead to the use of auditory display of images as an adjunctive procedure to visualization.

Martins et al. [511] conducted preliminary tests on the audification of MR images using selected areas corresponding to the gray and white matter of the brain, and to normal and infarcted tissues. By using the audification method, differences between the various tissue types were easily perceived by two radiologists; visual discrimination of the same areas while remaining

within their corresponding MR-image contexts was said to be difficult by the same radiologists. The results need to be confirmed with a larger study.

7.9 Application: Analysis of Breast Masses Using Texture and Gradient Measures

In addition to the textural changes caused by microcalcifications, the presence of spicules arising from malignant tumors causes disturbances in the homogeneity of tissues in the surrounding breast parenchyma. Based upon this observation, several studies have focused on quantifying the textural content in the mass ROI and mass margins to achieve the classification of masses versus normal tissue as well as benign masses versus malignant tumors.

Petrosian et al. [446] investigated the usefulness of texture features based upon GCMs for the classification of masses and normal tissue. With a dataset of 135 manually segmented ROIs, the methods indicated 89% sensitivity and 76% specificity in the training step, and 76% sensitivity and 64% specificity in the test step using the leave-one-out method. Kinoshita et al. [512] used a combination of shape factors and texture features based on GCMs. Using a three-layer feed-forward neural network, they reported 81% accuracy in the classification of benign and malignant breast lesions with a dataset of 38 malignant and 54 benign lesions.

Chan et al. [450], Sahiner et al. [451, 513], and Wei et al. [514, 515] investigated the effectiveness of texture features derived from GCMs for differentiating masses from normal breast tissue in digitized mammograms. One hundred and sixty-eight ROIs with masses and 504 normal ROIs were examined, and eight features including correlation, entropy, energy, inertia, inverse difference moment, sum average moment, sum entropy, and difference entropy were calculated for each region. All the ROIs were manually segmented by a radiologist. Using linear discriminant analysis, Chan et al. [450] reported an accuracy of 0.84 for the training set and 0.82 for a test set. Wei et al. [514, 515] reported improved classification results with the same dataset by applying multiresolution texture analysis. Sahiner et al. applied a convolutional neural network [513], and later used a genetic algorithm [513, 516] to classify the masses and normal tissue in the same dataset.

Analysis of the gradient or transition information present in the boundaries of masses has been attempted by a few researchers in order to arrive at benign-versus-malignant decisions. Kok et al. [517] used texture features, fractal measures, and edge-strength measures computed from suspicious regions for lesion detection. Huo et al. [518] and Giger et al. [519] extracted mass regions using region-growing methods and proposed two spiculation measures obtained from an analysis of radial edge-gradient information surrounding

the periphery of the extracted regions. Benign-versus-malignant classification studies performed using the features yielded an average efficiency of 0.85. Later on, the group reported to have achieved superior results with their computer-aided classification scheme as compared to an expert radiologist by employing a hybrid classifier on a test set of 95 images [520].

Highnam et al. [521] investigated the presence of a “halo” — an area around a mass region with a positive Laplacian — to indicate whether a circumscribed mass is benign or malignant. They found that the extent of the halo varies between the CC and MLO views for benign masses, but is similar for malignant tumors.

Guliatto et al. [276, 277] proposed fuzzy region-growing methods for segmenting breast masses, and further proposed classification of the segmented masses as benign or malignant based on the transition information present around the segmented regions; see Sections 5.5 and 5.11. Rangayyan et al. [163] proposed a region-based edge-profile acutance measure for evaluating the sharpness of mass boundaries; see Sections 2.15 and 7.9.2.

Many studies have focused on transforming the space-domain intensities into other forms for analyzing gradient and texture information. Claridge and Richter [522] developed a Gaussian blur model to characterize the transitional information in the boundaries of mammographic lesions. In order to analyze the blur in the boundaries and to determine the prevailing direction of linear patterns, a polar coordinate transform was applied to map the lesion into polar coordinates. A measure of spiculation was computed from the transformed images to discriminate between circumscribed and spiculated lesions as the ratio of the sum of vertical gradient magnitudes to the sum of horizontal gradient magnitudes.

Sahiner et al. [451, 516, 523] introduced the RBST method to transform a band of pixels surrounding the boundary of a segmented mass onto the Cartesian plane (see Figure 7.26). The band of pixels was extracted in the perpendicular direction from every point on the boundary. Texture features based upon GCMs computed from the RBST images resulted in an average efficiency of 0.94 in the benign-versus-malignant classification of 168 cases. Sahiner et al. reported that texture analysis of RBST images yielded better benign-versus-malignant discrimination than analysis of the original space-domain images. However, such a transformation is sensitive to the precise extraction of the band of pixels surrounding the ROI; the method may face problems with masses having highly spiculated margins.

Hadjiiski et al. [524] reported on the design of a hybrid classifier — adaptive resonance theory network cascaded with linear discriminant analysis — to classify masses as benign or malignant. They compared the performance of the hybrid classifier that they designed with a back-propagation neural network and linear discriminant classifiers, using a dataset of 348 manually segmented ROIs (169 benign and 179 malignant). Benign-versus-malignant classification using the hybrid classifier achieved marginal improvement in performance, with an average efficiency of 0.81. The texture features used in the classifier

were based upon GCMs and run-length sequences computed from the RBST images.


Giger et al. [525] classified manually delineated breast mass lesions in ultrasonographic images as benign or malignant using texture features, margin sharpness, and posterior acoustic attenuation. With a dataset of 135 ultrasound images from 39 patients, the posterior acoustic attenuation feature achieved the best benign-versus-malignant classification results, with an average efficiency of 0.84. Giger et al. reported to have achieved higher sensitivity and specificity levels by combining the features derived from both mammographic and ultrasonographic images of mass lesions as against using features computed from only the mammographic mass lesions.

Mudigonda et al. [275, 165] derived measures of texture and gradient using ribbons of pixels around mass boundaries, with the hypothesis that the transitional information in a mass margin from the inside of the mass to its surrounding tissues is important in discriminating between benign masses and malignant tumors. The methods and results of this work are described in the following sections. See Sections 6.7, 8.8, 12.11, and 12.12 for more discussion on the detection and analysis of breast masses.


7.9.1 Adaptive normals and ribbons around mass margins

Mudigonda et al. [165, 275] obtained adaptive ribbons around boundaries of breast masses and tumors that were drawn by an expert radiologist, in the following manner. Morphological dilation and erosion operations [526] were applied to the boundary using a circular operator of a specified diameter. Figures 7.24 and 7.25 show the extracted ribbons across the boundaries of a benign mass and a malignant tumor, respectively. The width of the ribbon in each case is 8 mm across the boundary (4 mm or 80 pixels on either side of the boundary at a resolution of 50 μm per pixel). The ribbon width of 8 mm was determined by a radiologist in order to take into account the possible depth of infiltration or diffusion of masses into the surrounding tissues.


In order to compute gradient-based measures and acutance (see Section 2.15), Mudigonda et al. developed the following procedure to extract pixels from the inside of a mass boundary to the outside along the perpendicular direction at every point on the boundary. A polygonal model of the mass boundary, computed as described in Section 6.1.4, was used to approximate the mass boundary with a polygon of known parameters. With the known equations of the sides of the polygonal model, it is possible to estimate the normal at every point on the boundary. The length of the normal at any point on the boundary was limited to a maximum of 80 pixels (4 mm) on either side of the boundary or the depth of the mass at that particular point. This is significant, especially in the case of spiculated tumors possessing sharp spicules or microlobulations, such that the extracted normals do not cross over into adjacent spicules or mass portions. The normals obtained as above for a benign mass and a malignant tumor are shown in Figures 7.24 and 7.25.



(a)




(b)




(c)

FIGURE 7.24


(a) A $1,000 \times 900$ section of a mammogram containing a circumscribed benign mass. Pixel size = $50 \mu m$. (b) Ribbon or band of pixels across the boundary of the mass extracted by using morphological operations. (c) Pixels along the normals to the boundary, shown for every tenth boundary pixel. Maximum length of the normals on either side of the boundary = 80 pixels or $4 mm$. Images courtesy of N.R. Mudigonda [166]. See also Figure 12.28.



(a)



(b)



(c)

FIGURE 7.25

(a) A 630×560 section of a mammogram containing a spiculated malignant tumor. Pixel size = $50 \mu m$. (b) Ribbon or band of pixels across the boundary of the tumor extracted by using morphological operations. (c) Pixels along the normals to the boundary, shown for every tenth boundary pixel. Maximum length of the normals on either side of the boundary = 80 pixels or $4 mm$. Images courtesy of N.R. Mudigonda [166]. See also Figure 12.28.

With an approach that is different from the above but comparable, Sahiner et al. [451] formulated the RBST method to map ribbons around breast masses in mammograms into rectangular arrays; see Figure 7.26. It was expected that variations in texture due to the spicules that are commonly present around malignant tumors would be enhanced by the transform, and lead to better discrimination between malignant tumors and benign masses. The rectangular array permitted easier and straightforward computation of texture measures.

7.9.2 Gradient and contrast measures

Due to the infiltration into the surrounding tissues, malignant breast lesions often permeate larger areas than apparent on mammograms. As a result, tumor margins in mammographic images do not present a clear-cut transition or reliable gradient information. Hence, it is difficult for an automated detection procedure to realize precisely the boundaries of mammographic masses, as there cannot be any objective measure of such precision. Furthermore, when manual segmentation is used, there are bound to be large inter-observer variations in the location of mass boundaries due to subjective differences in notions of edge sharpness. Considering the above, it is appropriate for gradient-based measures to characterize the global gradient phenomenon in the mass margins without being sensitive to the precise location of the mass boundary.

A modified measure of edge sharpness: The subjective impression of sharpness perceived by the HVS is a function of the averaged variations in intensities between the relatively light and dark areas of an ROI. Based upon this, Higgins and Jones [115] proposed a measure of acutance to compute sharpness as the mean-squared gradient along knife-edge spread functions of photographic films. Rangayyan and Elkadiki [116] extended this concept to 2D ROIs in images; see Section 2.15 for details. Rangayyan et al. [163] used the measure to classify mammographic masses as benign or malignant: acutance was computed using directional derivatives along the perpendicular at every boundary point by considering the inside-to-outside differences of intensities across the boundary normalized to unit pixel distance. The method has limitations due to the following reasons:

- Because derivatives were computed based on the inside-to-outside differences across the boundary, the measure is sensitive to the actual location of the boundary. Furthermore, it is sensitive to the number of differences (pixel pairs) that are available at a particular boundary point, which could be relatively low in the sharply spiculated portions of a malignant tumor as compared to the well-circumscribed portions of a benign mass. The measure thus becomes sensitive to shape complexity as well, which is not intended.
- The final acutance value for a mass ROI was obtained by normalizing the mean-squared gradient computed at all the points on the boundary

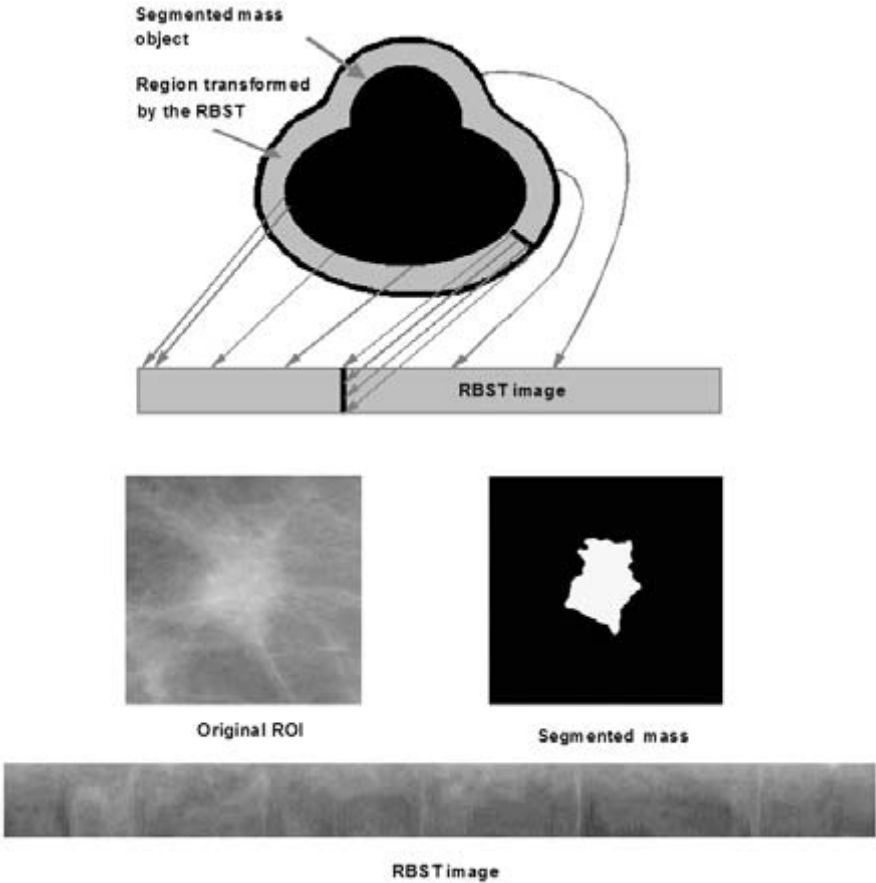


FIGURE 7.26

Mapping of a ribbon of pixels around a mass into a rectangular image by the rubber-band straightening transform [428, 451]. Figure courtesy of B. Sahiner, University of Michigan, Ann Arbor, MI. Reproduced with permission from B.S. Sahiner, H.P. Chan, N. Petrick, M.A. Helvie, and M.M. Goodsitt, “Computerized characterization of masses on mammograms: The rubber band straightening transform and texture analysis”, *Medical Physics*, 25(4): 516 – 526, 1995. © American Association of Medical Physicists.

with a factor dependent upon the maximum gray-level range and the maximum number of differences used in the computation of acutance. For a particular mass under consideration, this type of normalization could result in large differences in acutance values for varying numbers of pixel pairs considered.

Mudigonda et al. [165] addressed the above-mentioned drawbacks by developing a consolidated measure of directional gradient strength as follows. Given the boundary of a mass formed by N points, the first step is to compute the RMS gradient in the perpendicular direction at every point on the boundary with a set of successive pixel pairs as made available by the ribbon-extraction method explained in Section 7.9.1. The RMS gradient d_m at the m^{th} boundary point is obtained as

$$d_m = \sqrt{\frac{\sum_{n=0}^{(p_m-1)} [f_m(n) - f_m(n+1)]^2}{p_m}}, \quad (7.42)$$

where $f_m(n)$, $n = 0, 1, 2, \dots, p_m$, are the $(p_m + 1)$ pixels available along the perpendicular at the m^{th} boundary point, including the boundary point. The normal p_m is limited to a maximum of 160 pixels (80 pixels on either side of the boundary, with the pixel size being $50 \mu\text{m}$).

A modified measure of acutance based on the directional gradient strength A_g of the ROI is computed as

$$A_g = \frac{1}{N (f_{\max} - f_{\min})} \sum_{m=1}^N d_m, \quad (7.43)$$

where f_{\max} and f_{\min} are the local maximum and the local minimum pixel values in the ribbon of pixels extracted, and N is the number of pixels along the boundary of the ROI. Because RMS gradients computed over several pixel pairs at each boundary point are used in the computation of A_g , the measure is expected to be stable in the presence of noise, and furthermore, expected to be not sensitive to the actual location of the boundary. The factor $(f_{\max} - f_{\min})$ in the denominator in Equation 7.43 serves as an additional normalization factor in order to account for the changes in the gray-level contrast of images from various databases; it also normalizes the A_g measure to the range $[0, 1]$.

Coefficient of variation of gradient strength: In the presence of objects with fuzzy backgrounds, as is the case in mammographic images, the mean-squared gradient as a measure of sharpness may not result in adequate confidence intervals for the purposes of pattern classification. Hence, statistical measures need to be adopted to characterize the feeble gradient variations across mass margins. Considering this notion, Mudigonda et al. [165] proposed a feature based on the coefficient of variation of the edge-strength values computed at all points on a mass boundary. The stated purpose of this

feature was to investigate the variability in the sharpness of a mass around its boundary, in addition to the evaluation of its average sharpness with the measure A_g . Variance is a statistical measure of signal strength, and can be used as an edge detector because it responds to boundaries between regions of different brightness [527]. In the procedure proposed by Mudigonda et al., the variance (σ_w^2) localized in a moving window of an odd number of pixels (M) in the perpendicular direction at a boundary pixel is computed as

$$\sigma_w^2 = \frac{1}{M} \sum_{n=\lfloor -M/2 \rfloor}^{\lfloor M/2 \rfloor} [f_m(n) - \mu_w]^2, \quad (7.44)$$

where $M = 5$; $f_m(n)$, $n = 0, 1, 2, \dots, p_m$, are the pixels considered at the m^{th} boundary point in the perpendicular direction; and μ_w is the running mean intensity in the selected window:

$$\mu_w = \frac{1}{M} \sum_{n=\lfloor -M/2 \rfloor}^{\lfloor M/2 \rfloor} f_m(n). \quad (7.45)$$

The window is moved over the entire range of pixels made available at a particular boundary point by the ribbon-extraction method described in Section 7.9.1. The maximum of the variance values thus computed is used to represent the edge strength at the boundary point being processed. The coefficient of variation (G_{cv}) of the edge-strength values for all the points on the boundary is then computed. The measure is not sensitive to the actual location of the boundary within the selected ribbon, and is normalized so as to be applicable to a mixture of images from different databases.

7.9.3 Results of pattern classification

In the work of Mudigonda et al. [165, 166], four GCMs were constructed by scanning each mass ROI or ribbon in the 0° , 45° , 90° , and 135° directions with unit-pixel distance ($d = 1$). Five of Haralick's texture features, defined as F_1, F_2, F_3, F_5 , and F_9 in Section 7.3.2, were computed for the four GCMs, thus resulting in a total of 20 texture features for each ROI or ribbon.

A pixel distance of $d = 1$ is preferred to ensure large numbers of co-occurrences derived from the ribbons of pixels extracted from mass margins. Texture features computed from GCMs constructed for larger distances ($d = 3, 5$, and 10 pixels, with the resolution of the images being 50 or $62 \mu\text{m}$) were found to possess a high degree of correlation (0.9 and higher) with the corresponding features computed for unit-pixel distance ($d = 1$). Hence, pattern classification experiments were not carried out with the GCMs constructed using larger distances.

In addition to the texture features described above, the two gradient-based features A_g and G_{cv} were computed from adaptive ribbons extracted around

the boundaries of 53 mammographic ROIs, including 28 benign masses and 25 malignant tumors. Three leading features (with canonical coefficients greater than 1) including two texture measures of correlation ($d = 1$ at 90° and 45°), and a measure of inverse difference moment ($d = 1$ at 0°) were selected from the 20 texture features computed from the ribbons. The classification accuracy was found to be the maximum with the three features listed above. The two most-effective features selected for analyzing the mass ROIs included two measures of correlation ($d = 1$ at 90° and 135°).

Pattern classification experiments with 38 masses from the MIAS database [376] (28 benign and 10 malignant) indicated average accuracies of 68.4% and 78.9% using the texture features computed with the entire mass ROIs and the adaptive ribbons around the boundaries, respectively. This result supports the hypothesis that discriminant information is contained around the margins of breast masses rather than within the masses. With the extended database of 53 masses (28 benign and 25 malignant), and with features computed using the ribbons around the boundaries, the classification accuracies with the gradient and texture features as well as their combination were 66%, 75.5%, and 73.6%, respectively. (The area under the receiver operating characteristics curves were, respectively, 0.71, 0.80, and 0.81; see Section 12.8.1 for details on this method.) The gradient features were observed to increase the sensitivity, but reduce the specificity when combined with the texture features.

In a different study, Alto et al. [528, 529] obtained benign-versus-malignant classification accuracies of up to 78.9% with acutance (as in Equation 2.110), 66.7% with Haralick's texture measures, and 98.2% with shape factors applied to a different database of 57 breast masses and tumors. Although combinations of the features did not result in higher pattern classification accuracy, advantages were observed in experiments on content-based retrieval (see Section 12.12).

In experiments conducted by Sahiner et al. [428] with automatically extracted boundaries of 249 mammographic masses, Haralick's texture measures individually provided classification accuracies of up to only 0.66, whereas the Fourier-descriptor-based shape factor defined in Equation 6.58 gave an accuracy of 0.82 (the highest among 13 shape features, 13 texture features, and five run-length statistics). Each texture feature was computed using the RBST method [451] (see Figure 7.26) in four directions and for 10 distances. However, the full set of the shape factors provided an average accuracy of 0.85, the texture feature set provided the same accuracy, and the combination of shape and texture feature sets provided an improved accuracy of 0.89. These results indicate the importance of including features from a variety of perspectives and image characteristics in pattern classification.

See Sections 5.5, 5.11, and 8.8 for discussions on the detection of masses in mammograms; Sections 6.7 and 12.11 for details on shape analysis of masses; and Section 12.12 for a discussion on the application of texture measures for content-based retrieval and classification of mammographic masses.

7.10 Remarks

In this chapter, we have examined the nature of texture in biomedical images, and studied several methods to characterize texture. We have also noted numerous applications of texture analysis in the classification of biomedical images. Depending upon the nature of the images on hand, and the anticipated textural differences between the various categories of interest, one may have to use combinations of several measures of texture and contour roughness (see Chapter 6) in order to obtain acceptable results. Relating statistical and computational representations of texture to visually perceived patterns or expert opinion could be a significant challenge in medical applications. See Ojala et al. [530] for a comparative analysis of several methods for the analysis of texture. See Chapter 12 for examples of pattern classification via texture analysis.

Texture features may also be used to partition or segment multitextured images into their constituent parts, and to derive information regarding the shape, orientation, and perspective of objects. Haralick and Shapiro [440] (Chapter 9) describe methods for the derivation of the shape and orientation of 3D objects or terrains via the analysis of variations in texture.

Examples of oriented texture were presented in this chapter. Given the importance of oriented texture and patterns with directional characteristics in biomedical images, Chapter 8 is devoted completely to the analysis of oriented patterns.

7.11 Study Questions and Problems

Selected data files related to some of the problems and exercises are available at the site

www.enel.ucalgary.ca/People/Ranga/enel697

1. Explain the manner in which
 - (a) the variance,
 - (b) the entropy, and
 - (c) the skewnessof the histogram of an image can represent texture. Discuss the limitations of measures derived from the histogram of an image in the representation of texture.
2. What are the main similarities and differences between the histogram and a gray-level co-occurrence matrix of an image?
What are the orders of these two measures in terms of PDFs?

3. Explain why gray-level co-occurrence matrices need to be estimated for several values of displacement (distance) and angle.
 4. Explain how shape complexity and texture (gray-level) complexity complement each other. (You may use a tumor as an example.)
 5. Sketch two examples of fractals, in the sense of self-similar nested patterns, in biomedical images.
-

7.12 Laboratory Exercises and Projects

1. Visit a medical imaging facility and a pathology laboratory. Collect examples of images with
 - (a) random texture,
 - (b) oriented texture, and
 - (c) ordered texture.

Respect the priority, privacy, and confidentiality of patients.

Request a radiologist, a technologist, or a pathologist to explain how he or she interprets the images. Obtain information on the differences between normal and abnormal (disease) patterns in different types of samples and tests.

Collect a few sample images for use in image processing experiments, after obtaining the necessary permissions and ensuring that you carry no patient identification out of the laboratory.

2. Compute the log-magnitude Fourier spectra of the images you obtained in Exercise 1. Study the nature of the spectra and relate their characteristics to the nature of the texture observed in the images.
3. Derive the histograms of the images you obtained in Exercise 1. Compute the
 - (a) the variance,
 - (b) the entropy,
 - (c) the skewness, and
 - (d) kurtosis

of the histograms. Relate the characteristics of the histograms and the values of the parameters listed above to the nature of the texture observed in the images.

4. Write a program to estimate the fractal dimension of an image using the method given by Equation 7.39. Compute the fractal dimension of the images you obtained in Exercise 1. Interpret the results and relate them to the nature of the texture observed in the images.

Analysis of Oriented Patterns

Many images are composed of piecewise linear objects. Linear or oriented objects possess directional coherence that can be quantified and examined to assess the underlying pattern. An area that is closely related to directional image processing is texture identification and segmentation. For example, given an image of a human face, a method for texture segmentation would attempt to separate the region consisting of hair from the region with skin, as well as other regions such as the eyes that have a texture that is different from that of either the skin or hair. In texture segmentation, a common approach for identifying the differing regions is via finding the dominant orientation of the different texture elements, and then segmenting the image using this information. The subject matter of this chapter is more focused, and concerned with issues of whether there is coherent structure in regions such as the hair or skin. To put it simply, the question is whether the hair is combed or not, and if it is not, the degree of disorder is of interest, which we shall attempt to quantify. Directional analysis is useful in the effective identification, segmentation, and characterization of oriented (or weakly ordered) texture [432].

8.1 Oriented Patterns in Images

In most cases of natural materials, strength is derived from highly coherent, oriented fibers; an example of such structure is found in ligaments [35, 36]. Normal, healthy ligaments are composed of bundles of collagen fibrils that are coherently oriented along the long axis of the ligament; see Figure 1.8 (a). Injured and healing ligaments, on the other hand, contain scabs of scar material that are not aligned. Thus, the determination of the relative disorder of collagen fibrils could provide a direct indicator of the health, strength, and functional integrity (or lack thereof) of a ligament [35, 36, 37, 531]; similar patterns exist in other biological tissues such as bones, muscle fibers, and blood vessels in ligaments as well [414, 415, 532, 533, 534, 535, 536, 537, 538, 539, 540, 541, 542, 543].

Examples of oriented patterns in biomedical images include the following:

- Fibers in muscles and ligaments; see Figure 8.22.

- Fibroglandular tissue, ligaments, and ducts in the breast; see Figures 7.2 and 8.66.
- Vascular networks in ligaments, lungs, and the heart; see Figures 9.20 and 8.27.
- Bronchial trees in the lungs; see Figure 7.1.

Several more examples are presented in the sections to follow.

In man-made materials such as paper and textiles, strength usually relies upon the individual fibers uniformly knotting together. Thus, the strength of the material is directly related to the organization of the individual fibril strands [544, 545, 546, 547, 548, 549].

Oriented patterns have been found to bear significant information in several other applications of imaging and image processing. In geophysics, the accurate interpretation of seismic soundings or “stacks” is dependent upon the elimination of selected linear segments from the stacks, primarily the “ground roll” or low-frequency component of a seismic sounding [550, 551, 552]. Thorarinnsson et al. [553] used directional analysis to discover linear anomalies in magnetic maps that represent tectonic features.

In robotics and computer vision, the detection of the objects in the vicinity and the determination of their orientation relative to the robot are important in order for the machine to function in a nonstandard environment [554, 555, 556]. By using visual cues in images, such as the dominant orientation of a scene, robots may be enabled to identify basic directions such as up and down.

Information related to orientation has been used in remote sensing to analyze satellite maps for the detection of anomalies in map data [557, 558, 559, 560, 561, 562]. Underlying structures of the earth are commonly identified by directional patterns in satellite images; for example, ancient river beds [557]. Identifying directional patterns in remotely sensed images helps geologists to understand the underlying processes in the earth that are in action [553, 562]. Because man-made structures also tend to have strong linear segments, directional features can help in the identification of buildings, roads, and urban features [561].

Images commonly have sharp edges that make them nonstationary. Edges render image coding and compression techniques such as LP coding and DPCM (see Chapter 11) less efficient. By dividing the frequency space into directional bands that contain the directional image components in each band, and then coding the bands separately, higher rates of compression may be obtained [563, 564, 565, 566, 567, 568, 569]. In this manner, directional filtering can be useful in other applications of image processing, such as data compression.

8.2 Measures of Directional Distribution

Mardia [570] pointed out that the statistical measures that are commonly used for the analysis of data points in rectangular coordinate systems may lead to improper results if applied to circular or directional data. Because we do not usually consider directional components in images to be directed elements (or vectors), there should be no need to differentiate between components that are at angles θ and $\theta \pm 180^\circ$; therefore, we could limit our analysis to the semicircular space of $[0^\circ, 180^\circ]$ or $[-90^\circ, 90^\circ]$.

8.2.1 The rose diagram

The rose diagram is a graphical representation of directional data. Corresponding to each angular interval or bin, a sector (a petal of the rose) is plotted with its apex at the origin. In common practice, the radius of the sector is made proportional to the area of the image components directed in the corresponding angle band.

The area of each sector in a rose diagram as above varies in proportion to the square of the directional data. In order to make the areas of the sectors directly proportional to the orientation data, the square roots of the data elements could be related to the radii of the sectors. Linear histograms conserve areas and are comparatively simple to construct; however, they lack the strong visual association with directionality that is obtained through the use of rose diagrams. Several examples of rose diagrams are provided in the sections to follow.

8.2.2 The principal axis

The spatial moments of an image may be used to determine its principal axis, which could be helpful in finding the dominant angle of directional alignment. The moment of inertia of an image $f(x, y)$ is at its minimum when the moment is taken about the centroid (\bar{x}, \bar{y}) of the image. The moment of inertia of the image about the line $(y - \bar{y}) \cos \theta = (x - \bar{x}) \sin \theta$ passing through (\bar{x}, \bar{y}) and having the slope $\tan \theta$ is given by

$$m_\theta = \int_x \int_y [(x - \bar{x}) \sin \theta - (y - \bar{y}) \cos \theta]^2 f(x, y) dx dy. \quad (8.1)$$

In order to make m_θ independent of the choice of the coordinates, the centroid of the image could be used as the origin. Then, $\bar{x} = 0$ and $\bar{y} = 0$, and Equation 8.1 becomes

$$m_\theta = \int_x \int_y (x \sin \theta - y \cos \theta)^2 f(x, y) dx dy$$

$$= m_{20} \sin^2 \theta - 2 m_{11} \sin \theta \cos \theta + m_{02} \cos^2 \theta, \quad (8.2)$$

where m_{pq} is the $(p, q)^{\text{th}}$ moment of the image, given by

$$m_{pq} = \int_x \int_y x^p y^q f(x, y) dx dy. \quad (8.3)$$

By definition, the moment of inertia about the principal axis is at its minimum. Differentiating Equation 8.2 with respect to θ and equating the result to zero gives

$$m_{20} \sin 2\theta - 2 m_{11} \cos 2\theta - m_{02} \sin 2\theta = 0, \quad (8.4)$$

or

$$\tan 2\theta = \frac{2 m_{11}}{(m_{20} - m_{02})}. \quad (8.5)$$

By solving this equation, we can find the slope or the direction of the principal axis of the given image [11].

If the input image consists of directional components along an angle ϕ only, then $\phi \approx \theta$. If there are a number of directional components at different angles, then θ represents their weighted average direction. Evidently, this method cannot detect the existence of components in various angle bands, and is thus inapplicable for the analysis of multiple directional components. Also, this method cannot quantify the directional components in various angle bands.

8.2.3 Angular moments

The angular moment M_k of order k of an angular distribution is defined as

$$M_k = \sum_{n=1}^N \theta^k(n) p(n), \quad (8.6)$$

where $\theta(n)$ represents the center of the n^{th} angle band in degrees, $p(n)$ represents the normalized weight or probability of the data in the n^{th} band, and N is the number of angle bands. If we are interested in determining the dispersion of the angular data about their principal axis, the moments may be taken with respect to the centroidal angle $\bar{\theta} = M_1$ of the distribution. Because the second-order moment is at its minimum when taken about the centroid, we could choose $k = 2$ for statistical analysis of angular distributions. Hence, the second central moment M_2 may be defined as

$$M_2 = \sum_{n=1}^N [\theta(n) - \bar{\theta}]^2 p(n). \quad (8.7)$$

The use of M_2 as a measure of angular dispersion has a drawback: because the moment is calculated using the product of the square of the angular distance and the weight of the distribution, even a small component at a large

angular distance from the centroidal angle could result in a high value for M_2 . (See also Section 6.2.2.)

8.2.4 Distance measures

The directional distribution obtained by a particular method for an image may be represented by a vector $\mathbf{p}_1 = [p_1(1), p_1(2), \dots, p_1(N)]^T$, where $p_1(n)$ represents the distribution in the n^{th} angle band. The true distribution of the image, if known, may be represented by another vector \mathbf{p}_0 . Then, the Euclidean distance between the distribution obtained by the directional analysis method \mathbf{p}_1 and the true distribution of the image \mathbf{p}_0 is given as

$$\|\mathbf{p}_1 - \mathbf{p}_0\| = \sqrt{\sum_{n=1}^N [p_1(n) - p_0(n)]^2}. \quad (8.8)$$

This distance measure may be used to compare the accuracies of different methods of directional analysis.

Another distance measure that is commonly used is the Manhattan distance, defined as

$$|\mathbf{p}_1 - \mathbf{p}_0| = \sum_{n=1}^N |p_1(n) - p_0(n)|. \quad (8.9)$$

The distance measures defined above may also be used to compare the directional distribution of one image with that of another.

8.2.5 Entropy

The concept of entropy from information theory [127] (see Section 2.8) can be effectively applied to directional data. If we take $p(n)$ as the directional PDF of an image in the n^{th} angle band, the entropy H of the distribution is given by

$$H = - \sum_{n=1}^N p(n) \log_2[p(n)]. \quad (8.10)$$

Entropy provides a useful measure of the scatter of the directional elements in an image. If the image is composed of directional elements with a uniform distribution (maximal scatter), the entropy is at its maximum; if, however, the image is composed of directional elements oriented at a single angle or in a narrow angle band, the entropy is (close to) zero. Thus, entropy, while not giving the angle band of primary orientation or the principal axis, could give a good indication of the directional spread or scatter of an image [35, 36, 414, 415]. (See Figure 8.24.)

Other approaches that have been followed by researchers for the characterization of directional distributions are: numerical and statistical characterization of directional strength [535], morphological operations using a rotating

structural element [541], laser small-angle light scattering [538, 539, 549], and optical diffraction and Fourier analysis [532, 548, 558, 560].

8.3 Directional Filtering

Methods based upon the Fourier transform have dominated the area of directional image processing [36, 532, 550, 551, 552]. The Fourier transform of an oriented linear segment is a sinc function oriented in the direction orthogonal to that of the original segment in the spatial domain; see Figure 8.1. Based upon this property, we can design filters to select linear components at specific angles. However, a difficulty in using the Fourier domain for directional filtering lies in the development of high-quality filters that are able to select linear components without the undesirable effects of ringing in the spatial domain.

Schiller et al. [571] showed that the human eye contains orientation-selective structures. This motivated research on human vision by Marr [282], who showed that the orientation of linear segments, primarily edges, is important in forming the *primal sketch*. Several researchers, including Kass and Witkin [572], Zucker [573], and Low and Coggins [574] used oriented bandpass filters in an effort to simulate the human visual system's ability to identify oriented structures in images. Allen et al. [575] developed a very-large-scale integrated (VLSI) circuit implementation of an orientation-specific "retina".

Several researchers [36, 572, 573, 574] have used many types of simple filters with wide passbands at various angles to obtain a redundant decomposition or representation of the given image. Such representations were used to derive directional properties of the image. For example, Kass and Witkin [572] formed a map of flow lines in the given image, and under conformal mapping, obtained a transformation to regularize the flow lines onto a grid. The resulting transformation was used as a parameter representing the texture of the image. In this manner, various types of texture could be recognized or generated by using the conformal map specific to the texture.

Chaudhuri et al. [36] used a set of bandpass filters to obtain directional components in SEM images of ligaments; however, the filter used was relatively simple (see Sections 8.3.1 and 8.7.1). Generating highly selective filters in 2D is not trivial, and considerable research has been directed toward finding general rules for the formation of 2D filters. Bigün et al. [576] developed rules for the generation of least-squares optimal beam filters in multiple dimensions. Bruton et al. [577] developed a method for designing high-quality fan filters using methods from circuit theory. This method results in 2D recursive filters that have high directional selectivity and good roll-off characteristics, and is described in Section 8.3.3.




FIGURE 8.1

(a) A test image with a linear feature. (b) Log-magnitude Fourier spectrum of the test image in (a). (c) Another test image with a linear feature at a different angle. (d) Log-magnitude Fourier spectrum of the test image in (b). See also Figure 2.30.

8.3.1 Sector filtering in the Fourier domain


Fourier-domain techniques are popular methods for directional quantification of images [36, 532, 547, 550, 551, 552, 553, 557, 558, 559, 560, 562, 564, 565, 566, 567, 568, 569, 577, 578, 579, 580, 581, 582, 583, 584, 585, 586, 587, 588, 589]. The results of research on biological visual systems provide a biological base for directional analysis of images using filter-based methods [389, 563, 571, 572, 573, 575].

The Fourier transform is the most straightforward method for identifying linear components. The Fourier transform of a line segment is a sinc function oriented at $\pi/2$ radians with respect to the direction of the line segment in the spatial domain; see Figure 8.1. This fact allows the selective filtering of line segments at a specific orientation by filtering the transformed image with a bandpass filter.

Consider a line segment of orientation (slope) a and y -axis intercept b in the (x, y) plane, with the spatial limits $[-X, X]$ and $[-Y, Y]$. In order to obtain the Fourier transform of the image, we could evaluate a line integral in 2D along the line $y = ax + b$. To simplify the procedure, let us assume that the integration occurs over a square region with $X = Y$. Because the function $f(x, y)$ is a constant along the line, the term $f(x, y)$ in the Fourier integral can be normalized to unity, giving the equation $f(x, y) = 1$ along the line $y = ax + b$. Making the substitution $x = (y - b)/a$, we have the Fourier transform of the line image given by

$$\begin{aligned} F(u, v) &= \frac{1}{|a|} \int_{-Y}^Y \int_{-Y}^Y \exp \left\{ -j 2 \pi \left[u \frac{(y - b)}{a} + v y \right] \right\} dy dy \\ &= \frac{2Y}{|a|} \exp \left(j 2 \pi \frac{bu}{a} \right) \operatorname{sinc} \left[\left(\frac{u}{a} + v \right) Y \right]. \end{aligned} \quad (8.11)$$

From the result above, we can see that, for the image of a line, the Fourier transform is a sinc function with an argument that is a linear combination of the two frequency variables (u, v) , and with a slope that is the negative reciprocal of the slope of the original line. The intercept is translated into a phase shift of b/a in the u variable. Thus, the Fourier transform of the line is a sinc function oriented at 90° to the original line, centered about the origin in the frequency domain regardless of the intercept of the original line. This allows us to form filters to select lines solely on the basis of orientation and regardless of the location in the space domain. Spatial components in a certain angle band may thus be obtained by applying a bandpass filter in an angle band perpendicular to the band of interest and applying the inverse transform. If we include a spatial offset in the above calculation, it would only result in a phase shift; the magnitude spectrum would remain the same. Figure 8.2 illustrates the ideal form of the “fan” filter that may be used to select oriented segments in the Fourier domain.

**FIGURE 8.2**

Ideal fan filter in the Fourier domain to select linear components oriented between $+10^\circ$ and -10° in the image plane. Black represents the stopband and white represents the passband. The origin $(u, v) = (0, 0)$ is at the center of the figure.

Prior to the availability of high-speed digital processing systems, attempts at directional filtering used optical processing in the Fourier domain. Arseneault et al. [558] used optical bandpass filters to selectively filter contour lines in aeromagnetic maps. Using optical technology, Duvernoy and Chalasinska-Macukow [560] developed a directional sampling method to analyze images; the method involved integrating along an angle band of the Fourier-transformed image to obtain the directional content. This method was used by Dziedzic-Goclawska et al. [532] to identify directional content in bone tissue images. The need for specialized equipment and precise instrumentation limits the applicability of optical processing. The essential idea of filtering selected angle bands, however, remains valid as a processing tool, and is the basis of Fourier-domain techniques.

The main problem with Fourier-domain techniques is that the filters do not behave well with occluded components or at junctions of linear components; smearing of the line segments occurs, leading to inaccurate results when inverse transformed to the space domain. Another problem lies in the truncation artifacts and spectral leakage that can exist when filtering digitally, which leads to ringing in the inverse-transformed image. Ringing artifacts may be avoided by effective filter design, but this, in turn, could limit the spatial angle band to be filtered.

Considerable research has been reported in the field of multidimensional signal processing to optimize direction-selective filters [569, 576, 577, 580, 583, 585, 590]. Bigün et al. [576] addressed the problem of detection of orientation in a least-squares sense with FIR bandpass filters. This method has the added benefit of being easily implementable in the space domain. Bruton

et al. [577] proposed guidelines for the design of stable IIR fan filters. Hou and Vogel [569] developed a novel method of using the DCT for directional filtering: this method uses the fact that the DCT divides the spectrum into an upper band and a lower band. In 2D, such band-splitting divides the frequency plane into directional filter bands. By selecting coefficients of the DCT, the desired spectral components can be obtained. Because the DCT has excellent spectral reconstruction qualities, this results in high-quality, directionally selective filters. A limitation of this technique is that the method only detects the bounding edges of the directional components, because the band-splitting in the DCT domain does not include the DC component of the directional elements.

In the method developed by Chaudhuri et al. [36], a simple decomposition of the spectral domain into 12 equal angle bands was employed, at 15° per angle band. Each sector filter in this design is a combination of an ideal fan filter, a Butterworth bandpass filter, a ramp-shaped lowpass filter, and a raised cosine window as follows:

$$H(f_r) = \frac{(1 - \beta f_r)}{\left\{ \left[1 + \left(\frac{f_L}{f_r} \right)^{2p} \right] \left[1 + \left(\frac{f_r}{f_H} \right)^{2q} \right] \right\}^{1/2}} \cos^\alpha \left(\frac{\theta - \theta_o}{B} \pi \right), \quad (8.12)$$

where

β	= slope of the weighting function	= 0.7,
f_r	= normalized radial frequency	= $\sqrt{u^2 + v^2}$,
p	= order of the highpass filter	= 6,
q	= order of the lowpass filter	= 4,
f_H	= upper cutoff frequency (normalized)	= 0.5,
f_L	= lower cutoff frequency (normalized)	= 0.02,
θ	= angle of the Fourier transform sample	= $\text{atan}(v/u)$,
θ_o	= central angle of the desired angle band,	
B	= angular bandwidth, and	
α	= weighting factor	= 0.5.

The combined filter with $\theta = 135^\circ$ and $B = 15^\circ$ is illustrated in Figure 8.3.

Filtering an image with sector filters as above results in 12 component images. Each component image contains the linear components of the original image in the corresponding angle band.

Although the directional filter was designed to minimize spectral leakage, some ringing artifacts were observed in the results. To minimize the artifacts, a thresholding method was applied to accentuate the linear features in the image. Otsu's thresholding algorithm [591] (see Section 8.3.2) was applied in the study of collagen fiber images by Chaudhuri et al. [36].




FIGURE 8.3

Directional (sector) filter in the Fourier domain. The brightness is proportional to the gain [36]. Figure courtesy of W.A. Rolston [542].

8.3.2 Thresholding of the component images

Many methods are available for thresholding images with an optimal threshold for a given application [589, 591, 592]. The component images that result from the sector filters as described in Section 8.3.1 possess histograms that are smeared mainly due to the strong DC component that is present in most images. Even with high-quality filters, the DC component appears as a constant in all of the component images due to its isotropic nature. This could pose problems in obtaining an effective threshold to select linear image features from the component images. On the other hand, the removal of the DC component would lead to the detection of edges, and the loss of information related to the thickness of the oriented patterns.

Otsu's method of threshold selection [591] is based upon discriminant measures derived from the gray-level PDF of the given image. Discriminant criteria are designed so as to maximize the separation of two classes of pixels into a foreground (the desired objects) and a background.

Consider the gray-level PDF $p(l)$ of an image with L gray levels, $l = 0, 1, 2, \dots, L - 1$. If the PDF is divided into two classes C_0 and C_1 separated by a threshold k , then the probability of occurrence ω_i of the class C_i , $i = \{0, 1\}$, is given by

$$\omega_0(k) = P(C_0) = \sum_{l=0}^k p(l) = \omega(k), \quad (8.13)$$

$$\omega_1(k) = P(C_1) = \sum_{l=k+1}^{L-1} p(l) = 1 - \omega(k), \quad (8.14)$$

and the class mean levels μ_i for C_i , $i = \{0, 1\}$, are given by

$$\mu_0(k) = \sum_{l=0}^k l P(l|C_0) = \sum_{l=0}^k l \frac{p(l)}{\omega_0(k)} = \frac{\mu(k)}{\omega(k)}, \quad (8.15)$$

and

$$\mu_1(k) = \sum_{l=k+1}^{L-1} l P(l|C_1) = \sum_{l=k+1}^{L-1} l \frac{p(l)}{\omega_1(k)} = \frac{\mu_T - \mu(k)}{1 - \omega(k)}, \quad (8.16)$$

where

$$\omega(k) = \sum_{l=0}^k p(l) \quad (8.17)$$

and

$$\mu(k) = \sum_{l=0}^k l p(l) \quad (8.18)$$

are the cumulative probability and first-order moment of the PDF $p(l)$ up to the threshold level k , and

$$\mu_T = \sum_{l=0}^{L-1} l p(l) \quad (8.19)$$

is the average gray level of the image.

The class variances are given by

$$\sigma_0^2(k) = \sum_{l=0}^k [l - \mu_0(k)]^2 P(l|C_0) = \sum_{l=0}^k [l - \mu_0(k)]^2 \frac{p(l)}{\omega_0(k)}, \quad (8.20)$$

and

$$\sigma_1^2(k) = \sum_{l=k+1}^{L-1} [l - \mu_1(k)]^2 P(l|C_1) = \sum_{l=k+1}^{L-1} [l - \mu_1(k)]^2 \frac{p(l)}{\omega_1(k)}. \quad (8.21)$$

Using the discriminant criterion

$$\nu = \frac{\sigma_B^2(k)}{\sigma_T^2}, \quad (8.22)$$

where

$$\sigma_B^2(k) = \omega_0(k)[\mu_0(k) - \mu_T]^2 + \omega_1(k)[\mu_1(k) - \mu_T]^2, \quad (8.23)$$

and

$$\sigma_T^2 = \sum_{l=0}^{L-1} (l - \mu_T)^2 p(l), \quad (8.24)$$

Otsu's algorithm aims to find the threshold level k that maximizes the discriminant criterion ν given in Equation 8.22. Maximizing ν reduces to maximizing σ_B^2 , because the value σ_T^2 does not vary with the threshold value k . The optimal threshold value k^* is given as

$$k^* = \arg \left\{ \max_{0 \leq k \leq L-1} \sigma_B^2(k) \right\}. \quad (8.25)$$

Otsu's method of thresholding performs well in binarizing a large class of images.

Example: Chaudhuri et al. [36] applied the directional filtering procedure described above to a test pattern with line segments of various length, width, and gray level at four different angles, namely 0° , 45° , 90° , and 135° , as shown in Figure 8.4 (a). Significant overlap was included in order to test the performance of the filtering procedures under nonideal conditions. The log-magnitude Fourier spectrum of the test image is shown in part (b) of the figure; directional concentrations of energy are evident in the spectrum. The component image obtained using the filtering procedure for the angle band $125^\circ - 140^\circ$ in the Fourier domain is shown in Figure 8.4 (c). It is evident that only those lines oriented at 45° in the image plane have been passed, along with some artifacts. The corresponding binarized image, using the threshold value given by Otsu's method described above, is shown in part (d) of the figure. Parts (e) and (f) of the figure show the binarized components extracted from the test image for the angle bands $80^\circ - 95^\circ$ and $125^\circ - 140^\circ$ in the image plane.

A close inspection of the component images in Figure 8.4 indicates that regions of overlap of lines oriented at different directions contribute to each direction. The 135° component image in Figure 8.4 (f) has the largest error due to its low gray level and the large extent of overlap with the other lines. The areas of the line segments extracted by the filtering procedure had errors, with respect to the known areas in the original test image, of 3.0%, -4.3% , -3.0% , and -28.6% for the 0° , 45° , 90° , and 135° components, respectively. The results of application of methods as above for the directional analysis of collagen fibers in ligaments are described in Section 8.7.1.

8.3.3 Design of fan filters

Fan filters are peculiar to 2D filtering: there is no direct 1D analog of this type of filtering. This fact presents some difficulties in designing fan filters: because we cannot easily extend the well-established concepts that are used to design 1D filters. The main problem in the design of fan filters lies in forming the filter at the origin $(u, v) = (0, 0)$ or the DC point in the Fourier domain. At the DC point, the ideal fan filter structure has a knife edge, which makes the filter nonanalytic; that is, if one were to approach the origin in the Fourier domain from any point within the passband of the fan, the limit would ideally




Figure 8.4 (e)

(f)

FIGURE 8.4

(a) A test image with overlapping directional components at 0° , 45° , 90° , and 135° . (b) Log-magnitude Fourier spectrum of the test image. Results of directional filtering (with the angle bands specified in the image domain): (c) $35^\circ - 50^\circ$. (d) Result in (c) after thresholding and binarization. (e) $80^\circ - 95^\circ$ (binarized). (f) $125^\circ - 140^\circ$ (binarized). Reproduced with permission from S. Chaudhuri, H. Nguyen, R.M. Rangayyan, S. Walsh, and C.B. Frank, "A Fourier domain directional filtering method for analysis of collagen alignment in ligaments", *IEEE Transactions on Biomedical Engineering*, 34(7): 509 – 518, 1987. © IEEE.

be unity. On the other hand, the limit as one approaches the origin from a point within the stopband should be zero.

Various methods have been applied to overcome the problem stated above, such as the use of spectrum-shaping smoothing functions with the ideal fan filter; for example, Chaudhuri et al. [36] used the Butterworth and raised cosine functions, as described in Section 8.3.1. However, the performance of even the best spectrum-shaping function is limited by the tight spectral constraints imposed by the nonanalytic point at the origin. To obtain better spectral shaping, as in 1D, high-order FIR filters or low-order IIR filters may be used.

The discontinuity at the origin $(u, v) = (0, 0)$ is the main problem in the design of recursive fan filters. With nonrecursive or FIR filters, this problem does not result in instability. With IIR filters, instability can occur if the filters are not properly designed. Stability of filters is usually defined as bounded-input – bounded-output (BIBO) stability [577]. Filters that are BIBO stable ensure that all inputs that are not infinite in magnitude will result in outputs that are bounded.

2D filters are commonly derived from real, rational, continuous functions of the form

$$T(s_1, s_2) = \frac{Q(s_1, s_2)}{P(s_1, s_2)} = \frac{\sum_{m=0}^{M_2} \sum_{n=0}^{N_2} q_{mn} s_1^m s_2^n}{\sum_{m=0}^{M_1} \sum_{n=0}^{N_1} p_{mn} s_1^m s_2^n}, \quad (8.26)$$

where s_1 and s_2 are the Laplace variables; the function $T(s_1, s_2)$ is the Laplace-transformed version of the 2D partial differential equation that is related to the required filter response; $Q(s_1, s_2)$ is the numerator polynomial resulting from the Laplace transform of the forward differential forms expressed as a sum of products in s_1 and s_2 with the coefficients q_{mn} ; M_2 and N_2 represent the order of the polynomial Q in m and n , respectively; $P(s_1, s_2)$ is the denominator polynomial obtained from the Laplace transform of the backward differential forms expressed as a sum of products in s_1 and s_2 with the coefficients p_{mn} ; and M_1 and N_1 represent the order of the polynomial P in m and n , respectively. The corresponding frequency response function $T(u, v)$ is obtained by the substitution of $s_1 = j 2 \pi u$ and $s_2 = j 2 \pi v$.

The discontinuous requirement in the continuous prototype filter at the origin results in the filter transfer function $T(s_1, s_2)$ having a nonessential singularity of the second kind at the origin. A nonessential singularity of the second kind occurs when the numerator and the denominator polynomials, $P(s_1, s_2)$ and $Q(s_1, s_2)$ in Equation 8.26, approach zero at the same frequency location (a_1, a_2) , resulting in $T(a_1, a_2) = \frac{0}{0}$.

The discrete form of the function in Equation 8.26 is obtained through the 2D version of the bilinear transform in 1D, given as

$$s_i = \frac{(z_i - 1)}{(z_i + 1)} \quad \text{for } i = 1, 2, \quad (8.27)$$

to obtain the following discrete version of the filter:

$$H(z_1, z_2) = \frac{B(z_1, z_2)}{A(z_1, z_2)} = \frac{\sum_{m=0}^{M_2^*} \sum_{n=0}^{N_2^*} b_{mn} z_1^{-m} z_2^{-n}}{\sum_{m=0}^{M_1^*} \sum_{n=0}^{N_1^*} a_{mn} z_1^{-m} z_2^{-n}}, \quad (8.28)$$

where the orders of the polynomials M_1^* , N_1^* , M_2^* , and N_2^* are different from the corresponding limits of the continuous-domain filter in Equation 8.26 due to the bilinear transform.

Filter design using nonessential singularities: The 2D filter design method of Bruton and Bartley [587] views the nonessential singularity inherent to fan filters not as an obstacle in the design process, but as being necessary in order to generate useful magnitude responses. The method relies on classical electrical circuit theory, and views the input image as a surface of electrical potential. The surface of electrical potential is then acted upon by a 2D network of electrical components such as capacitors, inductors, and resistors; the components act as integrators, differentiators, and dissipators, respectively. The central idea is to construct a network of components that will not add energy to the input; that is, to make a completely passive circuit. The passiveness of the resulting circuit will result in no energy being added to the system (that is, the filter is “nonenergetic”), and thus will ensure that the filter is stable. Bruton and Bartley [587] showed that the necessary condition for a filter to be stable is that the admittance matrix that links the current and voltage surfaces must have negative Toeplitz symmetry with reactive elements supplied by inductive elements that satisfy the nonenergetic constraint.

The nonenergetic condition ensures that the filter is stable, because it implies that the filter is not adding any energy to the system. The maximum amount of energy that is output from the filter is the maximum amount put into the system by the input image. The derivation given by Bruton and Bartley [587] is the starting point of a numerical method for designing stable, recursive, 2D filters. The derivation shows that recursive filters can be built reliably, as long as the condition above on the admittance matrix is met.

Bruton and Bartley [587] provided the design and coefficients of a narrow, 15° fan-stop filter, obtained using a numerical optimization method with the

TABLE 8.1
Coefficients of the Discrete-domain Fan Filter with a 15° Fan Stopband [542].

b_{mn}	$n = 0$	$n = 1$	$n = 2$
$m = 0$	0.02983439380935332	-0.6855181788590949	0.7027763362367445
$m = 1$	-0.1469615281783627	3.397745073546105	-3.629041657524303
$m = 2$	0.2998008459584214	-6.767662643767763	7.49061181619684
$m = 3$	-0.3165448124171246	6.771378027945815	-7.725572280971142
$m = 4$	0.1724438585800683	-3.403226865621513	3.981678690012933
$m = 5$	-0.03857214742977072	0.6872844383634052	-0.82045337027416

a_{mn}	$n = 0$	$n = 1$	$n = 2$
$m = 0$	1.0000000000000000	-0.82545044546957	0.03722700706807863
$m = 1$	-4.476280705843249	3.791276128445935	-0.161179724642936
$m = 2$	8.03143251366382	-7.00124160940265	0.2870351311929377
$m = 3$	-7.220029589516617	6.499290024154175	-0.2623441075303727
$m = 4$	3.252431250257176	-3.03268003600527	0.122960282645262
$m = 5$	-0.5875259501210567	0.5687740686107076	-0.0236904653803231

Data courtesy of N.R. Bartley [587].

condition described above added. The method results in a recursive filter of small order with remarkable characteristics. A filter of fifth order in z_1 and second order in z_2 was designed using this method. The corresponding coefficients of the discrete function $H(z_1, z_2)$ as in Equation 8.28 are listed in Table 8.1. The coefficients in the numerator and denominator each add up to zero at $z_1 = 1$ and $z_2 = 1$, confirming that the filter conforms to the requirement of the knife-edge discontinuity.

Rotation of the filter and image: The fan filter design algorithm of Bruton and Bartley [587] provides filters only for a specific angle band — in the above case, for a 15° bandstop filter centered at 0° in the Fourier domain. In order to obtain filters with different central orientations, it is necessary to perform a rotation of the prototype filter. This may be achieved by using the following substitution in the analog prototype filter transfer function [542]:

$$\begin{aligned}
 s_1 &\leftarrow s_1 \cos \theta + s_2 \sin \theta \\
 s_2 &\leftarrow s_2,
 \end{aligned}
 \tag{8.29}$$

where θ is the amount of rotation desired. The discrete version of the filter is obtained through the bilinear transform. The rotation step above is not the usual rotational transformation for filters, but it is necessary to use this transformation in order to ensure that the filter is stable. [If the normal rotational transformation were to be used, s_2 would also be rotated as

$$s_2 \leftarrow -s_1 \sin \theta + s_2 \cos \theta. \quad (8.30)$$

Then, values of s_2 could turn out to be negative: this would indicate that there would be energy added to the system, which would make the filter unstable.]

Suppose that the prototype filter of the form shown in Equation 8.26, given by $T_0(s_1, s_2)$ and with the corresponding frequency response function given by $T_0(u, v)$, is bounded by the straight lines L_- and L_+ passing through the origin at angles of $-\theta_p$ and $+\theta_p$ with the central line of the filter $CL = 0^\circ$ where $u = 0$, as shown in Figure 8.5 (a). The lines L_- and L_+ are given by

$$\begin{aligned} u \cos \theta_p - v \sin \theta_p &= 0 : L_- \\ u \cos \theta_p + v \sin \theta_p &= 0 : L_+. \end{aligned} \quad (8.31)$$


As a result of the transformation in Equation 8.29, the center of the passband of the rotated frequency response $T_r(u, v)$ is given as $T_0(u', v') = T_0(u \cos \theta_c + v \sin \theta_c, v)$. Similarly, the straight lines L_- and L_+ are rotated to the straight lines given by

$$\begin{aligned} u \cos \theta_p \cos \theta_c + v (\sin \theta_c \cos \theta_p - \sin \theta_p) &= 0 : L_- \\ u \cos \theta_p \cos \theta_c + v (\sin \theta_c \cos \theta_p + \sin \theta_p) &= 0 : L_+; \end{aligned} \quad (8.32)$$

[see Figure 8.5 (b)].

A limitation to filter rotation as above is that rotating the filter by more than 45° would result in a loss of symmetry about the central line of the filter. The rotational warping effect may be compensated for in the prototype filter $T_0(s_1, s_2)$. In the work of Rolston [542], the prototype filter was rotated by 45° in either direction to obtain filters covering an angle band of 90° ($0^\circ - 45^\circ$ and $135^\circ - 180^\circ$ in the Fourier domain). Filtering in the range $45^\circ - 135^\circ$ was achieved by rotating the image by 90° before passing it through the same filters as above.

The fan filter as above has unit gain, which has some drawbacks as well as some advantages. An advantage is that features in the filtered component images have no more intensity than the corresponding features in the original image, so that image components that exist in the original image in a particular angle band will not be attenuated at all or will be attenuated only slightly. This also limits the number of iterations necessary to attenuate out-of-band components for a certain class of images. The unit gain is an advantage with images that have only a small depth of field, because a global threshold can be used for all component images to obtain a representation of the components that exist in each specific angle band.

**FIGURE 8.5**


(a) Original fan filter. (b) The fan filter after rotation by the transformation given in Equation 8.29. Figure courtesy of W.A. Rolston [542].

Example: A test image including rectangular patches of varying thickness and length at 0° , 45° , 90° , and 135° , with significant overlap, is shown in Figure 8.6 (a). The fan filter as described above, for the central angle of 90° , was applied to the test image. Because the fan filter is a fan-stop filter, the output of the filter was subtracted from the original image to obtain the fan-pass response. As evident in Figure 8.6 (b), the other directional components are still present in the result after one pass of the filter due to the nonideal attenuation characteristics of the filter. The filter, however, can provide improved results with multiple passes or iterations. The result of filtering the test image nine times is shown in Figure 8.6 (c). The result possesses good contrast between the desired objects and the background, and may be thresholded to reject the remaining artifacts.

8.4 Gabor Filters

Most of the directional, fan, and sector filters that have been used in the Fourier domain to extract directional elements are not analytic functions. This implies that filter design methods in 1D are not applicable in 2D. Such filters tend to possess poor spectral response, and yield images with not only the desired directional elements but also artifacts.

One of the fundamental problems with Fourier methods of directional filtering is the difficulty in resolving directional content at the DC point (the origin in the Fourier domain) [587]. The design of high-quality fan filters requires

**FIGURE 8.6**

(a) A test image with overlapping directional components at 0° , 45° , 90° , and 135° . Results of fan filtering at 90° after (b) one pass, (c) nine passes. Figure courtesy of W.A. Rolston [542].

conflicting constraints at the DC point: approaching the DC point from any location within the passband requires the filter gain to converge to unity; however, approaching the same point from any location in the stopband requires the filter gain to approach zero. In terms of complex analysis, this means that the filter is not analytic, or that it does not satisfy the Cauchy-Riemann equations. This prevents extending results in 1D to problems in 2D.

The Gabor filter provides a solution to the problem mentioned above by increasing the resolution at the DC point. Gabor filters are complex, sinusoidally modulated, Gaussian functions that have optimal localization in both the frequency and space domains [389]. Gabor filters have been used for texture segmentation and discrimination [381, 495, 542, 543, 593, 594, 595], and may yield better results than simple Fourier methods for directional filtering.

Time-limited or space-limited functions have Fourier spectra of unlimited extent. For example, the time-limited rectangular pulse function transforms into the infinitely long sinc function. On the other hand, the time-unlimited sine function transforms into a delta function with infinite resolution in the Fourier domain. Infinitely long functions cannot be represented in finite calculating machinery. Gabor [596] suggested the use of time-limited functions as the kernels of a transform instead of the unlimited sine and cosine functions that are the kernel functions of the Fourier transform. The functional nature of the Fourier transform implies that there exists an “uncertainty principle”, similar, but not identical, to the well-known Heisenberg uncertainty principle of quantum mechanics. Gabor showed that complex, sinusoidally modulated, Gaussian basis functions satisfy the lower bound on the fundamental uncertainty principle that governs the resolution in time and frequency, given by

$$\Delta t \Delta f \geq \frac{1}{4\pi}, \quad (8.33)$$

where Δt and Δf are time and frequency resolution, respectively. The uncertainty principle implies that there is a resolution limit between the spatial and the Fourier domains. Gabor proved that there are functions that can form the kernel of a transform that exactly satisfy the uncertainty relationship. The functions named after Gabor are Gaussian-windowed sine and cosine functions. By limiting the kernel functions of the Fourier transform with a Gaussian windowing function, it becomes possible to achieve the optimal resolution limit in both the frequency and time domains. The size of the Gaussian window function needs to be used as a new parameter in addition to the frequency of the sine and cosine functions.

Gabor functions provide optimal joint resolution in both the Fourier and time domains in 1D, and form a complete basis set through phase shift and scaling or dilation of the original (mother) basis function. The set of functions forms a multiresolution basis that is commonly referred to as a wavelet basis (formalized by Mallat [386]).

Daugman [389] extended Gabor functions to 2D as 2D sinusoidal plane waves of some frequency and orientation within a 2D Gaussian envelope. Ga-

bor functions have also been found to provide good models for the receptive fields of simple cells in the striate cortex [571, 389]; for this reason, there has been a significant amount of research conducted on using the functions for texture segmentation, analysis, and discrimination [381, 495, 542, 543, 593, 594, 595].

The extension of the principle above to 2D leads to space-limited plane waves or complex exponentials. Such an analysis was performed by Daugman [389]. The uncertainty relationship in 2D is given by

$$\Delta x \Delta y \Delta u \Delta v \geq \frac{1}{16\pi^2}, \quad (8.34)$$

where Δx and Δy represent the spatial resolution, and Δu and Δv represent the frequency resolution. The 2D Gabor functions are given as

$$\begin{aligned} h(x, y) &= g(x', y') \exp[-j 2\pi (Ux + Vy)], \\ (x', y') &= (x \cos \phi + y \sin \phi, -x \sin \phi + y \cos \phi), \end{aligned} \quad (8.35)$$

where (x', y') are the (x, y) coordinates rotated by an arbitrary angle ϕ ,

$$g(x, y) = \left(\frac{1}{2\pi\lambda\sigma^2} \right) \exp \left[-\frac{(x/\lambda)^2 + y^2}{2\sigma^2} \right] \quad (8.36)$$

is a Gaussian shaping window with the aspect ratio λ , and U and V are the center frequencies in the (u, v) frequency plane. An example of the real part of a Gabor kernel function is given in Figure 8.7 with $\sigma = 0.5$, $\lambda = 0.6$, $U = 1$, $V = 0$, and $\phi = 0$ (with reference to Equations 8.35 and 8.36). Another Gabor kernel function is shown in gray scale in Figure 8.8.

The imaginary component of the Gabor function is the Hilbert transform of its real component. The Hilbert transform shifts the phase of the original function by 90° , resulting in an odd version of the function.

The ‘‘Gabor transform’’ is not a transform as such; that is, there is usually no transform domain into which the image is transformed. The frequency domain is usually divided into a symmetric set of slightly overlapping regions at octave intervals. Examples of the ranges related to a few Gabor functions are shown in Figure 8.9; see also Figures 5.69, 8.57, and 8.68. It is evident that Gabor functions act as bandpass filters with directional selectivity.

8.4.1 Multiresolution signal decomposition

Multiresolution signal analysis is performed using a single prototype function called a wavelet. Fine temporal or spatial analysis is performed with contracted versions of the wavelet; on the other hand, fine frequency analysis is performed with dilated versions. The definition of a wavelet is flexible, and requires only that the function have a bandpass transform; thus, a wavelet at a particular resolution acts as a bandpass filter. The bandpass filters must




FIGURE 8.7

An example of the Gabor kernel with $\sigma = 0.5$, $\lambda = 0.6$, $U = 1$, $V = 0$, and $\phi = 0$ (with reference to Equations 8.35 and 8.36). Figure courtesy of W.A. Rolston [542].



FIGURE 8.8

An example of a Gabor kernel, displayed as an image. Figure courtesy of W.A. Rolston [542].




FIGURE 8.9

Division of the frequency domain by Gabor filters. Two sets of oval regions are shown in black, corresponding to the passbands of three filters in each set, oriented at 0° and 90° . In each case, the three regions correspond to three scales of the Gabor wavelets. There is a 90° shift between the angles of corresponding filter functions in the space and frequency domains. Figure courtesy of W.A. Rolston [542].

have constant relative bandwidth or constant quality factor. The importance of constant relative bandwidth of perceptual processes such as the auditory and visual systems has long been recognized [571]. Multiresolution analysis has also been used in computer vision for tasks such as segmentation and object recognition [284, 285, 288, 487]. The analysis of nonstationary signals often involves a compromise between how well transitions or discontinuities can be located, and how finely long-term behavior can be identified. This is reflected in the above-mentioned uncertainty principle, as established by Gabor.

Gabor originally suggested his kernel function to be used over band-limited, equally spaced areas of the frequency domain, or equivalently, with constant window functions. This is commonly referred to as the short-time Fourier transform (STFT) for short-time analysis of nonstationary signals [176, 31]. The 2D equivalent of the STFT is given by

$$F_S(x', y', u, v) = \int_{x=-\infty}^{\infty} \int_{y=-\infty}^{\infty} f(x, y) w(x - x', y - y') \times \exp[-j 2 \pi (ux + vy)] dx dy, \quad (8.37)$$

where w is a windowing function and f is the signal (image) to be analyzed. The advantage of short-time (or moving-window) analysis is that if the energy of the signal is localized in a particular part of the signal, it is also localized to a part of the resultant 4D space (x', y', u, v) . The disadvantage of this method is that the same window is used at all frequencies, and hence, the

resolution is the same at all locations in the resultant space. The uncertainty principle does not allow for arbitrary resolution in both of the space and frequency domains; thus, with this method of analysis, if the window function is small, the large-scale behavior of the signal is lost, whereas if the window is large, rapid discontinuities are washed out. In order to identify the fine or small-scale discontinuities in signals, one would need to use basis functions that are small in spatial extent, whereas functions of large spatial extent would be required to obtain fine frequency analysis. By varying the window function, one will be able to identify both the discontinuous and stationary characteristics of a signal. The notion of *scale* is introduced when the size of the window is increased by an order of magnitude. Such a multiresolution or multiscale view of signal analysis is the essence of the wavelet transform. Wavelet decomposition, in comparison to STFT analysis, is performed over regions in the frequency domain of constant relative bandwidth as opposed to a constant bandwidth.

In the problem of determining the directional nature of an image, we have the discontinuity in the frequency domain at the origin, or DC, to overcome. Wavelet analysis is usually applied to identify discontinuities in the spatial domain; however, there is a duality in wavelet analysis, provided by the uncertainty principle, that allows discontinuity analysis in the frequency domain as well. In order to analyze the discontinuity at DC, large-scale or dilated versions of the wavelet need to be used. This is the dual of using contracted versions of the wavelet to analyze spatial discontinuities.

The wavelet basis is given by

$$h_{x',y',\lambda_1,\lambda_2}(x,y) = \frac{1}{\sqrt{\lambda_1\lambda_2}} h\left(\frac{x-x'}{\lambda_1}, \frac{y-y'}{\lambda_2}\right) \quad (8.38)$$

where x', y', λ_1 , and λ_2 are real numbers, and h is the basic or mother wavelet. For large values of λ_1 and λ_2 , the basis function becomes a stretched or expanded version of the prototype wavelet or a low-frequency function, whereas for small λ_1 and λ_2 , the basis function becomes a contracted wavelet, that is, a short, high-frequency function.

The wavelet transform is then defined as

$$F_W(x', y', \lambda_1, \lambda_2) = \frac{1}{\sqrt{\lambda_1\lambda_2}} \int_{x=-\infty}^{\infty} \int_{y=-\infty}^{\infty} f(x, y) \times h\left(\frac{x-x'}{\lambda_1}, \frac{y-y'}{\lambda_2}\right) dx dy. \quad (8.39)$$

From this definition, we can see that wavelet analysis of a signal consists of the contraction, dilation, and translation of the basic mother wavelet, and computing the projections of the resulting wavelets on to the given signal.

8.4.2 Formation of the Gabor filter bank

In the method proposed by Bovik et al. [594], the given image is convolved with the complex Gabor kernel, and the maximum magnitude of the result is taken as an indicator to identify changes in the dominant orientation of the image. In the work of Rolston and Rangayyan [542, 543], this method was observed to fail in the presence of broad directional components. The real component of the Gabor filter acts as a matched filter to detect broad directional components, and thus, is better suited to the identification of such regions.

The parameters of Gabor filters that may be varied are as follows: With reference to Equations 8.35 and 8.36, the parameter σ specifies the spatial extent of the filter; λ specifies the aspect ratio of the filter that modulates the σ value. If $\lambda = 1$, the ϕ parameter in Equation 8.35 need not be specified, because $g(x, y)$ is then isotropic. In the frequency domain, such a filter results in an oriented filter occupying the middle subsection of the corresponding ideal fan filter, with the orientation being specified by $\tan^{-1}(V/U)$; see Figure 8.9. These parameters will then completely specify the Gabor filter bank.

In the directional analysis algorithm proposed by Rolston and Rangayyan [542, 543], only the real component of the Gabor wavelet is used, with $\lambda = 1/0.6$, $\sigma = 1.0$, and the primary orientation given by $\tan^{-1}(V/U) = 0^\circ, 45^\circ, 90^\circ, \text{ and } 135^\circ$. A given image is analyzed by convolving band-limited and decimated versions of the image with the same analyzing wavelet.

When a decimated image is convolved with a filter of constant spatial extent, relative to the original image, the filter is effectively scaled larger with respect to the decimated image. The advantage of this procedure is that filters with larger σ values, or with center frequencies closer to DC, can be simulated, instead of resorting to using filters of larger spatial extent. Filters with larger σ values correspond to portions of the frequency domain closer to the DC point. This effect is shown in Figure 8.9. The frequency plane is completely covered by the decimation and filtering operation. Each black oval in Figure 8.9 represents the frequency band or region that is being filtered by each decimation and filtering operation. The largest black oval at each orientation corresponds to one-to-one filtering, and the smaller ovals closer to the origin correspond to higher orders of decimation and filtering. Higher levels of decimation and filtering geometrically approach the DC point. Theoretically, arbitrary resolution of the DC point can be achieved using this method; however, the size of the original image imposes a limiting factor, because an image that has been digitized, for example, to 256×256 pixels, can only be decimated a few times before the details of interest are lost.

Another advantage of this method is that, because filtering is performed in the spatial domain, stable filters are obtained at DC, whereas, with strict Fourier-transform-based methods, the resolution at DC is reduced to one frequency increment or less, thereby making the Fourier-domain filters sensitive.

Because the filter bank works on decimated images, the computational load of convolution reduces geometrically at each successive stage of decimation.

8.4.3 Reconstruction of the Gabor filter bank output

In the directional filtering and analysis procedures proposed by Rolston and Rangayyan [542, 543], the given image is decimated and convolved at each of three scales with a filter of fixed size. Decimation and filtering at each scale results in equal energy across all of the scales due to the selection of the filter coefficients. Thus, after interpolation of the decimated and convolved images, the responses at the different scales can be added without scaling to obtain the overall response of the filter at the different scales.

After obtaining the responses to the filters at 0° , 45° , 90° , and 135° , a vector summation of the filter responses is performed, as shown in Figure 8.10. The vector summation is performed at each pixel in the original image domain to obtain a magnitude and phase (or angle) at each pixel.




FIGURE 8.10

Vector summation of the responses of Gabor filters at 0° , 45° , 90° , and 135° . Figure courtesy of W.A. Rolston [542].

The Gabor filters as described above do not have a perfect reconstruction condition. This results in a small amount of out-of-band energy interfering with the reconstruction, translated into artifacts in the reconstructed image. A thresholding operation can effectively remove such artifacts. Rolston and Rangayyan [542, 543] set the threshold as the maximum of the sum of the mean and standard deviation values across all of the component images.

Example: A test image including rectangular patches of varying thickness and length at 0° , 45° , 90° , and 135° , with significant overlap, is shown in Figure

8.11 (a). The output at each stage of the Gabor filter as described above for 0° is shown in same figure. It is evident that, with one-to-one filtering, the narrow horizontal elements have been successfully extracted; however, only the edges of the broader components are present in the result. As the decimation ratio is increased, the broader components are extracted, which indicates that the filtering is effective in low-frequency bands closer to the DC point.


See Sections 5.10.2, 8.9, and 8.10 for further discussions and results related to Gabor filters.

8.5 Directional Analysis via Multiscale Edge Detection

Methods for edge detection via multiscale analysis using LoG functions are described in Section 5.3.3. Liu et al. [37, 531] applied further steps to the edge stability map obtained by this method (see Figure 5.16) to detect linear segments corresponding to collagen fibers in SEM images of ligaments, which are described in the following paragraphs.

Estimating the area of directional segments: Directional analysis requires the estimation of the area covered by linear segments in specified angle bands. For this purpose, the pattern boundaries obtained by the relative stability index (see Equation 5.26) may be used for computing the area of coverage. The directional information of a pattern is given by the directions of the gradients along the detected pattern boundaries.

Figure 8.12 (a) depicts the approach of Liu et al. [37, 531] to area computation, where two pattern-covered regions are denoted by R_A and R_B . The arrows along the boundaries indicate the directions of the gradients, which are computed from the original image on a discretized grid depending upon the application. The use of gradients enables the definition of the region enclosed by the boundaries. It is seen from Figure 8.12 (a) that a linear segment can be identified by a pair of line segments running in opposite directions. In order to identify the region, Liu et al. [37, 531] proposed a piecewise labeling procedure that includes two steps: line labeling and region labeling. In the line-labeling procedure [597], the full plane is sectioned into eight sectors (see Figure 8.13), and a set of templates is defined for pixel classification. The relative stability index is scanned left to right and top to bottom. To each element in the relative stability index, a line label is assigned according to its match with one of the templates. A structure array is constructed to store the descriptions of the lines at both pixel and line levels. The structure array contains several description fields, such as the line starting location (x_s, y_s) ; the ending location (x_e, y_e) ; the orientation θ ; and a corner label, which is also a structure array, containing the corner location and the lines that form the corner [597].

**FIGURE 8.11**

(a) A test image with overlapping directional components at 0° , 45° , 90° , and 135° . Results of Gabor filtering at 0° after decimation at (b) one-to-one, (c) two-to-one, and (d) four-to-one. (e) Overall response at 0° after vector summation as illustrated in Figure 8.10. Figure courtesy of W.A. Rolston [542].




FIGURE 8.12

(a) Computation of the area covered by directional segments. The arrows perpendicular to the pattern boundaries represent gradient directions used for detecting the interior of the linear segment over which the area is computed. The directional information associated with the pattern is also stored for analysis. (b) Computation of occluded segments based upon the detected T-joints. The subscripts denote different regions, and the superscripts denote the line numbers. Reproduced with permission from Z.-Q. Liu, R.M. Rangayyan, and C.B. Frank, "Directional analysis of images in scale-space", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(11):1185–1192, 1991. © IEEE.

Once the line segments have been labeled, a set of region descriptors is generated, which includes paired line labels, their starting and ending locations, orientation, and the area of the region [see Figure 8.12 (a)]. In region labeling, a line (for example, $Line_A^1$) is paired with an adjacent line (for example, $Line_A^2$) having a direction that is in the sector opposite to that of $Line_A^1$ (see Figure 8.13). The area of the linear segment (R_A) is then computed by counting the number of pixels contained by the pair of line segments. The orientation of the linear segment is indicated by the orientation of the pair of line segments. For instance, if $Line_A^1$ and $Line_A^2$ form a pair, their associated region descriptor can be defined as

$$R\{A, [(xs, ys), (xe, ye), \theta]_1; [(xs, ys), (xe, ye), \theta]_2; \alpha\}, \quad (8.40)$$

where the subscripts 1 and 2 represent $Line_A^1$ and $Line_A^2$, respectively, and α is the area computed for the region R_A [see Figure 8.12 (a)].




FIGURE 8.13

The image plane is divided into eight sectors. $Line^1$ and $Line^2$ form a pair. Reproduced with permission from Z.-Q. Liu, R.M. Rangayyan, and C.B. Frank, "Directional analysis of images in scale-space", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(11):1185–1192, 1991. © IEEE.

Detection of occluded linear segments: It is often the case in natural images, particularly in SEM images of ligaments, that as linear patterns in-

intersect, some segments of a linear pattern will be occluded. Analysis based upon incomplete patterns will introduce errors. It is desirable that the occluded linear segments are detected and included in the final analysis. In the methods of Liu et al. [37, 531], a simple interpolation method is used for this purpose.

Occluded segments typically appear as T-junctions in an edge image. (See Chen et al. [598] for a discussion on various types of junctions in images with oriented patterns.) As described above, a corner structure array is generated along with the line structure array. T-junctions can be readily detected by inspecting the corners, and if necessary, linking lines according to the following procedure. The lines that form T-junctions with a common line [see Figure 8.12 (b)] are considered to be occluded line segments and are stored in a T-junction array structure:

$$T\{k, Line_A^1, Line_A^2; Line_B^1, Line_B^2; \dots; Line_*^k\}, \quad (8.41)$$

where k indicates the k^{th} T-junction structure, and the subscript $*$ indicates the region associated with the common line. After all the T-junction structures are constructed, they are paired by bringing together the T-junction structures with $Line_*^k$ that share the same region. Corresponding line elements in paired T-junction structures are then compared to detect lines that cut across the common region. This is performed by verifying if a line in one of the T-junction structures of the pair lies within a narrow cone-shaped neighborhood of the corresponding line in the other T-junction structure of the pair. If such a line pair is detected across a pair of T-junction structures, the lines are considered to be parts of a single line with an occluded part under the common region. Furthermore, if two such occluded lines form two regions (on either side of the common region), the two regions are merged by adding the occluded region, and relabeled as a single region. With reference to the situation depicted in Figure 8.12 (b), the above procedure would merge the regions labeled as R_D and R_E into one region, including the area occluded in between them by R_* .

Overview of the algorithm for directional analysis: In summary, the method proposed by Liu et al. [37, 531] for directional analysis via multiscale filtering with LoG functions (see Section 5.3.3) consists of the following steps:

1. Generate a set of zero-crossing maps (images).
2. Classify or authenticate the zero crossings.
3. Generate the adjusted zero-crossing maps from the original zero-crossing maps.
4. Generate a stability map from the set of adjusted zero-crossing maps.
5. Generate the relative stability index map from the stability map.

6. Compute the edge orientation from the relative stability index map and the original image.
7. Compute the orientational distribution of the segments identified.
8. Compute statistical measures to quantify the angular distribution of the linear patterns (such as entropy and moments, as described in Section 8.2).

The methods described above were tested with the image in Figure 8.4 (a). The areas of the line segments extracted by the procedures had errors, with respect to the known areas in the original test image, of -2.0% , -6.3% , -3.4% , and -40.6% for the 0° , 45° , 90° , and 135° components, respectively.

Liu et al. [37, 531] applied the procedures described above for the analysis of collagen remodeling in ligaments; this application is described in detail in Section 8.7.1.

8.6 Hough-Radon Transform Analysis

The Hough transform is a method of transforming an image into a parameter domain where it is easier to obtain the desired information in the image; see Section 5.6.1. The main drawback of the Hough transform is that it is primarily applicable to binary images; as a consequence, the results are dependent upon the binarization method used for segmenting the image. Rangayyan and Rolston [599, 542] proposed the use of a combination of the Hough transform and the Radon transform (see Section 9.1) that overcomes this drawback; their methods and results obtained are described in the following paragraphs.

8.6.1 Limitations of the Hough transform

With reference to Figure 8.14, we see that a straight line can be specified in terms of its orientation θ with respect to the x axis, and its distance ρ from the origin. In this form of parameterization, any straight line is bounded in angular orientation by the interval $[0, \pi]$ and bounded by the Euclidean distance to the farthest point of the image from the center of the image. The equation for an arbitrary straight-line segment in the image plane is given by

$$\rho = x \cos \theta + y \sin \theta. \quad (8.42)$$

For a specific point in the image domain (x_i, y_i) , we obtain a sinusoidal curve in the Hough domain (ρ, θ) . Each point (x_i, y_i) lying on a straight line with $\rho = \rho_0$ and $\theta = \theta_0$ in the image domain corresponds to a sinusoidal curve in

the (ρ, θ) domain specified by

$$\rho_0 = x_i \cos \theta_0 + y_i \sin \theta_0. \quad (8.43)$$

Through Equation 8.43, it is evident that for each point in the image domain, the Hough transform performs a one-to-many mapping, resulting in a modulated sum of sinusoids in the Hough domain.

The Hough transform is often referred to as a voting procedure, where each point in the image casts votes for all parameter combinations that could have produced the point. All of the sinusoids resulting from the mapping of a straight line in the image domain have a common point of intersection at (ρ_0, θ_0) in the Hough domain. Linear segments in the spatial domain correspond to large-valued points in the Hough domain; see Figures 5.39 and 5.40. Thus, the problem of determining the directional content of an image becomes a problem of peak detection in the Hough parameter space.

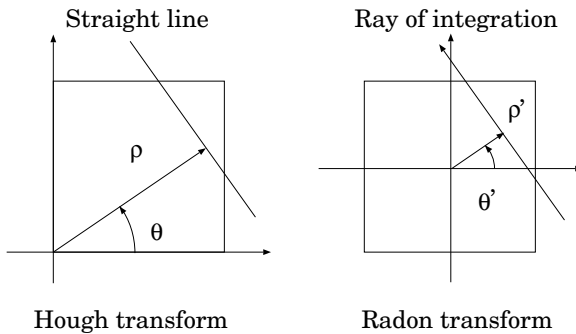


FIGURE 8.14

Parameters in the representation of a straight line in the Hough transform and a ray in the Radon transform. Reproduced with permission from R.M. Rangayyan and W.A. Rolston, “Directional image analysis with the Hough and Radon transforms”, *Journal of the Indian Institute of Science*, 78: 17–29, 1998. © Indian Institute of Science.

From the properties listed above, the Hough transform appears to be the ideal tool for detecting linear components in images. However, there are some limitations to this approach. The results are sensitive to the quantization intervals used for the angle θ and the distance ρ . Decreasing the quantization step for θ increases the computation time, because the calculation for ρ needs to be performed across each value of θ and each pixel. Another problem with this method is the “crosstalk” between multiple straight lines in the Hough domain. If the image contains several lines parallel to the x axis, they would correspond to several peak values in the Hough domain at differing ρ values

for $\theta = 90^\circ$. However, the Hough transform would also detect false linear segments for $\theta = 0^\circ$, which would show up as smaller peaks at a continuum of ρ values in the Hough domain; see Figure 8.15. This is caused by the fact that the Hough transform finds line segments at specific ρ values that are not necessarily contiguous. Another form of crosstalk can occur within a broad directional element: several straight lines may be perceived within a broad element with angles spread about the dominant orientation of the element, as well as at several other angles; see Figure 8.16.




FIGURE 8.15

Crosstalk between multiple lines causing the Hough transform to detect false lines. In the case illustrated, several short segments of vertical lines are detected, in addition to the true horizontal lines.

The Hough transform has the desirable feature that it handles the occlusion of directional components gracefully, because the size of the parameter peaks is directly proportional to the number of matching points of the component. The Hough transform also has the feature that it is robust to the addition of random pixels from poor segmentation, because random image points are unlikely to contribute coherently to a single point in the parameter space.

8.6.2 The Hough and Radon transforms combined

Deans [600] showed that there is a direct relationship between the Hough and Radon transforms. The Hough transform may be viewed as a special case of the Radon transform but with a different transform origin, and performed on a binary image. Typically, the Radon transform is defined with its transform origin at the center of the original image, whereas the Hough transform is defined with its transform origin at the location of the image where the row and column indices are zero. Thus, the distance ρ as in Equation 8.42 for a 256×256 image for the Hough transform would be calculated relative to the




FIGURE 8.16

False detection of straight lines at several angles (dashed lines) within a broad linear feature by the Hough transform.

$(0, 0)$ point in the original image, whereas, for the Radon transform, the ρ value would be calculated relative to the $(128, 128)$ point; see Figure 8.14.

In the method proposed by Rangayyan and Rolston [542, 599], a Hough-Radon hybrid transform is computed by updating the (ρ_i, θ_i) parameter point by adding the pixel intensity and not by incrementing by one as with the Hough transform. In this sense, brighter lines correspond to larger peaks in the Hough-Radon domain. The Hough-Radon space is indexed from 0° to 180° along one axis, and from $-N$ to $M\sqrt{2} + N\sqrt{2}$ for an image with M rows and N columns, as shown in Figure 8.17.

The generation of the Hough-Radon space produces relative intensities of the directional features in the given image. An example of the Hough-Radon space is shown in Figure 8.18 for a simple test pattern. In directional analysis, it would be of interest to obtain the number of pixels or the percentage of the image area covered by linear segments within a particular angle band. Therefore, it is necessary to form a shadow parameter space with the numbers of the pixels that are in a particular cell in the parameter space. The shadow parameter space is the Hough transform of the image with no accompanying threshold.

It is necessary to form both the Hough-Radon transform space and the Hough-transform shadow space, because performing only the Hough transform on an unthresholded image will produce, for most images, a transform with little information about the image. Computing the shadow parameter space is the same as performing the Hough transform on a thresholded image for all pixels with a gray level greater than zero. The Hough-Radon transform, however, facilitates the differentiation between the light and dark regions in an




FIGURE 8.17

Mapping of a straight line from the image domain to the Hough-Radon space. Reproduced with permission from R.M. Rangayyan and W.A. Rolston, “Directional image analysis with the Hough and Radon transforms”, *Journal of the Indian Institute of Science*, 78: 17–29, 1998. © Indian Institute of Science.

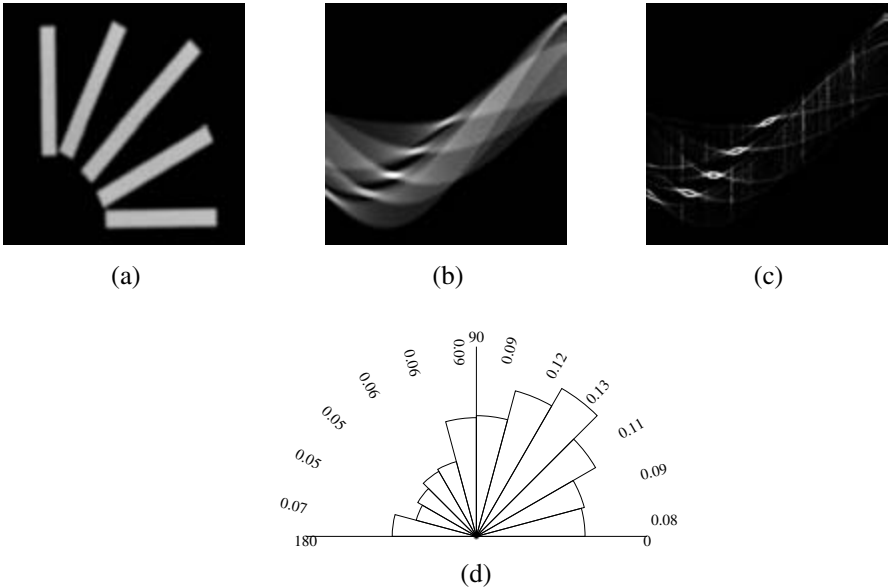


FIGURE 8.18

(a) A test image with five line segments. (b) The Hough-Radon space of the image. (c) Filtered Hough-Radon space. (d) Rose diagram of directional distribution. See also Figure 8.17. Reproduced with permission from R.M. Rangayyan and W.A. Rolston, “Directional image analysis with the Hough and Radon transforms”, *Journal of the Indian Institute of Science*, 78: 17–29, 1998. © Indian Institute of Science.

image, thus retaining all of the information about the image while performing the desired transform. The Hough transform is needed for further processing when deriving the numbers of pixels regardless of the related intensity.

From the result shown in Figure 8.18 (b), we can see the high level of crosstalk in the upper-right quadrant. From Figure 8.17, we see that this section maps to the angle band $[100^\circ, 165^\circ]$. This is due to the Hough transform's tendency to identify several lines of varying orientation within a broad linear segment, as illustrated in Figure 8.16: this is both a strength and a weakness of the Hough transform. A filtering procedure is described in the following subsection to reduce this effect.

8.6.3 Filtering and integrating the Hough-Radon space

Although the Hough-Radon transform is a powerful method for determining the directional elements in an image, it lacks, by itself, the means to eliminate elements that do not contribute coherently to a particular directional pattern. This is due to the transform being performed on all points in the given image: simple integration along a column of the transform space will include all the points in the image. A second step is needed to eliminate those pixels that do not contribute significantly to a particular pattern. Leavers and Boyce [601] proposed a simple 3×3 filter to locate maxima in the Hough space that correspond to connected collinearities in an "edge image" space.

The filter is derived from the (ρ, θ) parameterization of lines and the expected shape of the distribution of counts in the accumulator of the Hough space. For a linear element in an image, the expected shape is a characteristic "butterfly", which is a term commonly used to describe the typical falloff from a focal accumulator point as shown in Figure 8.17. It was shown by Leavers and Boyce [601] that, for any line in the image space, the extent of the corresponding butterfly in the Hough domain is limited to one radian or approximately 58° of the corresponding focal accumulator point.

The 2D filter

$$\begin{bmatrix} 0 & -2 & 0 \\ 1 & +2 & 1 \\ 0 & -2 & 0 \end{bmatrix} \quad (8.44)$$

provides a high positive response to a distribution that has its largest value at the focal point, and falls off to approximately 50% on either side, and vanishes rapidly above and below the focal point. A drawback of this filter is that it was designed for detecting peaks in the Hough space corresponding to lines of one pixel width. In the example shown in Figure 8.18 (b), we can see that the broad directional components in the test image correspond to broad peaks in the Hough-Radon domain. This results in the filter of Equation 8.44 detecting only the edges of the peaks in the Hough domain; an example of this effect is shown in Figure 8.18 (c).

The filter in Equation 8.44 is also sensitive to quantization of the θ increments. This can be seen in the vertical streaks of intensity in Figure 8.18 (c).

The vertical streaks occur at θ values that correspond to points where the value of ρ in Equation 8.43 approaches an integral value. Increasing the spatial extent of the filter would reduce the sensitivity of the filter to noise, as well as improve the ability of the filter to detect larger components in the Hough-Radon transform space. Detecting larger components in the Hough-Radon domain corresponds to detecting broad directional image components.

In the method proposed by Rangayyan and Rolston [542, 599], after the Hough-Radon transform has been filtered using the filter in Equation 8.44, the result is normalized to the range of 0.0 to 1.0 and then multiplied, point-by-point, with the shadow Hough transform mentioned earlier. This step is performed in order to obtain the relative strength of the numbers of pixels at each of the detected peaks. This step also reduces the accumulated quantization noise from the Hough-Radon transformation and the filtering steps. Although peaks may be detected in a region of the Hough-Radon domain, there may be few corresponding pixels in the original image that map to such locations. Multiplying noisy peaks by areas that contain few points in the original image will reduce the final count of pixels in the corresponding angle bands.

The final integration step is a simple summation along each of the columns of the filtered parameter space. Because the Hough transform generates a parameter space that is indexed in the column space from 0° to 180° , each of the columns represents a fraction of a degree depending upon the quantization interval selected for the transform. Also, because the Hough transform is a voting process, the peaks selected will contain some percentage of the pixels that are contained in the directional components.

Example: Figure 8.18 shows the results of the methods described above for a simple test image. The problem of the small-extent filter mentioned above is evident: the filter detects only the edges of the transformed components and neglects the central section of each of these components. From the rose diagram shown in Figure 8.18 (d), we can see that the filter has detected the relative distribution of the linear components; however, there is a large amount of smearing of the results, leading to poor differentiation between the angle bands. We can also see from the rose diagram that there is a large amount of crosstalk around 135° , where the result should be zero. This effect is probably due to a combination of crosstalk as well as quantization noise in the Hough-Radon transform.

For the ligament image shown in Figure 8.19 (a), the method has performed reasonably well. Regardless, the results contain artifacts due to the limitation of the algorithm with broad directional components, as described above.

See Rangayyan and Krishnan [360] for an application of the Hough-Radon transform for the identification of linear, sinusoidal, and hyperbolic frequency-modulated components of signals in the time-frequency plane.

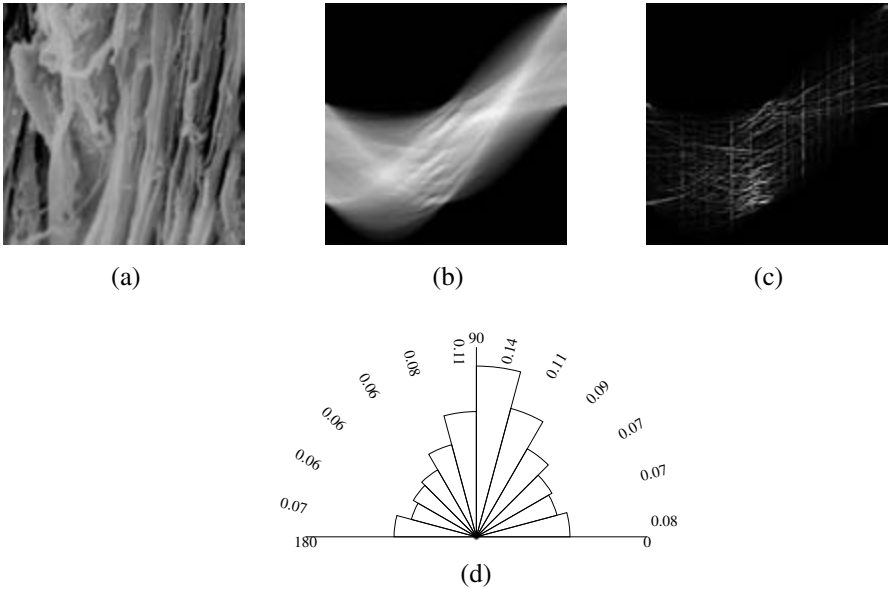


FIGURE 8.19

(a) An SEM image of a normal ligament with well-aligned collagen fibers. (b) The Hough-Radon space of the image. (c) Filtered Hough-Radon space. (d) Rose diagram of directional distribution. See also Figure 8.17. Reproduced with permission from R.M. Rangayyan and W.A. Rolston, "Directional image analysis with the Hough and Radon transforms", *Journal of the Indian Institute of Science*, 78: 17–29, 1998. © Indian Institute of Science.

8.7 Application: Analysis of Ligament Healing

The collagenous structure of ligaments [36]: Virtually all connective tissues in the human body have some type of fiber-filled matrix. The fibers consist of various sizes and shapes of chemically distinct proteins known as collagen [602], with augmentation by other fibrous materials such as elastin. There is a complex interaction between these materials and the nonfibrous “ground substance” in all tissues (water, proteoglycans, glycoproteins, and glycolipids), giving each tissue relatively unique mechanical properties. As with any composite fiber-reinforced material, the quantity and quality, as well as the organization of the reinforcing fibers, have considerable influence on the mechanical behavior of ligaments [603].

Ligaments are highly organized connective tissues that stabilize joints. Ligaments normally consist of nearly parallel arrangements of collagen fibers that are attached to bone on both sides of a joint, serve to guide the joint through its normal motions, and prevent its surfaces from becoming separated. Collagen fibers and their component fibrils make up the protenaceous “backbone” of ligaments, and provide the majority of their resistance to tensile loading. The spatial orientation of collagen fibrils is an important factor in determining tissue properties. Ligaments need to be loose enough to allow joints to move, but have to be tight enough to prevent the joint surfaces from separating.

Injuries to ligaments are common, with the normal, highly structured tissue being replaced by relatively disordered scar tissue. The scar tissue in ligaments has many quantitative and qualitative differences from the normal ligament [604], but, as with scar in other tissues [605], the relative disorganization of its collagen-fiber backbone may be among the most critical. The loose meshwork of the scar may not be able to resist tensile loads within the same limits of movement and deformation as a normal ligament. The injured or healing joint, therefore, may become loose or unstable.

The fine vascular anatomy of ligaments [414]: When a ligament is damaged by injury, the extent of the healing response determines whether and when normal function of the ligament will return. A critical factor thought to be important for the healing of a ligament is its blood supply, which exchanges oxygen, nutrients, and proteins with ligament tissue [606]. The nature of ligament vascularity has been qualitatively assessed in a few studies [607, 608, 609]. However, despite the potential importance of ligament blood-vessel anatomy, not much quantitative information is available on either normal or healing vascular anatomy of ligaments.

With respect to normal ligament vascularity, the medial collateral ligament (MCL) of the knee in a rabbit model commonly used for ligament healing studies has previously been characterized qualitatively [609]. Excluding its bony attachments, the MCL complex of the rabbit is composed of two main tissue types: epiligament and ligament tissue proper. The epiligament is a

thin layer of loose connective tissue surrounding the superficial surface of the ligament proper. Blood vessels in the normal (uninjured) ligament tissue proper appear sparse, and are oriented parallel to the long axis of the ligament in an organized fashion, whereas blood vessels in the normal epiligament appear more abundant, and are oriented in a less organized fashion [609]. In ligament scar tissue early after ligament injury, blood vessels have been described to be larger, more abundant, and more disorganized early on in the ligament healing process [607]. The need for a greater supply of materials to the ligament for early healing apparently leads to the formation of many new blood vessels, but with longer term maturation of healing tissue, the vascular supply decreases and vascularity may eventually return to normal [607].

Some of the qualitative and quantitative differences between normal and healing ligaments have been described in an animal model by Frank et al. [32]. Quantitative studies on collagen fiber organization were conducted by Chaudhuri et al. [36], Frank et al. [35], and Liu et al. [37]; the methods and results of these studies are described in Section 8.7.1. Eng et al. [414] and Bray et al. [415] conducted studies with the aim to develop a method to analyze quantitatively the variations in vascularity in normal and healing ligaments; to correlate such information with other aspects of ligament healing in a well-characterized ligament healing model; to predict ligament healing based upon the vascular response to injury; and to develop better methods to optimize ligament vascularity after injury. The related methods and results are discussed in Section 8.7.2.

8.7.1 Analysis of collagen remodeling

Tissue preparation and imaging: The animal model selected in the studies of Chaudhuri et al. [36], Frank et al. [35], and Liu et al. [37] was the ruptured and unrepaired MCL in the six-month-old female New Zealand white rabbit. Under general anesthesia and with sterile techniques, the right MCL was exposed through longitudinal medial incisions in the skin and fascia. The right MCL was completely ruptured by passing a 3D braided steel wire beneath the ligament, and failing the ligament with a strong upward pull on both ends of the suture. The left MCL was not ruptured and served as a normal control.

The injured MCL was allowed to heal for a specified period. The animal was then sacrificed by intravenous injection of 375 mg of phenobarbital, and the healing (right) and normal control (left) MCLs were harvested, as follows. The right and left MCLs were exposed through medial incisions in the skin and fascia. The MCLs were fixed *in situ* by dropping a fresh solution of 2.5% glutaraldehyde in 0.1 M cacodylate buffer with $pH = 7.4$ onto their surfaces. The MCLs were then removed at their insertions, placed in 2.5% glutaraldehyde in 0.1 M cacodylate buffer with $pH = 7.4$ for three hours, and dehydrated in increasing concentrations of ethanol (30%, 50%, 75%, and 100%). Each fixed and dehydrated ligament was then frozen quickly in liquid nitrogen, and fractured longitudinally to expose the internal collagen fiber and

component fibril arrangement along its length. The fractured tissue was then critically point dried, aligned, mounted, and sprayed with gold/ palladium. In each case, the longitudinal axis of the tissue was distinguishable at low magnification, so as to allow the orientation of high-magnification photographs relative to that axis.

Specimens were viewed under a Hitachi S-450 SEM. In order to select parts of the ligaments for imaging, pairs of (x, y) coordinates were obtained using a random number generator. In every image, the vertical axis of the photograph was aligned with the longitudinal axis of the original ligament tissue. A number of photographs were taken randomly in the midsubstance area of each of the healing and normal control MCLs at a magnification of 7,000. This magnification was experimentally chosen to give a good compromise between the resolution of the collagen fibrils and the area of the tissue being sampled. The resulting images were then digitized into 256×256 arrays.

Directional analysis: A sample image of a normal ligament is shown in Figure 8.20 (a). Parts (b) and (c) of the figure show two binarized component images obtained via directional filtering for the angle bands $75^\circ - 90^\circ$ and $0^\circ - 15^\circ$, using the sector-filtering methods described in Section 8.3.1. Directional components were obtained over 12 angle bands spanning the full range of $[0^\circ, 180^\circ]$. The fractional fiber-covered areas in the components are shown in the form a rose diagram in part (d) of the figure. The rose diagram indicates that most of the collagen fibers in the normal ligament tissue sample are aligned close to the long axis of the ligament (90° in the image plane).

A sample image of a one-week scar tissue sample is shown in Figure 8.21 (a). It is readily seen that the collagen fibers in the scar tissue do not have any dominant or preferred orientation. Two binarized component images for the angle bands $75^\circ - 90^\circ$ and $0^\circ - 15^\circ$ are shown in parts (b) and (c) of the figure. The rose diagram in part (d) of the figure shows that the angular distribution of collagen in the healing tissue is almost uniform or random.

Frank et al. [35] conducted a detailed assessment of collagen realignment in healing ligaments in response to three methods of treatment: immobilization of the affected joint for three weeks or six weeks, and no immobilization. Scar tissue samples were obtained from three rabbits each at three, six, and 14 weeks after injury, except in the second case that lacked the three-week samples. Sample images from the groups with no immobilization and immobilization for three weeks are shown in Figure 8.22. Sets of 10 images were obtained from each sample at randomly chosen locations. Composite rose diagrams were computed for each group, and are shown in Figure 8.23 for the groups with no immobilization and immobilization for three weeks. Figures 8.22 and 8.23 demonstrate the collagen remodeling or realignment process in healing ligaments.

Plots of the entropy of the rose diagrams for all the cases in the study are shown in Figure 8.24. The plots clearly demonstrate a reduction in entropy, indicating a return to orderly structure, as the healing time increases. Immobilization of the affected joint for three weeks after injury has resulted in





FIGURE 8.20

(a) A sample image showing collagen alignment in a normal ligament. Binarized directional components in the angle band (b) $75^\circ - 90^\circ$, and (c) $0^\circ - 15^\circ$. (d) Fractional fiber-covered areas in the form a rose diagram. Figure courtesy of S. Chaudhuri [610].

**FIGURE 8.21**

(a) A sample image showing collagen alignment in ligament scar tissue. Binarized directional components in the angle band (b) $75^\circ - 90^\circ$, and (c) $0^\circ - 15^\circ$. (d) Fractional fiber-covered areas in the form a rose diagram. Figure courtesy of S. Chaudhuri [610].

entropy values that are close to the values at 14 weeks in all cases, and well within the range for normal ligaments (the shaded region in Figure 8.24). The results indicate that immobilization of the affected joint for three weeks promotes the healing process, and that immobilization for the longer period of six weeks does not provide any further advantage.

Among the limitations of the methods described above, it should be noted that there is a certain degree of depth to the micrographs analyzed. The contribution of fibril components at varying depths in the different angle bands is affected by a number of factors: the intensity of the electron beam in the microscope, the depth of focus, photographic methods, and image digitization parameters. The final thresholding scheme applied to the filtered component images, being adaptive in nature, affects the various component images differently depending upon their content; this aspect cannot be controlled without introducing bias. Artifacts are also caused by tissue fixation and handling (spaces between fibrils, fiber disruption, and surface irregularities). Regardless, the results provide important quantitative information that can assist in the understanding of ligament structure and healing.

8.7.2 Analysis of the microvascular structure

Tissue preparation and imaging: In the works of Eng et al. [414] and Bray et al. [415], adult New Zealand White rabbits were used. Only 12-month-old females were used to reduce variability between animals due to age or gender differences. The selected rabbits received a “gap injury” in the right MCLs by surgical removal of a 4 *mm* segment of the tissue; see Figure 8.25. The remaining gap-injury site was marked by means of four small (6.0 nylon) sutures attached to the original cut ligament ends. The injured animals were allowed to heal for periods of three, six, 17, or 40 weeks. The MCLs from the injured animals were then removed as described below. Selected uninjured animals were used as normal controls. The right and left MCLs from these animals were removed by the procedure as described below.

The control and healing rabbits were sacrificed with an overdose of sodium pentobarbitol (Euthanyl 2.5 *cc* /4.5 *kg*). Using a constant volume pump, India-ink solution was perfused through the femoral arteries of each hind limb to vessels in the knee ligaments according to an established protocol [609]. When complete perfusion of the limbs was evident (by noting when the nail beds of claws were blackened), perfusion was stopped, and the entire hind limb was removed, placed in a container at 4° *C* for a minimum of four hours (to allow the ink solution to set), and subsequently the entire MCL was removed as shown in Figure 8.26. The left MCLs of the injured animals were not considered to be normal due to possible effects of injury to the opposite (contralateral) knee, and were used as contralateral control MCLs. The ligament scar material that developed over the gap-injury site on the right MCL was labeled as the midsubstance scar (see Figure 8.25). The two


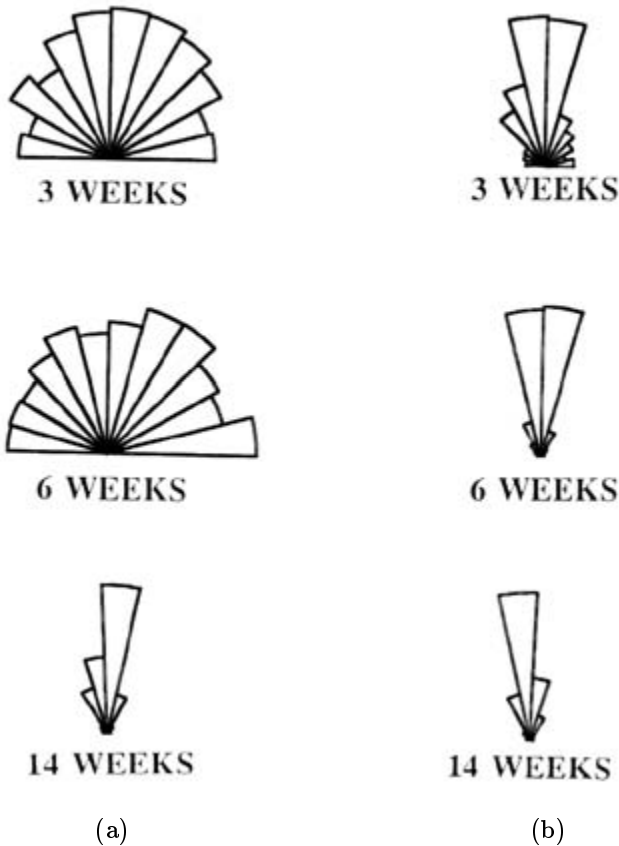


FIGURE 8.22

Sample images showing collagen alignment in ligament samples at three weeks, six weeks, and 14 weeks after injury: (a) without immobilization of the affected joint, (b) with immobilization of the affected joint for three weeks. Images courtesy of C.B. Frank. See also Figure 8.23.

**FIGURE 8.23**

Composite rose diagrams showing collagen realignment in ligament samples at three weeks, six weeks, and 14 weeks after injury: (a) without immobilization of the affected joint, (b) with immobilization of the affected joint for three weeks. See also Figure 8.22. Reproduced with permission from C.B. Frank, B. MacFarlane, P. Edwards, R. Rangayyan, Z.Q. Liu, S. Walsh, and R. Bray, "A quantitative analysis of matrix alignment in ligament scars: A comparison of movement versus immobilization in an immature rabbit model", *Journal of Orthopaedic Research*, 9(2): 219 - 227, 1991. © Orthopaedic Research Society.




FIGURE 8.24

Variation of the entropy of composite rose diagrams with collagen realignment in ligament samples at three weeks, six weeks, and 14 weeks after injury. The vertical bars indicate \pm one standard deviation about the corresponding means. “NON”: without immobilization of the affected joint; “3 IMM”: with immobilization of the affected joint for three weeks; “6 IMM”: with immobilization of the affected joint for six weeks. The shaded region indicates the range of entropy for normal ligament samples. See also Figures 8.23 and 8.22. Reproduced with permission from C.B. Frank, B. MacFarlane, P. Edwards, R. Rangayyan, Z.Q. Liu, S. Walsh, and R. Bray, “A quantitative analysis of matrix alignment in ligament scars: A comparison of movement versus immobilization in an immature rabbit model”, *Journal of Orthopaedic Research*, 9(2): 219 – 227, 1991. © Orthopaedic Research Society.

ligament sections directly connected to the midsubstance scar were labeled as the original ligament ends.

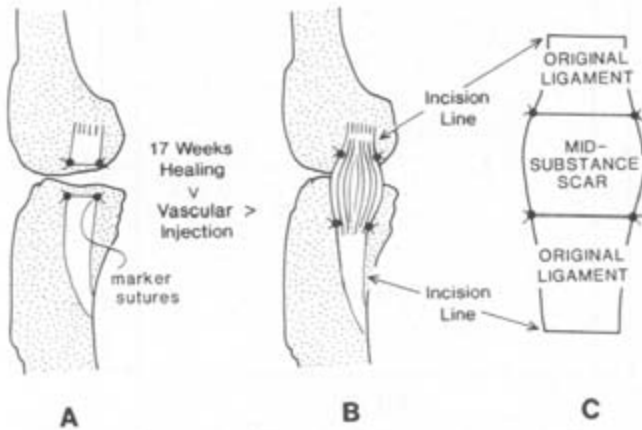


FIGURE 8.25

Gap-injury site in the ligament and the formation of scar. A: Gap injury created by removing a 4 mm section of the MCL. B: Scar after healing. C: Extracted ligament and its main regions. See also Figure 8.26. Reproduced with permission from K. Eng, R.M. Rangayyan, R.C. Bray, C.B. Frank, L. Anscomb, and P. Veale, "Quantitative analysis of the fine vascular anatomy of articular ligaments", *IEEE Transactions on Biomedical Engineering*, 39(3): 296 – 306, 1992. © IEEE.

Following removal, the ligament samples were fixed in 4% paraformaldehyde, frozen in an embedding medium, sagittally sectioned at 50 μm thickness using a Cryostat (Reichert-Jung 2800 Frigocut N), and contact-mounted maintaining proper orientation. This procedure resulted in black ink-filled blood vessels in ligament tissue, which when viewed under light microscopy showed black vessels on a clear background of ligament collagen. Every section was divided using a grid system, and under a microscope, the image of every tenth field was photographed. Photographs that did not contain visible blood vessels were discarded. The number of discarded (blank) photographs was taken into account when estimating the fractional volume of blood vessels in the ligament. The specimen magnification was measured to be 175 \times , and the scale of the digitized image was 2.32 μm per pixel.

Typical images of a normal ligament section and a 17-week healing ligament section are shown in Figure 8.27. It is evident that the normal ligament is relatively avascular, and that the blood vessels that exist are aligned along the length of the ligament (the horizontal direction of the image). On the






FIGURE 8.26

Ligament sectioning procedure for the imaging of vascular anatomy. A: knee joint. B: Extracted ligament and plane of sectioning. a: MCL complex. b: Ligament. c: Epiligament. d: Femur. e: Tibia. f: Sectioning (imaging) plane. See also Figure 8.25. Reproduced with permission from K. Eng, R.M. Rangayyan, R.C. Bray, C.B. Frank, L. Anscomb, and P. Veale, "Quantitative analysis of the fine vascular anatomy of articular ligaments", *IEEE Transactions on Biomedical Engineering*, 39(3): 296 – 306, 1992. © IEEE.

other hand, the scar tissue has a more abundant network of blood vessels to facilitate the healing process, with extensive branching and lack of preferred orientation.



(a)



(b)

FIGURE 8.27

Microvascular structure in ligaments: (a) normal; (b) 17-week scar. Images courtesy of R.C. Bray.

Binarization of the images: Images of ligament sections obtained as above contained ink-perfused blood vessels as well as collagen fibrils and other artifacts. In order to simplify the image analysis procedure, the gray-level images were thresholded to create binary images consisting of only two gray levels representing the blood vessels and the background. The gray-level histogram for a blood-vessel image was assumed to be bimodal, with the first

peak representing the pixels of blood vessels, and the second one representing the background pixels. Otsu's method (see Section 8.3.2) for threshold selection produced binary images with excessive artifacts. Threshold selection by histogram concavity analysis [611], a method that locates the locally significant minima and maxima in the gray-level histogram of the image and produces a list of possible thresholds, was investigated; however, it was difficult to choose the actual threshold to be used from the list. Another method tried was the Rutherford–Appleton threshold-selection algorithm [592], which computes a threshold by using gradient information from the image. The best threshold for the binarization of the blood-vessel images was obtained by using the Rutherford–Appleton algorithm to get a threshold estimate, followed by histogram concavity analysis to fine tune the final threshold value, as follows.

The derivatives of the given image $f(m, n)$ were obtained in the x and y directions as

$$d_x(m, n) = f(m, n + 1) - f(m, n - 1), \quad (8.45)$$

and

$$d_y(m, n) = f(m + 1, n) - f(m - 1, n). \quad (8.46)$$

The larger of the two derivatives was saved as

$$d(m, n) = \max[|d_x(m, n)|, |d_y(m, n)|]. \quad (8.47)$$

Two sums were computed over the entire image as

$$S_d = \sum \sum d(m, n), \quad (8.48)$$

and

$$S_{df} = \sum \sum d(m, n) f(m, n). \quad (8.49)$$

The Rutherford–Appleton threshold is given as

$$T_o = \frac{S_{df}}{S_d}. \quad (8.50)$$

Another potential threshold was determined by finding the position of maximal histogram concavity. A typical gray-level histogram consists of a number of significant peaks (local maxima) and valleys (local minima). Significant peaks may be identified by constructing a convex hull of the histogram, which is defined as the smallest convex polygon $\bar{h}(l)$ containing the given histogram $h(l)$, where l stands for the gray-level variable. The convex hull consists of straight-line segments joining the significant peaks in the histogram. The histogram concavity at any gray level is defined as the vertical distance between the convex hull and the histogram, that is, $[\bar{h}(l) - h(l)]$. Within each straight-line segment of the convex hull, the gray level at which the maximal concavity occurred was labeled as the optimal threshold for that segment.

Because the area covered by the blood vessels is small compared to the area covered by the background in the ligament section images, the gray-level

histogram was first scaled logarithmically to make the histogram peak representing the blood vessels and the background peak closer in height. A convex polygon of the scaled histogram was then constructed. The problem of choosing between the thresholds of each of the segments of the convex polygon was addressed by finding a threshold T_o using the Rutherford–Appleton algorithm described above. The threshold estimate T_o was found to lie between the background peak and the peak representing the blood-vessel pixels. The threshold representing the maximal histogram concavity within the convex hull segment joining these two peaks was chosen to be the threshold value T_c .

A threshold was also determined by finding the minimum point in the histogram between the peaks that represented the blood-vessel and background pixels. This threshold, labeled T_m , yielded a smaller value than T_c because of the height difference between the peaks.

Comparing the various methods of threshold selection described above, it was observed that T_c was often too high, resulting in an image with artifacts. The threshold T_m was often too low, resulting in the loss of blood-vessel pixels. By using the average of T_c and T_m as the final threshold, an acceptable compromise was reached.

A sample image of the vascular structure in a normal ligament is shown in Figure 8.28 (a). The histogram of the image with the various thresholds described above is shown in Figure 8.29. The binarized version of the image is shown in Figure 8.28 (b).

Skeletonization: Skeletonization makes directional analysis easier by reducing the binary blood-vessel patterns to their skeletal patterns with one-pixel-thick lines (see Section 6.1.6). In the work of Eng et al. [414], the binarized blood-vessel images as in Figure 8.28 (b) were reduced to their skeletons by the method described in Section 6.1.6; see Figures 8.30 and 6.13 for examples.


In order to assist the analysis of both the directionality and the volume of vascularization, an image array containing the diameter of the blood vessel at each skeleton point was formed, and referred to as the diameter-proportional skeleton of the image. The diameter at a skeleton point s_i was obtained as

$$\phi(x, y) = 2 \times \min[D(s_i, C)], \quad (8.51)$$

where C is the set of contour points of the binary image before skeletonization, and D is the Euclidean distance measure.

It is to be noted that during skeletonization, a line is shortened by one-half of its original thickness at each of its two ends. In order to correct for this during reconstruction and area estimation, pixels need to be added to the end points. Because most of the blood-vessel images were observed to have smooth contours, the ends of the line segments were assumed to be semicircular. The areas of such half circles were added at the end points.

Directional analysis: Skeletonization allows the use of the simple method of least-squares linear regression [612] to determine the angle of orientation



(a)



(b)

FIGURE 8.28

Microvascular structure in a normal ligament sample. (a) original image; (b) binarized image. See Figure 8.29 for details on the selection of the threshold for binarization. Reproduced with permission from K. Eng, R.M. Rangayyan, R.C. Bray, C.B. Frank, L. Anscomb, and P. Veale, "Quantitative analysis of the fine vascular anatomy of articular ligaments", *IEEE Transactions on Biomedical Engineering*, 39(3): 296 – 306, 1992. © IEEE.




FIGURE 8.29

Logarithmically scaled histogram of the image in Figure 8.28 (a), along with its convex hull and several possible thresholds for binarization. RATS: Rutherford–Appleton threshold-selection algorithm. Reproduced with permission from K. Eng, R.M. Rangayyan, R.C. Bray, C.B. Frank, L. Anscamb, and P. Veale, “Quantitative analysis of the fine vascular anatomy of articular ligaments”, *IEEE Transactions on Biomedical Engineering*, 39(3): 296 – 306, 1992. © IEEE.




FIGURE 8.30

Skeleton of the image in Figure 8.27 (b). See also Figure 6.13.

of each blood-vessel segment in the image. In the work of Eng et al. [414], from each point (x, y) in the skeleton image, a line segment consisting of $N = 11$ points was extracted, with the center point located at (x, y) . If (x_i, y_i) , $i = 1, 2, \dots, N$, represent the points in the line segment, the slope of the best-fitting straight line is given by

$$m = \frac{\sum_{i=1}^N x_i \sum_{i=1}^N y_i - \sum_{i=1}^N (x_i y_i)}{\left[\sum_{i=1}^N x_i \right]^2 - \sum_{i=1}^N (x_i)^2}. \tag{8.52}$$

It should be noted that, when the slope becomes large for a nearly vertical line segment, slope estimation as above becomes inaccurate due to increasing y -axis errors. This error can be obviated by adapting the least-squares formula to minimize the x -axis errors if the slope found by Equation 8.52 is greater than unity. The inverse of the slope is then given by

$$\frac{1}{m} = \frac{\sum_{i=1}^N x_i \sum_{i=1}^N y_i - \sum_{i=1}^N (x_i y_i)}{\left[\sum_{i=1}^N y_i \right]^2 - \sum_{i=1}^N (y_i)^2}. \tag{8.53}$$

The angle of the skeleton at the point (x, y) is then given by $\theta = \text{atan}(m)$.

The elemental area of the blood vessel at the point (x, y) is

$$A(x, y) = \phi(x, y) W(\theta), \tag{8.54}$$

where $\phi(x, y)$ is the vessel thickness at (x, y) as given by Equation 8.51, and

$$W(\theta) = \begin{cases} \frac{1}{\cos(\theta)} & \text{if } |\theta| < 45^\circ \\ \frac{1}{\sin(\theta)} & \text{if } |\theta| > 45^\circ \end{cases}. \tag{8.55}$$

The factor W as above (in pixels), accounts for the fact that diagonally connected pixels are farther apart than vertically or horizontally connected pixels. The elemental area was added to the corresponding angle of the histogram, and the process repeated for all points in the skeleton.

The overall accuracy of the directional analysis procedure as above was estimated to be $\pm 3^\circ$ by analyzing various test patterns. For this reason, the blood-vessel angular distributions were computed in bins of width 6° . Figure 8.31 shows composite rose diagrams obtained from 82 images from four normal ligaments and 115 images from three ligament scar samples at 17 weeks of healing. It is evident that the normal ligaments demonstrate a well-contained angular distribution of blood vessels (angular SD = 36.1° , entropy = 4.4 out of a maximum of 4.9), whereas the scar tissues show relatively widespread distribution (angular SD = 42.5° , entropy = 4.8).




FIGURE 8.31

Angular distributions of blood vessels in (a) normal ligaments (averaged over 82 images from four ligaments), and (b) 17-week scar tissues from three ligaments (115 images). Reproduced with permission from K. Eng, R.M. Rangayyan, R.C. Bray, C.B. Frank, L. Anscomb, and P. Veale, "Quantitative analysis of the fine vascular anatomy of articular ligaments", *IEEE Transactions on Biomedical Engineering*, 39(3): 296 – 306, 1992. © IEEE.

TABLE 8.2

Measures of Entropy and Standard Deviation (SD) of Composite Angular Distributions of Blood Vessels in Ligaments.

Tissue type	Ligaments	Images	Entropy	SD (°)	% Vasc.
NORMAL:					
Ligament	4	82	4.39	36.10	0.98
Epiligament	4	20	4.64	38.53	1.19
CONTRALATERAL:					
Ligament	3	93	4.33	34.79	1.05
Epiligament	3	36	4.79	42.98	2.40
SCAR:	3	115	4.79	42.52	2.50
ENDS:					
Ligament	3	80	4.59	36.55	2.24
Epiligament	3	20	4.78	44.08	3.10

The maximum possible value for entropy is 4.91. 'SCAR': midsubstance scar; 'ENDS': original ligament ends; see Figures 8.26 and 8.25. '% Vasc.': percentage of the analyzed tissue volume covered by the blood vessels detected. Reproduced with permission from K. Eng, R.M. Rangayyan, R.C. Bray, C.B. Frank, L. Anscomb, and P. Veale, "Quantitative analysis of the fine vascular anatomy of articular ligaments", *IEEE Transactions on Biomedical Engineering*, 39(3): 296 - 306, 1992. © IEEE.

In addition to the directional distributions and their statistics (entropy and angular dispersion or standard deviation), the relative volume of blood vessels in the various ligament samples analyzed were computed; see Table 8.2. Using the two-sample T-test, several assertions were arrived at about the relative volume and organization of blood vessels in normal and healing ligaments; see Table 8.3. Statistical analysis of the results indicated, with 96% confidence, that 17-week scars contain a greater volume of blood vessels than normal ligaments. Using entropy as a measure of chaos in the angular distribution of the blood-vessel segments, statistical analysis indicated, with 99% confidence, that blood vessels in 17-week scars are more chaotic than in normal ligaments.

A factor that affects the accuracy in the angular distributions derived as above is the width of the blood vessels. As the thickness of a blood vessel increases, more material is lost at the ends of the vessels during skeletonization. This loss, although corrected for by the addition of semicircular end pieces, could lead to reduced accuracy of the angular distribution. Sampling and quantization errors become significant when the thickness of blood vessels is small.

TABLE 8.3

Results of Statistical Comparison of the Relative Volume of Vascularization (V) and the Entropy of the Angular Distribution (H) of Various Ligament Samples.

Assertion	Confidence (%)
LIGAMENT:	
V (normal) < V (contralateral)	70
V (normal) < V (midsubstance scar)	96
V (normal) < V (original ligament ends)	85
V (original ligament ends) < V (midsubstance scar)	55
H (contralateral) < H (normal)	73
H (normal) < H (midsubstance scar)	99
H (normal) < H (original ligament ends)	53
H (original ligament ends) < H (midsubstance scar)	96
EPILIGAMENT:	
V (normal) < V (contralateral)	99
V (normal) < V (original ligament ends)	70
H (normal) < H (contralateral)	90
H (normal) < H (original ligament ends)	82

Reproduced with permission from K. Eng, R.M. Rangayyan, R.C. Bray, C.B. Frank, L. Anscumb, and P. Veale, "Quantitative analysis of the fine vascular anatomy of articular ligaments", *IEEE Transactions on Biomedical Engineering*, 39(3): 296 - 306, 1992. © IEEE.

It should be observed that blood vessels branch and merge in 3D within the ligament. The sectioning procedure used to obtain 2D slices imposes a limitation in the analysis: the segments of the blood vessels that traverse across the sectioning planes are lost in the procedures for directional analysis.

The increased blood-vessel volume and greater directional chaos observed in scar tissue as compared to normal ligaments indicate that the blood supply pattern is related to the healing process. Interestingly, after a 17-week healing period, increases in both the blood-vessel volume and dispersion were also observed in the MCL from the opposite knee (contralateral), as compared to the uninjured, normal MCL [414, 415]. The directional chaos of the contralateral MCL decreased in the ligament region, but increased in the epiligament region of the tissue as compared to the normal tissue. This may be attributed to a possible change in loading conditions on the contralateral limb after injury, or to a nervous response to injury transmitted by neural pathways [613]. As expected, both the vascularity and the directional chaos of the blood vessels in the original ligament ends increased as compared to the normal. This shows that injury to one portion of the tissue has some effect on the vascularity of connected ligament tissue.

8.8 Application: Detection of Breast Tumors

The detection of breast tumors is difficult due to the nature of mammographic images and features. Many studies have focused on image processing techniques for the segmentation of breast parenchyma in order to detect suspicious masses and reject false positives. The detection of masses requires the segmentation of all possible suspicious regions, which may then be subjected to a series of tests to eliminate false positives.

Early work by Winsberg et al. [614] using low-resolution imagery showed promise in detecting large solitary lesions in mammograms. Later on, several studies dealt with the detection of suspicious areas in xeromammograms. Ackerman and Gose [615] developed computer techniques for categorizing suspicious regions marked by radiologists based upon features in xeromammograms. Kimme et al. [616] proposed an automatic procedure for the detection of suspicious abnormalities on xeromammograms by identifying breast tissues and partitioning them into at most 144 sections per image. Ten normalized statistics for each section were used as texture features, and the classification of 2,270 mammographic sections of eight patients with six best-performing features yielded a false-positive rate of 26% and a false-negative rate of 0.6%. Hand et al. [617] and Semmlow et al. [618] reported on automated methods to detect suspicious regions in xeromammograms. The methods included routines for the detection of the breast boundary and the extraction of suspi-

cious areas. Classification of normal and abnormal regions was achieved using global and regional features. With a dataset of 60 xeroradiograms from 30 patients, the methods proposed by Hand et al. [617] correctly identified 87% of the suspicious areas presented, but resulted in a high false-positive rate.

Lai et al. [619] presented a method to detect circumscribed masses. They enhanced the contrast in mammograms using a selective median filter as a preprocessing step, and proposed a method based upon template matching to identify candidate regions. Finally, two tests (local neighborhood test and region histogram test) were applied to reduce the number of false positives. The method was effective in the detection of specific types of lesions, but was not generally applicable to all types of masses. The masses detected included both benign and malignant types; however, the results of detection were not reported separately for the individual mass types.

Multiscale methods: The features of masses in mammograms vary greatly in size and shape. Several computer techniques [620, 621] have, therefore, employed multiscale concepts for the detection of masses in mammographic images. Brzakovic et al. [620] proposed fuzzy pyramid linking for mass localization, and used intensity links of edge pixels, in terms of their position, at various levels of resolution of the image; relative stability of the links from one level of resolution to another was assumed. They reported 95% detection accuracy with 25 cases containing benign masses, malignant tumors, and normal images (the numbers of each type were not specified). They further reported to have achieved 85% accuracy in classifying the regions detected as benign, malignant, or nontumor tissue by using features based upon tumor size, shape, and intensity changes in the extracted regions. Details about the range of the size of the masses detected and their distribution in terms of circumscribed and spiculated types were not reported.

Chen and Lee [348] used multiresolution wavelet analysis and expectation-maximization (EM) techniques in conjunction with fuzzy C-means concepts to detect tumor edges. The method was tested on only five images, and classification of masses as benign or malignant was not performed. Li et al. [622] developed a segmentation method based upon a multiresolution Markov random field model, and used a fuzzy binary decision tree to classify the regions segmented as normal tissue or mass. The algorithm was reported to have achieved 90% sensitivity with two false-positive detections per image. Using a nonlinear multiscale approach, Miller and Ramsey [341] achieved an accuracy of 85% in detecting malignant tumors in a screening dataset.

Qian et al. [262, 270, 623] developed three image processing modules for computer-assisted detection of masses in mammograms: a tree-structured nonlinear filter for noise suppression, a multiorientation directional wavelet transform, and a multiresolution wavelet transform for image enhancement to improve the segmentation of suspicious areas. They performed wavelet decomposition of the original images, and used the lowpass-smoothed images for detecting the central core portions of spiculated masses. Wavelet-based adaptive directional filtering was performed on the highpass-detail images to

enhance the spicules in the mass margins. Chang and Laine [624] used coherence and orientation measures in a multiscale analysis procedure to enhance features and provide visual cues to identify lesions.

Density-based methods: Masses are typically assumed to be hyperdense with respect to their surroundings. However, due to the projection nature of mammograms, overlapped fibroglandular tissues could also result in high-intensity regions in the image, leading to their detection as false positives. Therefore, many studies [371, 617, 618, 625, 626] focused on detecting potential densities as an initial step, and incorporated modules to reject the false positives at a later stage. Woods and Bowyer [625] detected potential densities by using a local measure of contrast, and supplemented the method with a region-growing scheme to find the extents of the potential densities. They further used a set of features computed from each region in linear and quadratic neural classifiers to classify the regions segmented as masses or false positives. In a later work, Woods and Bowyer [627] examined five mass-detection algorithms and described a general framework for detection algorithms composed of two steps: pixel-level segmentation and region-level classification. Their review reveals some fundamental advantages in concentrating effort on pixel-level analysis for achieving higher detection accuracies.

Kok et al. [628] and Cerneaz [626] referred to a mass-detection approach that was developed by Guissin and Brady [629] based upon isointensity contours, and reported that the measures that Guissin and Brady used to reduce false positives are not suitable for mammography. Cerneaz and Brady [626, 630] proposed methods to remove curvilinear structures including blood vessels, milk ducts, and fibrous tissues prior to the detection of significant densities in mammograms. However, the assumption that such a step would not affect the spicules of tumors is questionable.

In the detection scheme proposed by Cerneaz [626], dense regions (blobs) in the mammogram were initially segmented by combining the methods proposed by Guissin and Brady [629] and Lindeberg [631]. Novelty analysis methods were then applied to separate mass regions from the plethora of dense regions thus segmented. Three of the five features used in the novelty analysis step included variance in the intensity of a blob's pixels, average blob height, and a measure of a blob's saliency. A set of 100 images (no details were mentioned about the nature of the distribution of the images) from the MIAS database [376] were used for evaluating the methods after reducing the resolution to $300 \mu\text{m}$. The detection accuracies were not stated explicitly. The detection approach proposed by Mudigonda et al. [275] and described in the subsections to follow possesses similarities with the method of Guissin and Brady [629] and Cerneaz and Brady [630, 626] in the initial stage of segmentation of isolated regions in the image.

Analysis of texture of mass regions: Petrick et al. [632] reported on the use of a two-stage adaptive density-weighted contrast enhancement filter in conjunction with a LoG edge detector for the detection of masses. In their approach, the original images at a resolution of $100 \mu\text{m}$ were downsampled

by a factor of eight to arrive at images of size 256×256 pixels and $800 \mu\text{m}$ resolution. Then, for each potential mass object, an ROI was extracted from the corresponding downsampled image using the bounding box of the object to define the region. A set of texture features based upon the GCMs yielded a true-positive detection rate of 80% at 2.3 false positives per image, and 90% detection accuracy at 4.4 false positives per image with a dataset of 168 cases [633].

Kobatake et al. [634] applied the iris filter for the detection of approximately rounded convex regions, and computed texture features based upon GCMs of the iris filter's output to isolate malignant tumors from normal tissue. The methods resulted in a detection rate of 90.4% with 1.3 false positives per image, with a dataset of 1,214 CR images containing 208 malignant tumors.

Kegelmeyer [635] developed a method to detect stellate lesions in mammograms, and computed Laws' texture features from a map of local edge orientations. A binary decision tree was used to classify the features. Detection results with five test images yielded a sensitivity of 83% with 0.6 false findings per image. Gupta and Undrill [636] used Laws' texture measures and developed a texture-based segmentation approach to identify malignant tumors.

Researchers have also used fractal measures for the detection and characterization of mammographic lesions. Priebe et al. [637] used texture and fractal features for the detection of developing abnormalities. Burdett et al. [471] used fractal measures and nonlinear filters for the characterization of lesion diffusion. Byng et al. [448] proposed measures of skewness of the image brightness histogram and the fractal dimension of image texture, which were found to be strongly correlated with a radiologist's subjective classifications of mammographic patterns.

Gradient-based analysis of mass regions: Several published algorithms for breast mass segmentation are based upon the analysis of the gradient orientation at edges to locate radiating lines or spicules [339, 518, 625, 638, 639, 640]. Karssemeijer and te Brake [641] reported on the use of scale-space operators and line-based pixel orientation maps to detect spiculated distortions. The map of pixel orientations in the image was used to construct two operators sensitive to star-like patterns of lines. A total of 50 images (nine spiculated masses, 10 cases of architectural distortion, and 31 normal cases) from the MIAS database [376] were analyzed. By combining the output from both the operators in a classifier, an accuracy of 90% was achieved in detecting stellate carcinomas at one false positive per image [641]. The mass cases studied belonged to the spiculated malignant category only, and did not include circumscribed malignant or circumscribed benign cases.

In a related study, te Brake and Karssemeijer [642, 643] used three pixel-based mass-detection methods, including the method they developed for detecting stellate distortions, to examine if the detection of masses could be achieved at a single scale. Experiments with simulated masses indicated that little could be gained by applying their methods at a number of scales. Us-

ing a dataset of 60 images (30 malignant and 30 normals) from the MIAS database [376], they reported that their gradient orientation method [339, 638, 641] performed better than two other methods based upon template matching and Laplacian filtering [642]. Karssemeijer and te Brake mentioned that their gradient orientation method requires the optimization of several parameters in order to achieve a good detection accuracy [641].

Kobatake and Yoshinaga [644] developed skeleton analysis methods using the iris filter to detect spicules of lesions. They used a modified Hough transform to extract radiating lines from the center of a mass region to discriminate between star-shaped malignant tumors and nonmalignant masses, and obtained an accuracy of 74% in detecting malignant tumors using a dataset of 34 CR images including 14 malignant, nine benign, and 11 normal cases.

Polakowski et al. [645] developed a model-based vision algorithm using DoG filters to detect masses and computed nine features based upon size, circularity, contrast, and Laws' texture features. A multilayer perceptron neural network was used for the classification of breast masses as benign or malignant. With a dataset of 36 malignant and 53 benign cases, they reported a detection sensitivity of 92% in identifying malignant masses, with 1.8 false positives per image.

Directional analysis of fibroglandular tissues: Several researchers [639, 646, 647, 648, 649, 650, 651, 652] have developed methods to analyze the orientation of fibroglandular tissues in order to detect mass regions and eliminate false positives. Zhang et al. [639] employed a Hough-spectrum-based technique to analyze the texture primitives of mammographic parenchymal patterns to detect spiculated mass regions and regions possessing architectural distortion. With a dataset of 42 images obtained from 22 patients, the methods yielded a sensitivity of 81% with 2.0 false positives per image. Parr et al. [640] studied Gabor enhancement of radiating spicules, and concluded that the linearity, length, and width parameters are significantly different for spicules in comparison to other linear structures in mammograms. Later on, the group reported on the detection of linear structures in digital mammograms by employing a multiscale directional line operator [647, 648, 650]. Two images indicating the probability of suspicion were derived by applying PCA and factor analysis techniques to model the orientations present in central core regions and the surrounding patterns of lesions, respectively. A sensitivity of 70% at 0.01 false positives per image was reported by combining the evidence present in both the probability images in a k-nearest-neighbor approach using a dataset of 54 mammograms containing 27 spiculated lesions and 27 normal mammograms [651]. The basis for the computation of the result mentioned above was not clarified: a value of 0.01 false positives per image with a dataset of 54 mammograms leads to a total of less than one false positive detection. The sizes of the abnormalities in their studies ranged from 5 mm to 30 mm, with a mean of 13.4 mm; 77% of the lesions tested were smaller than 15 mm. Mudigonda et al. [275] applied techniques of texture flow-field analysis [432, 653] to analyze oriented patterns in mammograms by

computing the angle of anisotropy or dominant orientation and the strength of flow-like information or coherence at every point in the image: the related methods and results are described in the following subsections.

Analysis of bilateral asymmetry: Several studies [371, 654, 655, 656, 657] have investigated the bilateral asymmetry between the left and right mammograms of an individual in order to localize breast lesions. Such studies have reported on the registration and unwarping transformations that are required for the comparison of different images. Apart from the anatomical differences that exist between the left and right breasts of an individual, registration methods will have to cope with the additional complexities arising due to the variation in the amounts of compression applied in the process of obtaining the images, as well as the differences in the angles of the projections. The transformations used in registration methods often face limitations in adequately addressing the above-mentioned complexities of the mammographic imaging process.

Lau and Bischof [371] applied a transformation based upon the outline of the breast to identify asymmetry in breast architecture. Using a B-spline model of the breast outline to normalize images, they compared features including brightness, roughness, and directionality to define asymmetry measures for breast tumor detection. Although they reported success in the detection of tumors, they warned that the method by itself is not reliable for clinical application, due mainly to the high rate of false negatives. Giger et al. [654] and Yin et al. [655] used similar warping procedures to align bilateral images to perform subtraction. Their method exploits the normal bilateral symmetry of healthy parenchyma to label regions of significant difference as potential masses. Nishikawa et al. [658] analyzed asymmetric densities to detect masses by performing nonlinear subtraction between the right and left breasts; their computer-aided scheme for the classification of masses using artificial neural networks achieved a higher accuracy than radiologists.

Miller and Astley [372, 659] proposed a method for the detection of asymmetry by comparing anatomically similar and homogeneous regions; the procedure included segmentation, classification as fat or nonfat region, texture energy, and shape features such as compactness, circularity, and eccentricity. Training and assessment were carried out on a leave-one-out basis with a dataset of 52 mammogram pairs, achieving 72% correct classification with a linear discriminant classifier. Ferrari et al. developed methods to achieve segmentation of the outline of the breast [279], the pectoral muscle [278], and the fibroglandular disc [280] in mammograms, and analyzed the glandular tissue patterns by applying directional filtering concepts to detect bilateral asymmetry [381]. Their methods, described in Section 8.9, avoid the application of registration procedures that are associated with most of the above-mentioned methods to analyze asymmetry. Instead, they performed a global analysis of the directional distributions of the glandular tissues using Gabor wavelets to characterize the asymmetry in tissue flow patterns.

Analysis of prior mammograms: It is known that a small but significant number of cancer cases detected in screening programs have prompts visible in earlier screening examinations [660]. These cases represent screening errors or limitations, and may occur due to the lack of an adequate understanding of the perceptual features of early breast abnormalities as apparent on mammograms. The above observations have prompted many researchers [661, 662, 663, 664, 665, 666, 667, 668] to analyze the previous or prior mammograms taken as part of routine screening and follow-up studies of patients diagnosed with cancer, in an effort to detect the disease in the prior mammograms. Brzakovic et al. [661] developed methods for detecting changes in mammograms of the verified positive group in the regions labeled as abnormal by a medical expert, by comparing them with the features derived from the corresponding regions of the previous screenings. Their procedures included segmentation, partitioning into statistically homogeneous regions, and region-based statistical analysis in order to achieve registration and perform comparison between mammograms acquired at different instances of screening.

Sameti et al. [664, 669] studied the structural differences between the regions that subsequently formed malignant masses on mammograms, and other normal areas in images taken in the last screening instance prior to the detection of tumors. Manually identified circular ROIs were transformed into their optical density equivalents, and further divided into three discrete regions representing low, medium, and high optical density. Based upon the discrete regions, a set of photometric and texture features was extracted. They reported that in 72% of the 58 breast cancer cases studied, it was possible to realize the differences between malignant mass regions and normal tissues in previous screening images.

Petrick et al. [632, 667] studied the effectiveness of their mass-detection method in the detection of masses in prior mammograms. The dataset used included 92 images (54 malignant and 38 benign) from 37 cases (22 malignant and 15 benign). Their detection methods achieved a “by film” mass-detection sensitivity of 51% with 2.3 false positives per image. They achieved a slightly better accuracy of 57% in detecting only malignant tumors. Their detection scheme attempts to segment salient densities by employing region growing after enhancement of contrast in the image. Such an intensity-based segmentation approach fails to detect the developing densities in the previous screening images due to the inadequate contrast of mass regions before the masses are actually formed.

Several semiautomated segmentation schemes [632, 633, 276, 518, 407] have used manually segmented ROIs in order to search for masses in a specified region of the breast. Such methods find limited practical utility in a screening program.

The female breast is a complex organ made up of fibrous, glandular, fatty, and lymphatic tissues. The differences in density information of the breast tissues are captured in a mammogram in the form of intensity and textural

variations. Mudigonda et al. [275] proposed an unsupervised segmentation approach to localize suspicious mass regions in mammographic images. The approach aims to isolate the spatially interconnected structures in the image to form regions concentrated around prominent intensities. It would then be possible to extract high-level information characterizing the physical properties of mass regions, and to short-list suspicious ROIs for further analysis. A block diagram of this approach is shown in Figure 8.32; the various steps of the detection algorithm are explained in detail in the following subsections.

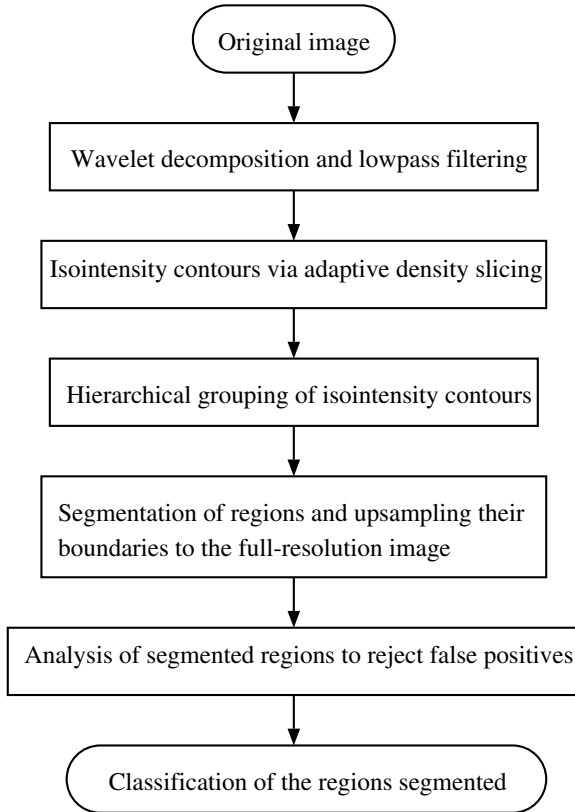


FIGURE 8.32

Block diagram of the mass-detection algorithm. Figure courtesy of N.R. Mudigonda [166].


8.8.1 Framework for pyramidal decomposition

Malignant tumors, due to their invasive nature, possess heterogeneous density distributions and margins causing distortion in the orientation of the surrounding tissues. In order to detect such structures as single entities, prior smoothing of the image is required. Mudigonda et al. [275] employed recursive wavelet decomposition and Gaussian smoothing operations in a multiresolution pyramidal architecture as preprocessing steps to achieve the required level of smoothing of the image.


A pyramidal representation of the given image was obtained by iterative decimation operations on the full-resolution image, thereby generating a hierarchy of subimages with progressively decreasing bandwidth and increasing scale [670, 671]. Wavelet decomposition divides the frequency spectrum of the original image f into its lowpass-subband-equivalent image f_L and highpass-equivalent detail image f_H at different scales. The lowpass-subband image at each scale, produced by decimating its preceding higher-resolution image present in the hierarchy by an octave level, was further smoothed by a 3×3 Gaussian kernel, and the resulting image was stretched to the range of 0 – 60 in pixel value. The wavelet used was a symlet of eighth order. Symlets are compactly supported wavelets with the least asymmetry and the highest number of vanishing moments for a given support width [672]. Figure 8.33 shows plots of the decomposition lowpass kernels used with symlets, at two different scales. The wavelet decomposition was performed recursively to three octave levels using the symlets mentioned above.

The preprocessing steps of wavelet decomposition and Gaussian smoothing operations described above successively and cumulatively modulate the intensity patterns of mass regions to form smooth hills with respect to their surroundings in low-resolution images. Figure 8.34 (a) shows a $1,024 \times 1,024$ section of a mammogram containing two circumscribed benign masses. Parts (b) – (d) of the figure show the corresponding low-resolution images after the first, second, and third levels of decomposition, respectively. The effects of the preprocessing steps mentioned above may be observed in the low-resolution images.

The choice of the wavelet, the width of the kernel used for lowpass filtering, and the degree or scale factor of decomposition can influence the smoothed results. The preprocessing operations described above were employed to arrive at an estimate of the extent of isolated regions in a low-resolution image, and studies were not performed with different sets of choices. However, satisfactory smoothed results were obtained with the wavelet chosen due to its symmetry. A scale factor of three, which causes the decomposition of the original $50 \mu m/pixel$ images to a resolution of $400 \mu m/pixel$, was found to be effective on most of the images tested: decomposition to a higher scale resulted in over-smoothing of images and merged multiple adjoining regions into single large regions, whereas a scale factor of two yielded insignificant regions due to

**FIGURE 8.33**

Plots of symlet decomposition lowpass filters at two scales. Figure courtesy of N.R. Mudigonda [166].

**FIGURE 8.34**

(a) A $1,024 \times 1,024$ section of a mammogram containing two circumscribed benign masses. Pixel size = $50 \mu m$. Image width = $51 mm$. Low-resolution images obtained by wavelet filtering: (b) After the first level of decomposition; 512×512 pixels, $100 \mu m$ per pixel. (c) After two levels of decomposition; 256×256 pixels, $200 \mu m$ per pixel. (d) After three levels of decomposition; 128×128 pixels, $400 \mu m$ per pixel. The intensity of the filtered images has been enhanced by four times for display purposes. Figure courtesy of N.R. Mudigonda [166].

insufficient smoothing. However, some researchers [632] have performed mass detection after reducing images to a resolution of $800 \mu\text{m}/\text{pixel}$.

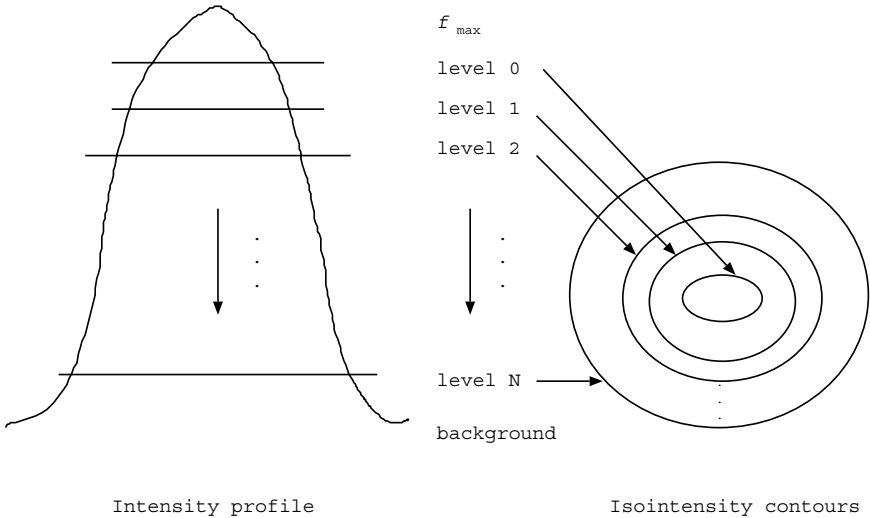
8.8.2 Segmentation based upon density slicing

The recursive smoothing and decimation operations described above result in a gradual modulation of intensity information about the local intensity maxima present in various isolated regions in the low-resolution image. As a result, the intensity levels are expected to assume either unimodal or bimodal histogram distributions. The next step in the algorithm is to threshold the image at varying levels of intensity to generate a map of iso-intensity contours [673]. The purpose of this step is to extract concentric groups of closed contours to represent the isolated regions in the image.

The density-slicing or intensity-slicing technique slices the given image (represented as a 2D intensity function) by using a plane that is placed parallel to the coordinate plane of the image [526]. A level curve (also known as an iso-intensity curve) is then formed by extracting the boundary of the area of intersection of the plane and the intensity function. Figure 8.35 shows a schematic illustration of the density-slicing operation. Each level curve obtained using the procedure explained above is guaranteed to be continuous and closed. The number of levels of thresholding, starting with the maximum intensity in the image, and the step-size decrement for successive levels, were adaptively computed based upon the histogram distribution of the image under consideration, as explained below.

Let f_{max} represent the maximum intensity level in the low-resolution image (which was scaled to 60), and let f_{th} be the threshold representing the mass-to-background separation, which is to be derived from the histogram. It is assumed that the application of the preprocessing smoothing operations results in exponentially decreasing intensity from the central core region of a mass to its background, represented as $f_{\text{th}} = f_{\text{max}} \exp[-\mu N]$, where N is the number of steps required for the exponentially decreasing intensity function to attain the background level represented by f_{th} , $N = (f_{\text{max}} - f_{\text{th}})$, and μ is the intended variation in step size between the successive levels of thresholding. The step size μ may be computed through a knowledge of the parameters f_{th} and N . The threshold f_{th} was derived from the histogram, and corresponds to the intensity level representing the maximum number of occurrences when the histogram assumes a unimodal distribution.

It is essential to set bounds for f_{th} so as not to miss the detection of masses with low-density core regions, while maintaining the computational time of the algorithm at a reasonable level. A large threshold value might miss the grouping of low-intensity mass regions, thereby affecting the sensitivity of the detection procedure; on the other hand, a low value would result in a large map of iso-intensity contours, and increase the computational load on the algorithm in further processing of the contours. A smaller threshold could also result in large numbers of false detections. Initial estimates of f_{th} derived from the

**FIGURE 8.35**

Schematic illustration of the density-slicing operation. f_{\max} represents the maximum intensity in the image, and levels $0, 1, 2, \dots, N$ represent a set of N threshold values used for density slicing. Figure courtesy of N.R. Mudigonda [166].

corresponding histograms of low-resolution images were observed to range between 50% and 90% of f_{\max} , and N was observed to range between 10 and 30. The parameter f_{th} was adaptively selected based upon the histogram as explained below:

1. If $0.5 f_{\max} < f_{\text{th}} \leq 0.9 f_{\max}$, f_{th} could be assumed to represent the mass-to-background transition, and the same threshold value is retained.
2. If $f_{\text{th}} > 0.9 f_{\max}$, the mass regions that are to be detected in the image are expected to be merged with the surrounding background, and no distinct central core regions would be present. In such cases, f_{th} is considered to be $0.9 f_{\max}$, and N is set to 30 (the maximum number of levels of thresholding considered) to limit the step-size increments of the level function to a low value. These steps facilitate close tracking of difficult-to-detect mass-to-background demarcation.
3. If $f_{\text{th}} \leq 0.5 f_{\max}$, f_{th} might not represent the true mass-to-background transition, and hence, is ignored. An alternative search for f_{th} is initiated so that the value obtained will lie in the upper half of the histogram distribution.

The steps described above realize a domain of isointensity contours in the low-resolution image.

8.8.3 Hierarchical grouping of iso-intensity contours

The next step in the algorithm is to perform grouping and elimination operations on the framework of closed contours generated in the low-resolution image, considering their parent-child nodal relations in a family-tree architecture. A schematic representation of such a hierarchical grouping procedure is shown in Figure 8.36, which depicts the segmentation of the low-resolution image into three isolated regions based upon three concentric groups of contours.

The strategy adopted was to short-list at first the possible central dense-core portions, which are usually small in size but of higher density (represented by f_{\max} in each group of contours in Figure 8.36), and to identify the immediate low-density parent members encircling them. The process was continued until all the members in the available set of closed contours in the image were visited. Each of the closed contours was assigned to a specific group or family of concentric contours based upon nodal relations, thus leading to segmentation of the image into isolated regions. A concentric group of contours represents the propagation of density information from the central core portion of an object in the image into the surrounding tissues. In some images with dense and fatty backgrounds, the outermost contour members were observed to contain multiple regions of dissimilar structures. For this reason, a specified number of outer contours were discarded to separate the groups of contours representing adjacent structures.

The outermost contour in each family or group and the family count in terms of the number of contours present could be useful in the analysis of the regions segmented in order to reject false positives. Masses, irrespective of their size, were observed to result in a higher family count as compared to elongated glandular tissues. By setting a threshold on the family count, chosen to be five, dense glandular structures could be avoided from further analysis. A lower threshold value for the minimum-allowable family count was observed to affect the specificity in terms of the number of false positives detected, but not affect the sensitivity of the detection procedure. Finally, the outermost contour from each of the short-listed groups was upsampled to the full-resolution image to form the corresponding segmented area.

8.8.4 Results of segmentation of masses

The results of application of the algorithm to the image shown in Figure 8.34 (a) are presented in Figure 8.37 (a). The contour map and the outermost contours detected are shown superimposed on the low-resolution image (at a scale of three). Figure 8.38 shows the histogram of the corresponding low-resolution and smoothed image. Figure 8.37 (b) shows the contours (white) upsampled from the low-resolution image in Figure 8.37 (a) to the full-resolution image of Figure 8.34 (a); the corresponding contours (black) that were manually drawn by an expert radiologist are overlaid for comparison.

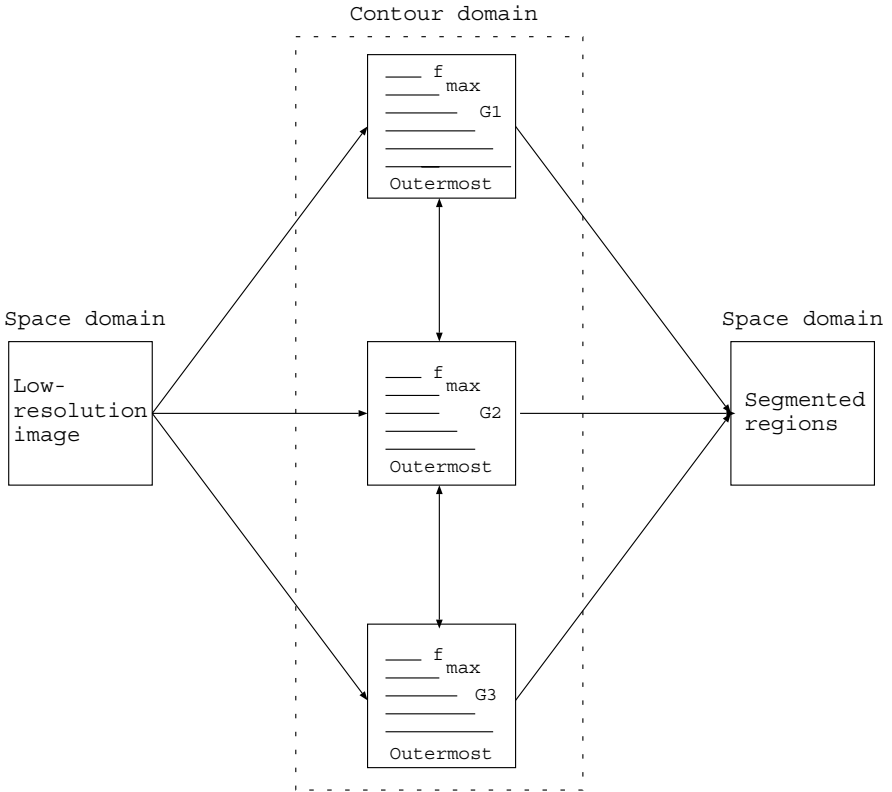


FIGURE 8.36

Schematic representation of hierarchical grouping of contours. G1, G2, and G3 are groups of contours that represent isolated regions in the image. Reproduced with permission from N.R. Mudigonda, R.M. Rangayyan, and J.E.L. Desautels, "Detection of breast masses in mammograms by density slicing and texture flow-field analysis", *IEEE Transactions on Medical Imaging*, 20(12): 1215 – 1227, 2001. © IEEE.


Figure 8.39 shows a similar set of results of application of the mass-detection algorithm to a $1,024 \times 1,024$ section of a mammogram containing a spiculated malignant tumor.

As can be seen from Figures 8.37 and 8.39, the upsampled contours in the full-resolution image contain the most significant portions of the corresponding masses, and are in close agreement with the corresponding areas manually delineated by the radiologist. The results of application of the methods discussed above to full-size mammograms are presented in the subsections to follow, along with methods based upon texture flow-field principles for detailed analysis of the various regions segmented in order to reject false positives.

The mass-detection algorithm was tested on segments of size up to $2,048 \times 2,048$ pixels of 39 mammographic images (28 benign and 11 malignant) from the MIAS database [376], with a spatial resolution of $50 \mu\text{m} \times 50 \mu\text{m}$. In 29 of the 39 cases (19 benign and 10 malignant), the segmented regions were in agreement with the corresponding regions that were manually identified by the radiologist. In six images, including five images with circumscribed benign masses and an image with a spiculated malignant tumor, the regions segmented by the algorithm were not in agreement with the corresponding regions manually delineated by the radiologist. The radiologist indicated that he encountered difficulty while tracing the boundaries of the masses in some images from the MIAS database. In the remaining four images where the method failed, all belonging to the spiculated benign category, the mass portions are merged in fatty and glandular background. In these images, the technique failed to generate a contour map with the specified minimum number of concentric closed contours to be able to delineate the mass regions. In two of the images, including a circumscribed benign mass and a spiculated benign mass, the masses are located close to the edges of the images, and the process of generation of concentric closed contours was impeded.

Overall, the mass-detection algorithm performed well on images containing malignant tumors, and successfully segmented tumor areas that were in agreement with the corresponding regions identified manually by the radiologist. However, the method encountered limited success in images with benign masses. In the detection scheme, only the contour map generated in the lowest-resolution image is analyzed to segment the mass regions, and the information available in the intermediate-resolution images along the hierarchy is not considered. Establishment of reliable intensity links through the intermediate-resolution images may result in improved detection results.

Benign-versus-malignant pattern classification was carried out using the BMDP 7M stepwise discriminant analysis program [674] with texture features computed based upon averaged GCMs for the 29 masses (19 benign and 10 malignant) that were successfully segmented by the mass-detection procedure. (See Sections 7.3.2 and 7.9.1 for details on the computation of texture features using adaptive ribbons.) Four effective features including entropy, second moment, second difference moment, and correlation were short-listed. The



(a)






Figure 8.37 (b)

FIGURE 8.37

(a) Groups of isointensity contours and the outermost contour in each group in the third low-resolution image of the mammogram section of Figure 8.34 (d). (b) The contours (white) of two masses (indicated by arrows) and two false positives detected in the full-resolution image of Figure 8.34 (a), with the corresponding contours (black) of the masses drawn independently by a radiologist. Reproduced with permission from N.R. Mudigonda, R.M. Rangayyan, and J.E.L. Desautels, "Segmentation and classification of mammographic masses", *Proceedings of SPIE Volume 3979, Medical Imaging 2000: Image Processing*, pp 55 – 67, 2000. © SPIE.

**FIGURE 8.38**

Histogram of the low-resolution and smoothed image shown in Figure 8.37 (a). Reproduced with permission from N.R. Mudigonda, R.M. Rangayyan, and J.E.L. Desautels, "Segmentation and classification of mammographic masses", *Proceedings of SPIE Volume 3979, Medical Imaging 2000: Image Processing*, pp 55 – 67, 2000. © SPIE.



(a)





Figure 8.39 (b)



(c)




Figure 8.39 (d)

FIGURE 8.39

(a) A $1,024 \times 1,024$ section of a mammogram containing a spiculated malignant tumor. Pixel size = $50 \mu\text{m}$. Image width = 51 mm . (b) Group of iso-intensity contours and the outermost contour in the group in the third low-resolution image. (c) Histogram of the low-resolution and smoothed image shown. (d) The contour (white) of the spiculated malignant tumor detected in the full-resolution image, superimposed with the corresponding contour (black) drawn independently by a radiologist. Reproduced with permission from N.R. Mudigonda, R.M. Rangayyan, and J.E.L. Desautels, "Segmentation and classification of mammographic masses", *Proceedings of SPIE Volume 3979, Medical Imaging 2000: Image Processing*, pp 55 – 67, 2000. © SPIE.

GCM-based texture features computed from the mass ribbons resulted in an average classification efficiency of 0.80.

8.8.5 Detection of masses in full mammograms

Masses containing important signs of breast cancer may be difficult to detect as they often occur in dense glandular tissue. Successful identification of such difficult-to-detect masses often results in a large number of false positives. Rejection of false positives forms an important part of algorithms for mass detection [341, 629, 632, 634, 635, 638, 643, 644, 645, 664, 675].

In the algorithm proposed by Mudigonda et al. [275] to detect masses, the pyramidal decomposition approach described in the preceding subsections was extended for application to full mammograms. Furthermore, the orientation information present in the margins of the regions detected was analyzed using texture flow-field principles to reject false positives. The approach, described below, is significant in the following aspects:

- The methods constitute a comprehensive automated scheme for the detection of masses, analysis of false positives, and classification of mammographic masses as benign or malignant. The detection methods proposed are not limited to the analysis of any specific category of masses; instead, they cover a wide spectrum of masses that include malignant tumors as well as benign masses of both the circumscribed and spiculated categories.
- As described at the beginning of this section, many of the recently published research works [639, 641, 643, 650, 651, 652] have noted the significance of analyzing the oriented information in mammograms in identifying regions that correspond to abnormal distortions in the images. Such methods require precise computation of orientation estimates. Mudigonda et al. [275] introduced methods to analyze oriented textural information in mammograms using flow-field principles, and proposed features to differentiate mass regions from other dense tissues

in the images in order to reduce the number of false positives. The flow-field methods employed use a strong analytical basis in order to provide optimal orientation estimates [653].

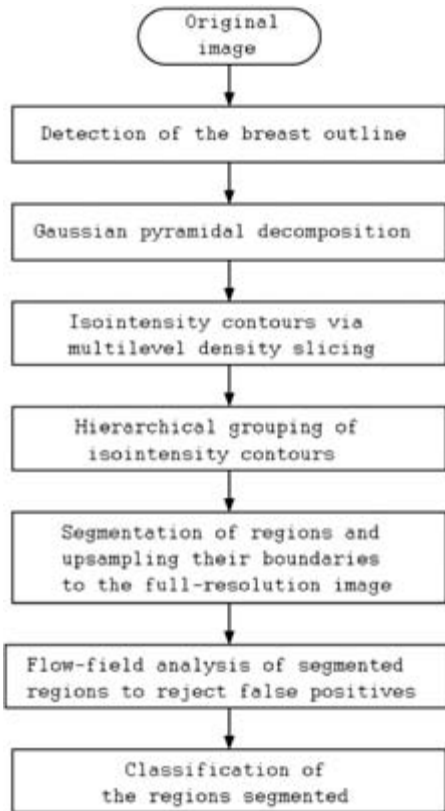
- As discussed in Section 7.9.1, the analysis of textural information in ribbons of pixels across the boundaries of masses has been found to be effective in the benign-versus-malignant discrimination of masses [165, 451, 676]. The studies cited above computed textural features using ribbons of pixels of a fixed width of 8 mm. In the work of Mudigonda et al. [275], a method is proposed to estimate the widths of the ribbons to be able to adapt to variations in the size and shape of the detected mass regions. The adaptive ribbons of pixels extracted are used to classify the regions detected as masses or false positives at first, and subsequently to discriminate between benign masses and malignant tumors.
- The features used for the classification of masses and false positives are based upon specific radiographic characteristics of masses as described by an expert radiologist.

The block diagram shown in Figure 8.40 lists the various steps of the detection algorithm, which are explained in detail in the following paragraphs.

Detection of the breast boundary: In order to limit further processing to only the breast region, an approximate outline of the breast was detected initially by employing the following steps. The image was smoothed with a separable Gaussian kernel of width 15 pixels (pixel width = 200 μm), and quantized to 64 gray levels. The method proposed by Schunck [677] and Rao and Schunck [653] was used to generate Gaussian kernels. Figure 8.41 shows a plot of the Gaussian kernel used.

A map of iso-intensity contours was generated by thresholding the image using a threshold close to zero. From the map of iso-intensity contours, a set of closed contours was identified by employing the chain code [526]. The contour containing the largest area was then considered to be the outline of the breast. Figure 8.42 illustrates a mammogram of size 1,024 \times 1,024 pixels with a spiculated malignant tumor. The outline of the breast detected in the mammogram of Figure 8.42 is shown in Figure 8.43. The method successfully detected the outlines of all of the 56 images tested. The method worked successfully with images lacking skin-air boundaries all around as well.

Detection of salient densities: Gaussian pyramidal decomposition was employed to achieve the required smoothing instead of wavelet decomposition that was used in the case of detection of masses in sectional images of mammograms as discussed in Section 8.8.1. Gaussian decomposition, when tested with some of the previously used sectional mammographic images, provided comparable detection results in terms of the boundaries of the regions segmented. Gaussian smoothing using separable kernels has the advantage of ease of implementation, and also provides computational advantages, particularly with full-size mammograms.

**FIGURE 8.40**

Block diagram of the algorithm for the detection of masses in full mammograms. Reproduced with permission from N.R. Mudigonda, R.M. Rangayyan, and J.E.L. Desautels, "Detection of breast masses in mammograms by density slicing and texture flow-field analysis", *IEEE Transactions on Medical Imaging*, 20(12): 1215 – 1227, 2001. © IEEE.




FIGURE 8.41

Plot of a Gaussian kernel with the support width of 15 pixels. The width at half-maximum height is five pixels. Figure courtesy of N.R. Mudigonda [166].

The original 8 *b* images with a spatial resolution of $200\ \mu\text{m}$ were subsampled to a resolution of $400\ \mu\text{m}$ after performing smoothing with a separable Gaussian kernel of width five pixels. The width of the Gaussian kernel at half-maximum height is about $400\ \mu\text{m}$, and hence, is not expected to cause excessive smoothing of mass regions because mass features in mammograms typically span a few millimeters.

The preprocessing steps described above are essential in order to capture the complete extent of mass features as single large regions, so as to facilitate adequate inputs for further discrimination between masses and false positives. Masses were assumed to be hyperdense, or at least of the same density, with respect to their background. The preprocessing steps described above may not be effective with masses that do not satisfy this assumption.

Multilevel thresholding: In the procedure of Mudigonda et al. [275], the low-resolution image is initially reduced to 64 gray levels in intensity and thresholded at $N = 30$ levels starting from the maximum intensity level $f_{\text{max}} = 64$, with a step-size decrement of $\mu = 0.01 f_{\text{max}}$. The purpose of this step is to extract concentric groups of closed contours to represent the isolated regions in the image as explained earlier. The above-mentioned parameters were chosen based upon the observation of the histograms of several low-resolution images. The histogram of the low-resolution image obtained by preprocessing the mammogram in Figure 8.42 is shown in Figure 8.44. As indicated in the figure, the intensity level at which the masses and other dense




FIGURE 8.42

A mammogram (size $1,024 \times 1,024$ pixels, $200 \mu m$ per pixel) with a spiculated malignant tumor (radius = 2.28 cm). Case mdb184 from the MIAS database [376]. Reproduced with permission from N.R. Mudigonda, R.M. Rangayyan, and J.E.L. Desautels, "Detection of breast masses in mammograms by density slicing and texture flow-field analysis", *IEEE Transactions on Medical Imaging*, 20(12): 1215 – 1227, 2001. © IEEE.




FIGURE 8.43

The map of isointensity contours extracted in the smoothed and subsampled version (size 512×512 pixels, $400 \mu m$ per pixel) of the mammogram shown in Figure 8.42. The breast outline detected is superimposed. In some cases, several contours overlap to produce thick contours in the printed version of the image. Reproduced with permission from N.R. Mudigonda, R.M. Rangayyan, and J.E.L. Desautels, "Detection of breast masses in mammograms by density slicing and texture flow-field analysis", *IEEE Transactions on Medical Imaging*, 20(12): 1215 – 1227, 2001. © IEEE.

tissues appear to merge with the surrounding breast parenchyma is around the minimum threshold level of 44.




FIGURE 8.44

Histogram of the low-resolution image corresponding to the mammogram in Figure 8.42. Reproduced with permission from N.R. Mudigonda, R.M. Rangayyan, and J.E.L. Desautels, “Detection of breast masses in mammograms by density slicing and texture flow-field analysis”, *IEEE Transactions on Medical Imaging*, 20(12): 1215 – 1227, 2001. © IEEE.

Figure 8.43 shows the map of iso-intensity contours obtained by density slicing or multilevel thresholding of the mammogram shown in Figure 8.42. (Observe that multiple concentric contours may fuse into thick contours in the printed version of the image.)

Grouping of iso-intensity contours: The scheme represented in Figure 8.36 was adopted to perform a two-step grouping and merging operation on the individual contours possessing a minimum circumference of 2 mm (five pixels at 400 μm), to arrive at groups of concentric iso-intensity contours. Initially, the contour members with intensity values ranging from $0.8 f_{\text{max}}$ to f_{max} , with $f_{\text{max}} = 64$, were grouped to form a set of regions corresponding to high intensities in the image, and then the remaining contour members were grouped into a separate set. The undesired merging of adjoining regions

was controlled by monitoring the running family count of each group for any abrupt fluctuations in terms of its family count. The information from both the sets of groups of contours was combined by establishing correspondences among the outermost members of the various groups present in each set to arrive at the final set of segmented regions in the low-resolution image. The largest contour in each group thus finalized with a minimum family count of two members was upsampled into the full-resolution image to form the corresponding segmented area. The segmented regions identified in the full-resolution image were analyzed using the features described above to identify true-positive and false-positive regions.

8.8.6 Analysis of mammograms using texture flow-field

In a mammogram of a normal breast, the fibroglandular tissues present oriented and flow-like or anisotropic textural information. Mudigonda et al. [275] proposed features to discriminate between masses and the strongly oriented fibroglandular tissues based upon the analysis of oriented texture in mammograms. The method proposed by Rao and Schunck [432, 653], briefly described in the following paragraphs, was used to characterize flow-like information in the form of intrinsic orientation angle and coherence images. The intrinsic angle image reveals the direction of anisotropy or flow orientation of the underlying texture at every point in the image. Coherence is a measure of the degree or strength of anisotropy in the direction of flow.

Rao and Schunck [653] and Rao [432] made a qualitative comparison of the performance of their method with the method proposed by Kass and Witkin [678], and reported that their method achieved superior results in characterizing flow-field information. However, it appears that their implementation of Kass and Witkin's scheme differed from the method that was originally proposed by Kass and Witkin. Regardless, the method of Rao and Schunck has a strong analytical basis with fewer assumptions.

The methodology to derive the intrinsic images begins with the computation of the gradient information at every point in the image by preprocessing the image with a gradient-of-Gaussian filter of a specified width. The impulse response of a 2D Gaussian smoothing filter $g(x, y)$ of width σ is given by

$$g(x, y) = \exp \frac{-(x^2 + y^2)}{2\sigma^2}, \quad (8.56)$$

where the scale factor has been ignored. The impulse response of the gradient-of-Gaussian filter $h(x, y)$ tuned to a specified orientation Θ can be obtained using $g(x, y)$ as

$$h(x, y) = \left[\frac{\partial g}{\partial x}, \frac{\partial g}{\partial y} \right] \bullet [\cos \Theta, \sin \Theta], \quad (8.57)$$

where \bullet represents the dot product. At each point in the given image, the filter $h(x, y)$, upon convolution with the image, yields the maximal response in

the orientation (Θ) that is perpendicular to the orientation of the underlying texture (that is, the angle of anisotropy). Based upon the above, and with the assumption that there exists a dominant orientation at every point in the given image, Rao and Schunck [653] derived the optimal solution to compute the angle of anisotropy ψ_{pq} at a point (p, q) in the image as described below.

Let G_{mn} and θ_{mn} represent the gradient magnitude and gradient orientation at the point (m, n) in an image, respectively, and $P \times P$ be the size of the neighborhood around (p, q) used for computing ψ_{pq} . The gradient magnitude is computed as

$$G_{mn} = \sqrt{G_x^2(m, n) + G_y^2(m, n)}, \tag{8.58}$$

where $G_x(m, n)$ and $G_y(m, n)$ represent the outputs of the gradient-of-Gaussian filter at (m, n) in the x and y directions, respectively. The gradient orientation is computed as

$$\theta_{mn} = \arctan \left(\frac{G_y(m, n)}{G_x(m, n)} \right). \tag{8.59}$$

The projection of G_{mn} on to the gradient orientation vector at (p, q) at angle θ_{pq} is $G_{mn} \cos(\theta_{mn} - \theta_{pq})$, as illustrated schematically in Figure 8.45.

Based upon the discussion above, the sum-of-squares S of the projections of the gradient magnitudes computed at the various points of the neighborhood in a reference orientation specified by Θ is given by

$$S = \sum_{m=1}^P \sum_{n=1}^P G_{mn}^2 \cos^2(\theta_{mn} - \Theta). \tag{8.60}$$

The sum S varies as the orientation Θ is varied, and attains its maximal value when Θ is perpendicular to the dominant orientation that represents the underlying texture in the given set of points. Differentiating S with respect to Θ yields

$$\frac{dS}{d\Theta} = 2 \sum_{m=1}^P \sum_{n=1}^P G_{mn}^2 \cos(\theta_{mn} - \Theta) \sin(\theta_{mn} - \Theta). \tag{8.61}$$

By setting $\frac{dS}{d\Theta} = 0$ and further simplifying the result, we obtain the solution for $\Theta = \Theta_{pq}$ that maximizes S at the point (p, q) in the image as

$$\Theta_{pq} = \frac{1}{2} \arctan \left(\frac{\sum_{m=1}^P \sum_{n=1}^P G_{mn}^2 \sin 2\theta_{mn}}{\sum_{m=1}^P \sum_{n=1}^P G_{mn}^2 \cos 2\theta_{mn}} \right). \tag{8.62}$$

The second derivative $\frac{d^2S}{d\Theta^2}$ is given by

$$\frac{d^2S}{d\Theta^2} = -2 \sum_{m=1}^P \sum_{n=1}^P G_{mn}^2 \cos(2\theta_{mn} - 2\Theta). \tag{8.63}$$




FIGURE 8.45

Schematic illustration of the projection of the gradient magnitude for computing the dominant orientation angle and coherence (the scheme of Rao and Schunck [653]). G_{pq} and θ_{pq} indicate the gradient magnitude and orientation at (p, q) , respectively. The corresponding parameters at (m, n) are G_{mn} and θ_{mn} . The size of the neighborhood shown is $P \times P = 5 \times 5$ pixels. Figure courtesy of N.R. Mudigonda [166].

The value of Θ_{pq} that is obtained using Equation 8.62 represents the direction of the maximal gradient output, because the second derivative shown in Equation 8.63 is negative at $\Theta = \Theta_{pq}$ when the texture has only one dominant orientation. The estimated orientation angle of flow ψ_{pq} at (p, q) in the image is then

$$\psi_{pq} = \Theta_{pq} + \frac{\pi}{2}, \quad (8.64)$$

because the gradient vector is perpendicular to the direction of flow. The angles computed as above range between 0 and π radians.

In order to analyze mammograms with the procedure described above, the original image was initially smoothed using a separable Gaussian kernel [677] of a specified width, and the gradients in the x and y directions were computed from the smoothed image using finite differences in the respective directions. The choice of the width of the Gaussian affects the gradient computation; a width of 2.2 mm (11 pixels) was used by Mudigonda et al. [275], in relation to the range of the size of features related to breast masses. The filter has a width of about 1 mm at its half-maximum height. This filter size is appropriate given that mammograms may demonstrate lumps that are as small as 3 mm in diameter.

The gradient estimates computed as above were smoothed using a neighborhood of size 15×15 pixels (3×3 mm), the width of which was chosen to be larger than the Gaussian that was initially used to compute the gradient estimates. Figure 8.46 shows the intrinsic angle image of the mammogram shown in Figure 8.42. The bright needles, overlaid on the image, indicate the underlying dominant orientation at points spaced every fifth row and fifth column, computed using Equation 8.62. Needles have been plotted only for those pixels where the coherence, computed as follows, is greater than zero.

The coherence γ_{pq} at a point (p, q) in the given image was computed as the cumulative sum of the projections of the gradient magnitudes of the pixels in a window of size $P \times P$, in the direction of the dominant orientation at the point (p, q) under consideration, as

$$\gamma_{pq} = G_{pq} \frac{\sum_{m=1}^P \sum_{n=1}^P G_{mn} \cos(\theta_{mn} - \psi_{pq})}{\sum_{m=1}^P \sum_{n=1}^P G_{mn}}. \quad (8.65)$$

The result was normalized with the cumulative sum of the gradient magnitudes in the window, and multiplied with the gradient magnitude at the point under consideration in order to obtain high coherence values at the points in the image having high visual contrast. The coherence image computed for the mammogram in Figure 8.42 is shown in Figure 8.47. It can be observed that glandular tissues, ligaments, ducts, and spicules corresponding to architectural distortion possess high coherence values.




FIGURE 8.46

Intrinsic angle information (white lines) for the mammogram shown in Figure 8.42. The boundaries (black) represent the mass and false-positive regions segmented at the initial stage of the mass-detection algorithm. The breast outline detected is superimposed. Reproduced with permission from N.R. Mudigonda, R.M. Rangayyan, and J.E.L. Desautels, "Detection of breast masses in mammograms by density slicing and texture flow-field analysis", *IEEE Transactions on Medical Imaging*, 20(12): 1215 – 1227, 2001. © IEEE.



FIGURE 8.47

Intrinsic coherence image of the mammogram shown in Figure 8.42. Reproduced with permission from N.R. Mudigonda, R.M. Rangayyan, and J.E.L. Desautels, “Detection of breast masses in mammograms by density slicing and texture flow-field analysis”, *IEEE Transactions on Medical Imaging*, 20(12): 1215 – 1227, 2001. © IEEE.

8.8.7 Adaptive computation of features in ribbons

The regions detected by the method described above vary greatly in size and shape. For this reason, a method was devised to compute adaptively the width of the ribbon for the derivation of features (see Section 7.9.1), or equivalently, the diameter of the circular morphological operator for a particular region based upon the region's size and shape.

Figure 8.48 shows a schematic representation of the method used to compute adaptively the size of the ribbon. Initially, the diameter of the bounding circle enclosing a given candidate region was found by computing the maximal distance between any two points on its boundary. Then, the areas of the region (A_r) and the bounding circle (A_c) enclosing the region were computed. The width of the ribbon was computed as

$$R_w = R_c \frac{A_r}{A_c}, \quad (8.66)$$

where R_c is the radius of the bounding circle. The ratio $\frac{A_r}{A_c}$ is a simple measure of the narrowness and shape complexity of the region. The size of the ribbon computed above was limited to a maximum of 8 mm or 40 pixels. The regions for which the sizes of ribbons computed was less than 0.8 mm or four pixels were rejected, and not processed further in the false-positive analysis stage. The ribbons of pixels (white) extracted across the boundaries (black) of the various regions detected in the image shown in Figure 8.42 are illustrated in Figure 8.49.




FIGURE 8.48

Schematic representation of the adaptive computation of the width of the ribbon. A_r : area of the candidate region, A_c : area of the bounding circle, and R_c : radius of the bounding circle. Figure courtesy of N.R. Mudigonda [166].




FIGURE 8.49

Ribbons of pixels (white) extracted adaptively across the boundaries (black) of the regions detected in the mammogram shown in Figure 8.42. Reproduced with permission from N.R. Mudigonda, R.M. Rangayyan, and J.E.L. Desautels, "Detection of breast masses in mammograms by density slicing and texture flow-field analysis", *IEEE Transactions on Medical Imaging*, 20(12): 1215 – 1227, 2001. © IEEE.

Features for mass-versus-false-positive classification: In order to classify the regions detected as true masses or false positives, the following features were proposed by Mudigonda et al. [275], based upon certain well-established radiological notions about breast masses as apparent on mammograms.

- Contrast (C_{fg}) : Masses in mammograms may be presumed to be hyperdense, or at least isodense, with respect to their surroundings. For this reason, the contrast (C_{fg}) of a region was computed as the difference between the mean intensities of the foreground region or ROI, and a background region defined as the region enclosed by the extracted ribbon of pixels excluding the ROI. Regions possessing negative contrast values were rejected from further analysis.
- Coherence ratio (γ_r) : The interior regions of masses are expected to be less coherent than their edges. The ratio (γ_r) of the mean coherence of the ROI (excluding the ribbon of pixels) to the mean coherence in the ribbon of pixels was used as a feature in pattern classification.
- Entropy of orientation estimates (H_o) : The orientation of spicules in the margins of spiculated masses is usually random. Furthermore, the orientation estimates computed in the margins of circumscribed masses could cover a wide range of angles between zero and π radians, and may not possess any dominant orientation. On the contrary, fibroglandular tissues are highly directional. For these reasons, the entropy (H_o) of the orientation estimates was computed in the ribbon of pixels of each region detected for use as a feature in pattern classification.
- Variance of coherence-weighted angle estimates (σ_h^2) : The fourth feature was based upon the coherence-weighted angular histogram, which was computed for a particular region by incrementing the numbers of occurrence of angles with the magnitudes of coherence values computed at the respective points, after resampling the angle values in the ribbon regions to $Q = 6$ equally spaced levels between zero and π . This is equivalent to obtaining a cumulative sum of the coherence estimates of the points belonging to each bin as the height of the corresponding bin in the histogram. The histogram distributions obtained as above were normalized with the cumulative sum of the coherence values computed in the ribbons of the respective regions, and the variance (σ_h^2) was computed as

$$\sigma_h^2 = \frac{1}{Q} \sum_{i=1}^Q (\alpha_i - \mu_h)^2, \quad (8.67)$$

where α_i , $i = 1, 2, \dots, Q$, are the normalized values of the heights of the Q bins of the histogram formed by the coherence-weighted angle

estimates, and μ_h is the average height of the bins of the histogram:

$$\mu_h = \frac{1}{Q} \sum_{i=1}^Q \alpha_i . \quad (8.68)$$

Features for benign-versus-malignant classification: The efficacy in benign-versus-malignant classification of the true-positive mass regions successfully segmented by the mass-detection algorithm was evaluated by using a set of five GCM-based texture features: entropy, second moment, difference moment, inverse difference moment, and correlation (see Section 7.3.2 for details). The features were computed in the ribbon of pixels extracted adaptively from each segmented mass margin as described above. The GCMs constructed by scanning each mass ribbon in the 0° , 45° , 90° , and 135° directions were averaged to obtain a single GCM, and the five texture features were computed for the averaged GCM for each ribbon.

8.8.8 Results of mass detection in full mammograms

Mudigonda et al. [275] tested their methods with a total of 56 images (each of size $1,024 \times 1,024$ pixels at a resolution of $200 \mu m$) including 30 benign masses, 13 malignant tumors, and 13 normal cases selected from the Mini-MIAS [376] database. The dataset included circumscribed and spiculated cases in both of the benign and malignant categories. The mean values of the sizes of the masses were $1.07 \pm 0.77 \text{ cm}$ and $1.22 \pm 0.85 \text{ cm}$ for the benign and malignant categories, respectively. The radius of the smallest mass (malignant) was 0.34 cm , and that of the largest mass (benign) was 3.9 cm . The center of abnormality and an approximate radius of each mass are indicated in the database. The circular demarcation of masses as done in the database is not useful for confirming the results of mass detection, because such a demarcation may also include normal fibroglandular tissues in the ROIs, particularly in spiculated cases. Hence, only the center of the abnormality as indicated for each mass in the database was used to confirm the result of mass detection.

The mass-detection algorithm successfully detected all of the 13 malignant tumors in the database used. However, the algorithm met with limited success in detecting benign masses. In 11 (five circumscribed and six spiculated) of the 30 benign cases tested, the algorithm failed to detect the masses. The overall detection accuracy was 74% with a total of 43 cases.

A close inspection of some of the benign cases in which the algorithm failed to detect the mass revealed the following details. In three of the four dense-glandular masses (cases labeled as mdb244, mdb290, and mdb315 in the Mini-MIAS database [376]), two fatty-glandular masses (mdb017 and mdb175), and a fatty mass (mdb069), the masses do not have prominent central core regions and possess poor contrast with respect to their background. Contrast enhancement [123] of such mammograms prior to the detection step

could improve the performance of the detection algorithm. In a mammogram containing a fatty mass (mdb190), the high intensity in the pectoral muscle region affected the multilevel thresholding process of the detection algorithm. This calls for methods to detect precisely the pectoral muscle [278] (see Section 5.10), and a two-step detection procedure: initially, the mammographically apparent regions corresponding to suspicious lymph nodes could be searched for inside the pectoral muscle region, and in the second stage, the region of the breast excluding the pectoral muscle area could be searched for the possible presence of masses.

In two other cases (mdb193 and mdb191), the masses did not satisfy the hyperdense or isodense assumption. Successful detection of masses in such cases may require additional methods based upon the asymmetry between the right and the left mammograms, in order to detect regions possessing architectural distortion [381, 595, 679, 680, 681].

In the method proposed by Mudigonda et al. [275], no region was rejected based upon its size during the initial stage, because masses can be present with any size. Instead, the emphasis was on reducing false positives by using texture flow-field features. As a result, a large number of false-positive regions, at approximately 11 per image, were detected by the algorithm along with the true mass regions during the initial stage of detection.

Mass-versus-false-positive classification: The four features C_{fg} , γ_r , H_o , and σ_h^2 , described in Section 8.8.7, were computed in the ribbons of the candidate regions that were detected in all of the 56 cases tested, and used in a linear discriminant classifier to identify the true mass regions and false positives. The MIAS database contains an unusual number of spiculated benign cases [345]. In order to study the effect of such an atypical distribution of cases on the accuracy of the detection method, pattern classification experiments were carried out in two stages: at first, a mass-versus-normal-tissue classification was conducted with the 671 regions detected in the 56 cases tested. Next, malignant-tumor-versus-normal-tissue classification was performed using the features computed from the 343 regions detected in the 13 malignant and the 13 normal cases tested.

Pattern classification was carried out using the BMDP 7M stepwise discriminant analysis program with the leave-one-out scheme [674]. In datasets with limited numbers of cases, it is difficult to form separate sets of data for training and testing purposes. The leave-one-out cross-validation scheme helps to obtain the least-biased estimates of classification accuracy in such situations. The overall classification efficiency in the classification of malignant tumors versus normal tissue was 0.9, and that for discriminating between masses (both benign and malignant) and normal tissue was 0.87.

The linear discriminant function obtained at equal prior probability values for the group with tumors and the group with normal cases was encoded to arrive at a mass-versus-false-positive decision for each segmented region. Figure 8.50 indicates the final set of regions that were retained for the image in Figure 8.42 after the detection and false-positive analysis stages. The region

detected inside the pectoral muscle area as shown in Figure 8.50 was suspected to be a lymph node affected by the invasive carcinoma. Such areas may be prompted to a radiologist for studying possible nodal involvement, particularly in cases with no localizing signs of the disease.

Figure 8.51 shows a mammogram with a spiculated malignant tumor (indicated by an arrow, radius = 0.54 cm) that is smaller and less obvious than the tumor shown in Figure 8.42. Figures 8.52 and 8.53 show the results of detection for the image shown in Figure 8.51 before and after the false-positive analysis stage, respectively. The tumor has been successfully detected, along with one false positive.

The mass-versus-normal-tissue classification experiment, involving the 32 mass regions (19 benign and 13 malignant) that the algorithm successfully detected and 639 false positives from a total of 56 images (including 13 normal cases), resulted in an overall classification efficiency of 0.87, with a sensitivity of 81% at 2.2 false positives per image. A total of six masses (four benign and two malignant) were misclassified as normal tissue. However, if the fact that the algorithm missed 11 benign masses during the initial stage of detection itself is taken into consideration, the true detection sensitivity of the algorithm with the database of 30 benign and 13 malignant masses reduces to 60% (26/43).

In the case of malignant-tumor-versus-normal-tissue classification, a high overall classification efficiency of 0.9 was achieved; the dataset included 13 malignant tumors and 330 false positives from a total of 26 images (including 13 normal cases). A sensitivity of 85% was obtained at 2.46 false positives per image. Although all of the 13 tumors were successfully detected in the initial stage, two of the malignant tumors that were detected were misclassified later as normal tissue, yielding a small proportion (2/13) of false negatives.

In a related work, te Brake and Karssemeijer [643] compared three mass-detection schemes to detect malignant tumors of small size (radius smaller than 6 mm), medium size (radius between 6 mm and 10 mm), and large size (radius greater than 10 mm). The method based upon gradient-orientation maps [641] was reported to have achieved the best results in the detection of masses in the small and medium categories. The study of te Brake and Karssemeijer focused only on malignant tumors and did not include benign masses.

Petrick et al. [632] reported similar trends in detection results using a dataset of 25 mammograms containing 14 benign and 11 malignant cases. A sensitivity of 96% was reported at 4.5 false positives per image. Most of the other previous studies [634, 643, 645] in the related field reported on the detection of only malignant tumors, and did not specifically consider the detection of benign masses.

Benign-versus-malignant classification: The effectiveness of the segmentation results in benign-versus-malignant pattern classification was verified using the five GCM-based texture features computed based upon averaged GCMs as explained above for the 32 cases (19 benign and 13 malignant)




FIGURE 8.50

Adaptive ribbons of pixels (white) and boundaries (black) of the regions retained in the mammogram shown in Figure 8.42 after the false-positive analysis stage. The larger region corresponds to the malignant tumor; the other region is a false positive. See also Figure 8.49. Reproduced with permission from N.R. Mudigonda, R.M. Rangayyan, and J.E.L. Desautels, "Detection of breast masses in mammograms by density slicing and texture flow-field analysis", *IEEE Transactions on Medical Imaging*, 20(12): 1215 – 1227, 2001. © IEEE.




FIGURE 8.51

A mammogram (size $1,024 \times 1,024$ pixels, $200 \mu\text{m}$ per pixel) with a spiculated malignant tumor (pointed by the arrow, radius = 0.54 cm). Case mdb144 from the MIAS database [376]. Reproduced with permission from N.R. Mudigonda, R.M. Rangayyan, and J.E.L. Desautels, "Detection of breast masses in mammograms by density slicing and texture flow-field analysis", *IEEE Transactions on Medical Imaging*, 20(12): 1215 – 1227, 2001. © IEEE.




FIGURE 8.52

Ribbons of pixels (white) extracted adaptively across the boundaries (black) of the regions detected in the mammogram shown in Figure 8.51. Reproduced with permission from N.R. Mudigonda, R.M. Rangayyan, and J.E.L. Desautels, "Detection of breast masses in mammograms by density slicing and texture flow-field analysis", *IEEE Transactions on Medical Imaging*, 20(12): 1215 – 1227, 2001. © IEEE.




FIGURE 8.53

Adaptive ribbons of pixels (white) and boundaries (black) of the regions retained in the mammogram shown in Figure 8.51 after the false-positive analysis stage. The larger region corresponds to the malignant tumor; the other region is a false positive. See also Figure 8.52. Reproduced with permission from N.R. Mudigonda, R.M. Rangayyan, and J.E.L. Desautels, "Detection of breast masses in mammograms by density slicing and texture flow-field analysis", *IEEE Transactions on Medical Imaging*, 20(12): 1215 – 1227, 2001. © IEEE.

that were successfully segmented by the mass-detection procedure. Pattern classification was carried out using the BMDP stepwise logistic regression program [674]. The five GCM-based texture features resulted in an overall classification efficiency of 0.79. The results obtained confirm that the mass regions segmented in images of resolution $200\ \mu\text{m}$ possess adequate discriminant information to permit their classification as benign or malignant with texture features. Similar benign-versus-malignant classification results were obtained using partial images of the same cases, but with $50\ \mu\text{m}$ resolution. It appears that the detection and classification of masses may be successfully performed using images of resolution $200\ \mu\text{m}$ with the techniques described above.

8.9 Application: Bilateral Asymmetry in Mammograms

Asymmetry between the left and right mammograms of a given subject is an important sign used by radiologists to diagnose breast cancer [54]. Analysis of asymmetry can provide clues about the presence of early signs of tumors (parenchymal distortion, small asymmetric bright spots and contrast, etc.) that are not evaluated by other methods [372]. Several works have been presented in the literature addressing this problem [371, 372, 682, 683, 684], with most of them applying some type of alignment of the breast images before performing asymmetry analysis. However, alignment procedures applied to mammograms have to confront many difficult problems, such as the natural asymmetry of the breasts of a given subject, the absence of good corresponding points between the left and right breast images to perform matching, and the distortions inherent to breast imaging.

Procedures for systematic analysis were proposed by Lau and Bischof [371] and Miller and Astley [372] to perform comparison of the corresponding anatomical regions between the left and right breast images of an individual in terms of shape, texture, and density. Lau and Bischof [371] also proposed a directional feature to quantify oriented patterns.

Ferrari et al. [381] proposed a procedure based upon directional analysis using Gabor wavelets in order to analyze the possible presence of global disturbance between the left and right mammograms of an individual in the normally symmetrical flow of mammary structures. The analysis was focused on the fibroglandular disc of the mammograms, segmented in a preprocessing step [280, 375]. The methods and results of Ferrari et al. are presented in the following paragraphs.

8.9.1 The fibroglandular disc

As indicated by the proceedings of the recent International Workshops on Digital Mammography [685, 686, 687, 688], several researchers are developing image processing methods to detect early breast cancer. Most of the techniques proposed perform analysis of the whole mammogram, without taking into account the fact that mammograms have different density patterns and anatomical regions that are used by radiologists in diagnostic interpretation. In fact, mammographic images are complex and difficult to analyze due to the wide variation in the density and the variable proportion of fatty and fibroglandular tissues in the breast [689]. Based upon these observations, a few researchers have proposed methods to segment and also to model mammograms in terms of anatomical regions [375, 383, 690, 691].

Miller and Astley [372] investigated the visual cues utilized by radiologists and the importance of a comparison of the corresponding anatomical structures in order to detect asymmetry between the left and right mammograms of an individual. However, in their work, the anatomical segmentation approach and the possible methodologies to segment the mammograms were not the main issue. Aylward et al. [383] devised a modeling system to segment a given mammographic image into five major components: background, uncompressed fat (periphery of the breast close to the skin-air boundary), fat, dense tissue, and muscle; the system combined geometric and statistical techniques. A few other segmentation techniques for mammograms have been presented in the literature; however, the focus has not been on anatomical segmentation but on specific problems such as density correction of peripheral breast tissue [368, 369], localization of the nipple on mammograms [370, 682], and quantification of breast density [356] and its association with the risk of breast cancer [382, 448, 689, 692, 693, 694, 695].

The fibroglandular disc is an anatomical region of the breast characterized by dense tissues, ligaments, and milk ducts. Normally, it has the shape of a disc or a cone, and goes through the interior of the breast from the region near the chest wall to the nipple [696]. Segmentation of the fibroglandular disc could form an important stage in techniques for the detection of breast cancer that use asymmetry between the left and right mammograms of the same subject, or for monitoring breast density changes in screening programs.

According to Caulkin et al. [697], it has been noticed clinically that breast cancer occurs most frequently in the upper and outer quadrant of the breast, and that the majority of cancers are associated with glandular rather than fatty tissues. A common procedure used by radiologists in screening programs for the detection of breast cancer is comparison between the left and right fibroglandular discs of the mammograms of the same subject.

Several works reported in the literature have been directed to address the problem of automatic quantification of breast density and its association with the risk of breast cancer [382, 448, 689, 692, 693, 694, 695]. Most of such works propose an index or a set of values for the quantification of breast tissue

density. However, only a few works have attempted to address the problem of detection and segmentation of the fibroglandular disc [356, 375, 383, 690, 691] for subsequent analysis.

Ferrari et al. [280] proposed a method to segment the fibroglandular disc in mammograms. In this method, prior to the detection of the fibroglandular disc, the breast boundary and the pectoral muscle are detected using the methods described in Sections 5.9 and 5.10. The fibroglandular disc is detected by defining a breast density model. The parameters of the model are estimated by using the EM algorithm [698] and the minimum-description length (MDL) principle [699]. Then, a reference value computed by using information from the pectoral muscle region is used along with the breast density model in order to identify the fibroglandular disc. The details of the methods are described in the following subsections.

8.9.2 Gaussian mixture model of breast density

The breast density model used by Ferrari et al. [280] is based upon a Gaussian mixture model [700] (see also Section 9.9.3) estimated by using the gray-level intensity distribution that represents categories or classes with different density values in mammograms. Except for the first category, the categories are related to types of tissue that may be present in the breast.

Different from the model proposed by Aylward et al. [383], which fixes at five the number of tissue classes, the model used by Ferrari et al. [280] was formulated with the hypothesis that the number of tissue classes in the effective region of the breast (after extracting the pectoral muscle) may vary from two to four among the following possibilities:

1. Uncompressed fatty tissues — represented by fatty tissues localized in the periphery of the breast.
2. Fatty tissues — composed by fatty tissues that are localized next to the uncompressed fatty tissues, and surround the denser areas of the fibroglandular disc.
3. Nonuniform density tissues — including the density region that surrounds the high-density portions of the fibroglandular disc extending close to the chest wall.
4. High-density tissues — represented by the high-density portions of the fibroglandular disc.

The hypothesis described above is based upon the fact that breast tissues may naturally vary from one person to another, or even for the same person during her life time due to aging [701] or hormone-replacement therapy [702]. A high-density and high-fat breast, for example, will likely present only two categories. It was assumed that the data (the gray-level values in a segmented

mammogram) are generated by a Gaussian mixture model with a one-to-one correspondence between the mixture model components and the (observed) tissue classes (see also Section 9.9). Thus, the marginal probability of having a gray level x is the sum of the probability over all mixture components, and it is represented by a linear superposition of multiple weighted Gaussians as

$$p(x|\bar{\Theta}) = \sum_{i=1}^K W_i p(x|\bar{\theta}_i), \quad (8.69)$$

where the x values represent the gray-level values in the image; W_i are the normalized mixing parameters ($\sum_{i=1}^K W_i = 1$ with $0 \leq W_i \leq 1$); $p(x|\bar{\theta}_i)$ is the Gaussian PDF parameterized by $\bar{\theta}_i = [\mu_i, \sigma_i]$ that is, the mean value μ_i and the standard deviation σ_i of the i^{th} Gaussian kernel; the vector $\bar{\Theta}$ represents the collection of the parameters of the mixture model ($W_1, W_2, \dots, W_K, \bar{\theta}_1, \bar{\theta}_2, \dots, \bar{\theta}_K$); and K is the number of Gaussian kernels (that is, tissue categories). The Gaussian kernel is represented as

$$p(x|\bar{\theta}_i) = \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left[-\frac{(x - \mu_i)^2}{2\sigma_i^2}\right]. \quad (8.70)$$

In the case of using features other than the gray-level values of the image, such as texture features, a multivariate Gaussian must be used instead of a univariate Gaussian. In this case, the mean value and the standard deviation of the gray-level values are replaced by the mean vector and the covariance matrix of the feature vectors, respectively. In the model as above, the Bayesian assumption is made: that the PDF associated with a pixel in the image is independent of that of the other pixels given a class of tissue, and furthermore, independent of its position in the image. The estimation of the parameters is performed by using the EM algorithm, which is an iterative procedure that maximizes the log-likelihood of the parameters of the model for a dataset representing a PDF [703]. In the EM algorithm, the estimation of the model parameters is performed in two consecutive steps: the E-step and the M-step. In the E-step, the current set of parameters is used to compute the model. The model is then assumed to be correct and the most likely distribution of the data with respect to the model is found. In the M-step, the parameters of the model are reevaluated with respect to the new data distribution by maximizing the log-likelihood, given as

$$\log L(\bar{\Theta}|\chi) = \log \prod_{i=1}^N p(x_i|\bar{\Theta}), \quad (8.71)$$

where N is the number of pixels in the effective region of the breast (which is the region demarcated by the breast boundary without the pectoral muscle), and χ represents the data sample. The procedure is iterated until the values of $\log L(\bar{\Theta}|\chi)$ between two consecutive estimation steps increase by less than 1%, or the number of iterations reaches a specified limit (200 cycles).

Initialization of the model parameters: The parameters of the model were first initialized by setting the center and weight of each Gaussian as $\mu_i = \eta$ and $W_i = 1/K$, where $i = 1, 2, \dots, K$ is the index of the Gaussian kernel, and η is a random value within the range defined by the minimum and maximum gray-level values present in the effective area of the breast. The variance σ_i^2 of each Gaussian was initialized to the nearest distance to the other Gaussian kernels. If the variance σ_i^2 became less than unity during the maximization step (the M-step), it was reinitialized with a large random value. This procedure was intended to avoid shrinkage of the variance to a small value. The EM estimation procedure was initialized and repeated three times in order to minimize the chance of convergence to a local minimum.

Model selection: Besides the initialization of the parameters, another difficulty with the mixture model lies in choosing the number of components that best suits the number of clusters (or density categories) present in the image. Several methods have been proposed in the literature to estimate the number of components in a mixture of Gaussians, among which Akaike's information criterion, Bayesian information criterion, and MDL are the most commonly used methods [704, 705]. These estimators are motivated from different theories that translate, in practice, to different penalty factors in the formulation used to select the best model. The MDL criterion is based upon an information-theoretic view of induction as data compression. It is equivalent to the Bayesian information criterion, which gives a Bayesian interpretation. Akaike's information criterion is derived from a different theoretical perspective: it is an optimal selection rule in terms of prediction error; that is, the criterion identifies a finite-dimensional model that, while approximating the data provided, has good prediction properties. However, Akaike's information criterion tends to yield models that are large and complex. In general, the MDL criterion outperforms Akaike's and Bayesian criteria.

As discussed above, the number of density categories in the breast density model can vary from two to four. Because no reliable prior information is available about the number of tissue categories present in a given mammogram, the MDL principle [699] was used to select the number K of the Gaussian kernels of the model. The MDL principle deals with a tradeoff between the maximum-likelihood (or minimum-error) criterion for fitting the model to a dataset and the complexity of the model being designed [706]. Thus, if the models designed by using two different values of K fit the data equally well, the simpler model is used. The value of K was chosen so as to maximize the quantity

$$\log L(\bar{\Theta}|\chi) - \frac{N(K)}{2} \log K, \quad (8.72)$$

where $N(K) = K(2d + 1)$ is the number of free parameters in the mixture model with K Gaussian kernels. The value of K ranges from two to four, and $d = 1$ represents the dimension of the feature space.

8.9.3 Delimitation of the fibroglandular disc

After computing the parameters of the Gaussian mixture model, the maximum-likelihood method was applied to the original image to produce a K -level image that encoded, at each pixel, a cluster membership with the highest likelihood among the K estimated Gaussian kernels. Figure 8.54 shows the effective region of the image mdb042 used for the model estimation process and the frequency distribution plots of the resulting Gaussian mixture components.


According to Karssemeijer [382], the density of the pectoral muscle is an important item of information that can be used as a reference in the interpretation of densities in the breast tissue area, because regions of similar brightness or density will most likely correspond to fibroglandular tissue. Based upon this observation and the breast density model described above, a postprocessing stage was developed in order to determine the cluster region in the K -level image, if one existed, that agreed with the fibroglandular disc. In this stage, the K -level cluster was classified as the fibroglandular region if $\mu_K \geq \mu_P - \sigma_P$, where μ_P and σ_P are, respectively, the mean and standard deviation of the gray-level values of the pectoral muscle region, and μ_K is the mean gray level of the cluster K computed from the effective region of the given image. The threshold value ($\mu_P - \sigma_P$) used to determine the fibroglandular disc was arrived at based upon experiments, because a direct comparison between the densities of the pectoral muscle and the fibroglandular tissue, in terms of physical parameters, would be difficult due to several influential factors of the image-acquisition protocol [707].

Figures 8.54 (e) and (f) illustrate, respectively, the K -level image ($K = 4$ by MDL) resulting from the mixture model designed, and the fibroglandular disc identified, for the mammogram in Figure 8.54 (a). The results for another mammogram are provided in Figure 8.55, with $K = 3$ by MDL. A simplified description of the methods is as follows:

1. Initialize the Gaussian mixture model parameters $\bar{\Theta}(\mu_i, \sigma_i^2, W_i, i = 1, 2, \dots, K)$.
2. Repeat:
 - (a) E-step: Compute the model $p(x|\bar{\Theta})$ by maximizing the log-likelihood and assuming the parameter vector Θ to be correct.
 - (b) M-step : Reevaluate Θ based upon the new data distribution computed in the previous step.

Until $\log L(\bar{\Theta}|\chi) - \frac{N(K)}{2} \log K$ increases by less than 1%.

3. Obtain the K -level image by encoding in each pixel the cluster membership with the highest likelihood.



(a)





Figure 8.54 (b)



(c)





Figure 8.54 (d)



(e)




Figure 8.54 (f)

FIGURE 8.54

(a) Original mammographic image mdb042 from the Mini-MIAS database [376]. (b) Breast contour and pectoral muscle edge detected automatically. (c) Effective region of the mammogram obtained after performing the segmentation steps. (d) Histogram of the effective area of the mammogram and the mixture of Gaussian components. (e) Four-level image resulting from the EM algorithm. (f) Fibroglandular disc obtained after thresholding. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, R.A. Borges, and A.F. Frère, "Segmentation of the fibro-glandular disc in mammograms using Gaussian mixture modelling", *Medical and Biological Engineering and Computing*, 42: 378 – 387, 2004. © IFMBE.


4. Delimit the fibroglandular disc based upon the density of the pectoral muscle.

Evaluation of the results of segmentation: In the work of Ferrari et al. [280], 84 images randomly chosen from the Mini-MIAS database [376] were used to test the segmentation of the fibroglandular disc. All images were MLO views with $200\ \mu\text{m}$ sampling interval and 8 *b* gray-level quantization. In order to reduce the processing time, all images were downsampled with a fixed sampling distance such that the original images with a matrix size of $1,024 \times 1,024$ pixels were reduced to 256×256 pixels. The results obtained with the downsampled images were mapped to the original mammograms for subsequent analysis and display.

The results obtained were evaluated in consultation with a radiologist experienced in mammography. The test images were displayed on a computer monitor. By using the Gimp program [380], the contrast, brightness, and zoom options were provided for improved visualization and assessment of the results of the segmentation procedure.

Because the delineation of the fibroglandular disc is a difficult and time-consuming task, and also because the segmentation method may provide disjoint regions, the results were evaluated by using the following subjective procedure. The segmented and original images were simultaneously presented to the radiologist on a computer monitor. The radiologist visually compared the two images and assigned one of five categories to the result of segmentation, as follows:

1. Excellent: Agreement between the segmented disc and the observed disc on the mammogram is higher than 80%.
2. Good: Agreement between the segmented disc and the observed disc on the mammogram is between 60 and 80%.



(a)





Figure 8.55 (b)



(c)




Figure 8.55 (d)

FIGURE 8.55

(a) Breast contour and pectoral muscle edge superimposed on the original image mdb008 [376]. (b) Histogram of the effective area of the mammogram and the mixture of Gaussian components. (c) Three-level image resulting from the EM algorithm. (d) Fibroglandular disc obtained after thresholding. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, R.A. Borges, and A.F. Frère, “Segmentation of the fibro-glandular disc in mammograms using Gaussian mixture modelling”, *Medical and Biological Engineering and Computing*, 42: 378 – 387, 2004. © IFMBE.

3. Average: Agreement between the segmented disc and the observed disc on the mammogram is between 40 and 60%.
4. Poor: Agreement between the segmented disc and the observed disc on the mammogram is between 20 and 40%.
5. Complete failure: Agreement between the segmented disc and the observed disc on the mammogram is less than 20%.

The overall results of segmentation of the fibroglandular disc were considered to be promising. Approximately 81% of the cases (68 images) were rated as acceptable for CAD purposes (Categories 1 and 2). In Category 1, the results for 10 images were considered by the radiologist to be underestimated, and the results for 11 images to be overestimated. In Category 2, the result for one image was considered to be underestimated, and the results for two images to be overestimated. The results for about 19% of the cases (16 images in Categories 3, 4, and 5) were considered to be unsatisfactory.

In spite of the attractive features of the EM algorithm, such as reliable convergence, well-established theoretical basis, and easy implementation, a major drawback is the problem of convergence to local maxima, which is related to the initialization of the parameters of the model. A more efficient minimization technique, such as the deterministic annealing EM algorithm (DAEM) [708], may give better performance. The DAEM algorithm uses the principle of maximum entropy and analogy with statistical mechanics to reformulate the maximization step of the EM algorithm. The log-likelihood function is replaced by a posterior PDF that is parameterized by β , with $1/\beta$ corresponding to the temperature in analogy to the annealing process. This formulation of the PDF provides robustness to DAEM and minimizes the possibility of convergence to local maxima.

The application of the DAEM algorithm to the images in the work of Ferrari et al. [280] provided limited improvement in terms of the fit between the sums of the Gaussians in the mixture models and the corresponding true

histograms. However, the change in terms of the final result of segmentation of the fibroglandular disc with respect to the result of the EM algorithm was not significant.

The stochastic EM algorithm [709] is a variant of EM with the aim to avoid the dependence of the final results upon the initialization. In the stochastic EM algorithm, after the E-step, a stochastic classification step (the S-step) is added to obtain partitions of the data according to posterior distributions. In the absence of any other information, random initialization is applied. The M-step uses the partitions to compute new estimates of the mean vector and the covariance matrix. The stochastic EM algorithm presents the following improvements in comparison with the standard EM algorithm: it is adequate to know an upper bound on the number of classes; the solution is essentially independent of the initialization; and the speed of convergence is appreciably improved.

Other features such as texture may also be used as additional information in order to improve the results. Because the segmentation method is based upon the classification of each individual pixel in the image, the resulting images are often not compact, and a follow-up procedure may be necessary in order to demarcate the convex hull of the fibroglandular disc.

8.9.4 Motivation for directional analysis of mammograms

Most of the concepts used in image processing and computer vision for oriented pattern analysis have their roots in neurophysiological studies of the mammalian visual system. Campbell and Robson [710] suggested that the human visual system decomposes retinal images into a number of filtered images, each of which contains intensity variations over a narrow range of frequency and orientation. Marcelja [711] and Jones and Palmer [390] demonstrated that simple cells in the primary visual cortex have receptive fields that are restricted to small regions of space and are highly structured, and that their behavior corresponds to local measurements of frequency.

According to Daugman [389, 712], a suitable model for the 2D receptive field profiles measured experimentally in mammalian cortical simple cells is the parameterized family of 2D Gabor functions (see Section 8.4). Another important characteristic of Gabor functions or filters is their optimal joint resolution in both space and frequency, which suggests that Gabor filters are appropriate operators for tasks requiring simultaneous measurement in the two domains [495]. Except for the optimal joint resolution possessed by the Gabor functions, the DoG and difference of offset Gaussian filters used by Malik and Perona [713] have similar properties.

Gabor filters have been presented in several works on image processing [495, 500, 714]; however, most of these works are related to the segmentation and analysis of texture. Rolston and Rangayyan [542, 543] proposed methods for directional decomposition and analysis of linear components in images using multiresolution Gabor filters; see Section 8.4.

Multiresolution analysis by using Gabor filters has natural and desirable properties for the analysis of directional information in images; most of these properties are based upon biological vision studies as described above. Other multiresolution techniques have also been used with success in addressing related topics, such as texture analysis and segmentation, and image enhancement. Chang and Kuo [715], for instance, developed a new method for texture classification that uses a tree-structured wavelet transform for decomposing an image. Image decomposition was performed by taking into account the energy of each subimage instead of decomposing subsignals in the low-frequency channels.

Laine and Fan [488] presented a new method for computing features for texture segmentation based upon a multiscale representation. They used a discrete wavelet packet frame to decompose an image at different levels of resolution. The levels were analyzed by using two algorithms for envelope detection based upon the Hilbert transform and zero crossings. Laine and Fan stated that the selection of the filters for feature extraction should take into account three important features: symmetry, frequency response, and boundary accuracy. However, it should be noted that the Gabor function is the only function that can achieve optimal joint localization (tradeoff between boundary accuracy and good frequency response).

Li et al. [716] performed analysis and comparison of the directional selectivity of wavelet transforms by using three different types of frequency decomposition: rectangular, hexagonal, and radial-angular decomposition. The methods were applied to detect spiculated breast masses and lung nodules. According to Li et al., the best results were achieved by using radial-angular decomposition with both mammographic and chest radiographic images. Qian et al. [623] presented three image processing modules for mass detection in digital mammography; a tree-structured filter module for noise suppression, a directional wavelet transform (DWT) technique for the decomposition of mammographic directional features, and a tree-structured wavelet transform for image enhancement. By making the parameters of the three methods adaptive, Qian et al. improved the results obtained in previous related works [716, 717, 718]. In the DWT method, which is related to the technique proposed by Ferrari et al. [381], the number of filter orientations was adaptively selected in the range 4 – 32, corresponding to angular bandwidth in the range $45^\circ - 5.63^\circ$. The adaptive DWT module provided the best results, with an overall classification efficiency of 0.91.

Chang and Laine [624] proposed a method designed with the goal of enhancing mammograms based upon information regarding orientation. They used a set of separable steerable filters to capture subtle features at different scales and orientations spanning an over-complete multiscale representation computed via a fast wavelet transform. By using the filtered output images, they generated a coherence image and phase information that were used by a nonlinear operator applied to the wavelet coefficients to enhance the directional information in digital mammograms. Finally, the enhanced image was

obtained from the modified wavelet coefficients via an inverse fast wavelet transform. (See Section 7.7 for further discussion on related topics.)

Inspired by studies on the Gabor function and the related topics mentioned above, Ferrari et al. [381] proposed a scheme based upon a bank of self-similar Gabor functions and the KLT to analyze directional components of images. The method was applied to detect global signs of asymmetry in the fibroglandular discs of the left and right mammograms of a given subject. The related methods are described in the following paragraphs.

8.9.5 Directional analysis of fibroglandular tissue

Ferrari et al. [381] used the formulation of 2D Gabor functions as a Gaussian modulated by a complex sinusoid, specified by the frequency of the sinusoid W and the standard deviations σ_x and σ_y of the Gaussian envelope as

$$\psi(x, y) = \frac{1}{2\pi\sigma_x\sigma_y} \exp \left[-\frac{1}{2} \left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} \right) + j 2\pi Wx \right]. \quad (8.73)$$

Despite this general form, there is no standard and precise definition of a 2D Gabor function, with several variations appearing in the literature [500, 714, 496]; see Sections 8.4 and 8.10.3. Most of these variations are related to the use of different measures of width for the Gaussian envelope and the frequency of the sinusoid [719]. Based upon neurophysiological studies [711, 390] and wavelet theory [386, 720], the Gabor function, normalized in an appropriate way, can be used as a mother wavelet to generate a family of nonorthogonal Gabor wavelets [385]. However, as pointed out by Jain and Farrokhnia [500], although the Gabor function can be an admissible wavelet, by removing the DC response of the function, it does not result in an orthogonal decomposition, which means that a wavelet transform based upon the Gabor wavelet includes redundant information. A formal mathematical derivation of 2D Gabor wavelets, along with the computation of the frame bounds for which this family of wavelets forms a tight frame, is provided by Lee [385]. Despite the lack of orthogonality presented by the Gabor wavelets, the Gabor function is the only function that can achieve the theoretical limit for joint resolution of information in both the space and the frequency domains.

Ferrari et al. [381] used the phrase “Gabor wavelet representation” to represent a bank of Gabor filters normalized to have DC responses equal to zero and designed in order to have low redundancy in the representation. The Gabor wavelet representation used was as proposed by Manjunath and Ma [384]. The Gabor wavelets were obtained by dilation and rotation of $\psi(x, y)$ as in Equation 8.73 by using the generating function

$$\psi_{m,n}(x, y) = a^{-m} \psi(x', y'), \quad a > 1, \quad m, n = \text{integers}, \quad (8.74)$$

$$x' = a^{-m} [(x - x_0) \cos \theta + (y - y_0) \sin \theta],$$

$$y' = a^{-m} [-(x - x_0) \sin \theta + (y - y_0) \cos \theta],$$

where (x_0, y_0) is the center of the filter in the spatial domain, $\theta = \frac{n\pi}{K}$, K is the total number of orientations desired, and m and n indicate the scale and orientation, respectively. The scale factor a^{-m} in Equation 8.74 is meant to ensure that the energy is independent of m . Examples of the Gabor wavelets used in the work of Ferrari et al. [381] are shown in Figure 8.56.

Equation 8.73 can be written in the frequency domain as


$$\Psi(u, v) = \frac{1}{2\pi\sigma_u\sigma_v} \exp \left\{ -\frac{1}{2} \left[\frac{(u - W)^2}{\sigma_u^2} + \frac{v^2}{\sigma_v^2} \right] \right\}, \quad (8.75)$$

where $\sigma_u = \frac{1}{2\pi\sigma_x}$ and $\sigma_v = \frac{1}{2\pi\sigma_y}$. The design strategy used is to project the filters so as to ensure that the half-peak magnitude supports of the filter responses in the frequency spectrum touch one another, as shown in Figure 8.57. In this manner, it can be ensured that the filters will capture most of the information with minimal redundancy.

Other related methods proposed in the literature use either complex-valued Gabor filters [496] or pairs of Gabor filters with quadrature-phase relationship [495]. In the formulation of Ferrari et al. [381], the Gabor wavelet representation uses only real-valued, even-symmetric filters oriented over a range of 180° only, as opposed to the full 360° range commonly described in the literature. Because the Gabor filters are used to extract meaningful features from real images (and hence with Hermitian frequency response [387]), the response to the even-symmetric filter components will remain unchanged for filters oriented 180° out of phase and the odd-symmetric component will be negated. Thus, based upon this fact, and also on psychophysical grounds provided by Malik and Perona [713], Manjunath and Ma [384] ignored one half of the orientations in their Gabor representation, as illustrated in Figure 8.57. In order to ensure that the bank of Gabor filters designed as above becomes a family of admissible 2D Gabor wavelets [385], the filters $\psi(x, y)$ must satisfy the admissibility condition of finite energy [386], which implies that their Fourier transforms are pure bandpass functions having zero response at DC. This condition was achieved by setting the DC gain of each filter as $\Psi(0, 0) = 0$, which ensures that the filters do not respond to regions with constant intensity.

The approach described above results in the following formulas for computing the filter parameters σ_u and σ_v :

$$a = \left(\frac{U_h}{U_l} \right)^{\frac{1}{S-1}}, \quad (8.76)$$

**FIGURE 8.56**

Examples of Gabor wavelets in the space domain, with four orientations ($\theta = 0^\circ, 45^\circ, 90^\circ,$ and 135°) and four scales ($\sigma_x = 11, 5, 2, 1,$ and $\sigma_y = 32, 16, 7, 4$ pixels). The size of each wavelet image shown is 121×121 pixels. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, and A.F. Frère, "Analysis of asymmetry in mammograms via directional filtering with Gabor wavelets", *IEEE Transactions on Medical Imaging*, 20(9): 953 – 964, 2001. © IEEE.





Figure 8.57 (a)



(b)

FIGURE 8.57

Examples of Gabor filters in the frequency domain. Each ellipse represents the range of the corresponding filter response from 0.5 to 1.0 in squared magnitude. The plots (a) and (b) illustrate two ways of dividing the frequency spectrum by changing the U_l , U_h , S , and K parameters of the Gabor representation. Plot (a) represents the filter bank used in the work of Ferrari et al. [381] for the analysis of mammograms. The redundancy in the representation is minimized by ensuring that the half-peak magnitude supports of the filter responses touch one another. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, and A.F. Frère, "Analysis of asymmetry in mammograms via directional filtering with Gabor wavelets", *IEEE Transactions on Medical Imaging*, 20(9): 953 – 964, 2001. © IEEE.

$$\sigma_u = \frac{(a-1)U_h}{(a+1)\sqrt{2\ln 2}}, \quad (8.77)$$

$$\sigma_v = \frac{\tan\left(\frac{\pi}{2K}\right) \left[U_h - \left(\frac{\sigma_u^2}{U_h}\right) 2\ln 2 \right]}{\sqrt{2\ln 2 - \frac{(2\ln 2)^2 \sigma_u^2}{U_h^2}}}, \quad (8.78)$$

where U_l and U_h denote the lower and upper center frequencies of interest. The K and S parameters are, respectively, the number of orientations and the number of scales in the desired multiresolution decomposition procedure. The frequency of the sinusoid W is set equal to U_h , and $m = 0, 1, \dots, S-1$.

Because of the lack of orthogonality of the Gabor wavelets, the computation of the expansion coefficients becomes difficult. This task, however, is trivial when using a set of orthogonal functions, because the expansion coefficients, given by

$$c_{m,n} = \langle f(x, y), \psi_{m,n}(x, y) \rangle = \int_x \int_y f(x, y) \psi_{m,n}(x, y) dx dy, \quad (8.79)$$

are the projections of the image $f(x, y)$ onto the same set of functions, where \langle, \rangle denotes the inner product. In this case, the analysis and synthesis windows are the same, and the original image can be reconstructed as

$$f(x, y) = \sum_m \sum_n \langle f(x, y), \psi_{m,n}(x, y) \rangle \psi_{m,n}(x, y). \quad (8.80)$$

However, the joint localization and orthogonality of the set of functions are properties that cannot be met simultaneously. Much work has been done to overcome the problem of the lack of orthogonality of Gabor wavelets, with most of them using dual frames or biorthogonal functions [720], iterative methods [712], or adjustment of the phase-space sampling in order to obtain a reasonably tight frame [385]. In the dual-frame approach, for instance, the set of projection coefficients $c_{m,n} = \langle f(x, y), \tilde{\psi}_{m,n}(x, y) \rangle$ of the dual frame can be obtained by minimizing the cost function

$$\kappa = \left| f(x, y) - \sum_m \sum_n c_{m,n} \psi_{m,n}(x, y) \right|^2, \quad (8.81)$$

where $\tilde{\psi}_{m,n}$ is the dual frame.

In directional filtering and analysis, the interest lies in image analysis without the requirement of exact reconstruction (synthesis) of the image. Therefore, instead of using the wavelet coefficients, Ferrari et al. [381] used the magnitude of the filter response, computed as

$$a_{m,n} = |f(x, y) * \psi_{m,n}^{\text{even}}(x, y)|, \quad (8.82)$$

where $\psi_{m,n}^{\text{even}}(x, y)$ indicates only the even-symmetric part of the complex Gabor filter, and $*$ represents 2D convolution.

By adjusting the parameters U_l and U_h in the Gabor representation of Manjunath and Ma [384], the range of the frequency spectrum to be used for multiresolution analysis may be selected. The area of each ellipse indicated in Figure 8.57 represents the spectrum of frequencies covered by the corresponding Gabor filter. Once the range of the frequency spectrum is adjusted, the choice of the number of scales and orientation may be made in order to cover the range of the spectrum as required. The choice of the number of scales (S) and orientations (K) used in the work of Ferrari et al. [381] for processing mammographic images was based upon the resolution required for detecting oriented information with high selectivity [388, 389]. The frequency bandwidths of the simple and complex cells in the mammalian visual system have been found to range from 0.5 to 2.5 octaves, clustering around 1.2 octaves and 1.5 octaves, and their angular bandwidth is expected to be smaller than 30° [390, 389]. By selecting $U_l = 0.05$, $U_h = 0.45$, $S = 4$, and $K = 12$ for processing mammographic images, Ferrari et al. [381] indirectly set the Gabor representation to have a frequency bandwidth of approximately one octave and an angular bandwidth of 15° . The effects of changing the U_l , U_h , S , and K parameters of the Gabor representation as above on frequency localization are shown in Figure 8.57.

The directional analysis method proposed by Ferrari et al. [381] starts by computing the Gabor wavelets using four scales ($S = 4$) and twelve directions ($K = 12$), with $U_l = 0.05$ and $U_h = 0.45$ cycles/pixel. The parameters U_l and U_h were chosen according to the scales of the details of interest in the mammographic images. Differing from other Gabor representations [500, 714, 385], it should be noted that the parameters U_l , U_h , S , and K in the representation described above have to be adjusted in a coordinated manner, by taking into account the desirable frequency and orientation bandwidths. (In the Gabor representation of Manjunath and Ma [384], the filters are designed so as to represent an image with minimal redundancy; the ellipses as in Figure 8.57 touch one another.)

The filter outputs for each orientation and the four scales were analyzed by using the KLT (see Section 11.6.2). The KLT was used to select the principal components of the filter outputs, preserving only the most relevant directional elements present at all of the scales considered. Results were then combined as illustrated in Figure 8.58, in order to allow the formation of an S -dimensional vector (\bar{x}) for each pixel from each set of the corresponding pixels in the filtered images ($S = 4$).

The vectors corresponding to each position in the filter responses were used to compute the mean vector μ and the covariance matrix σ . The eigenvalues and eigenvectors of the covariance matrix were then computed and arranged in descending order in a matrix \mathbf{A} such that the first row of \mathbf{A} was the eigenvector corresponding to the largest eigenvalue, and the last row was the eigenvector corresponding to the smallest eigenvalue. The first N principal components corresponding to 95% of the total variance were then selected, and used to represent the oriented components at each specific orientation. The principal

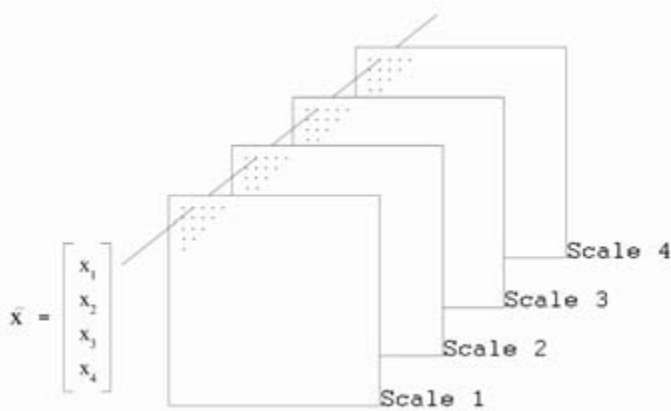


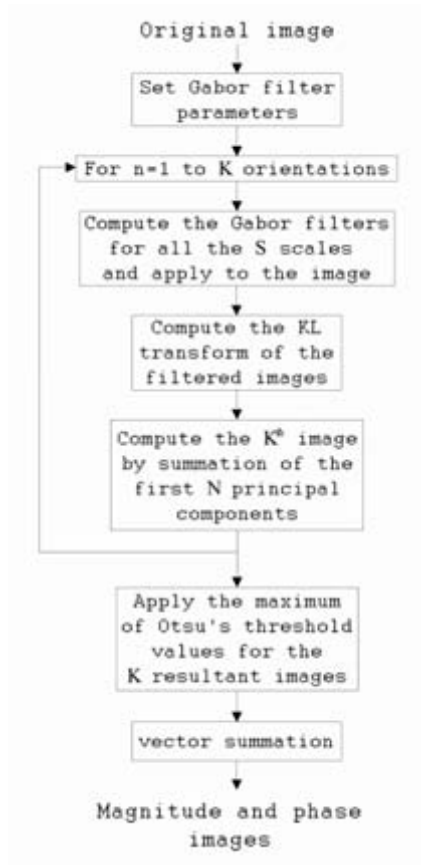
FIGURE 8.58

Formation of the vector $\mathbf{x} = \bar{\mathbf{x}}$ from the corresponding pixels of the same orientation and four scales. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, and A.F. Frère, “Analysis of asymmetry in mammograms via directional filtering with Gabor wavelets”, *IEEE Transactions on Medical Imaging*, 20(9): 953 – 964, 2001. © IEEE.

components were computed as $\mathbf{y} = \mathbf{A}(\mathbf{x} - \boldsymbol{\mu})$. Analysis of variance was performed by evaluating the eigenvalues of the matrix \mathbf{A} . The KLT method is optimal in the sense that it minimizes the MSE between the vectors \mathbf{x} and their resulting approximations \mathbf{y} . The result of application of the KLT to all orientations, as described above, is a set of K images, where K is the number of orientations.

Because Gabor wavelets are nonorthogonal functions, they do not have a perfect reconstruction condition. This fact results in a small amount of out-of-band energy interfering with the reconstruction, which is translated into artifacts in the reconstructed image. Although reconstruction of the original image from the filtered images is not required in directional analysis, the effects of spectral leakage need to be removed. For this purpose, the images resulting from the KLT were thresholded by using the maximum of Otsu’s threshold value [591] (see Section 8.3.2) computed for the K images.

Phase and magnitude images, indicating the local orientation and intensity, were composed by vector summation of the K filtered images [387], as illustrated in Figure 8.10. Rose diagrams were composed from the phase and magnitude images to represent the directional distribution of the fibroglandular tissue in each mammogram. The complete algorithm for directional analysis based upon Gabor filtering is summarized in Figure 8.59.

**FIGURE 8.59**

Block diagram of the procedure for directional analysis using Gabor wavelets. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, and A.F. Frère, "Analysis of asymmetry in mammograms via directional filtering with Gabor wavelets", *IEEE Transactions on Medical Imaging*, 20(9): 953 – 964, 2001. © IEEE.

8.9.6 Characterization of bilateral asymmetry


Figures 8.60 and 8.61 show two pairs of images from the Mini-MIAS database [376]. All of the images in this database are $1,024 \times 1,024$ pixels in size, with $200 \mu\text{m}$ sampling interval and 8 *b* gray-level quantization. Figures 8.60 (a) and (b) are, respectively, the MIAS images mdb043 and mdb044 representing the right and left MLO mammograms of a woman, classified as a normal case. Figures 8.61 (a) and (b) show the MIAS images mdb119 and mdb120, classified as a case of architectural distortion.

Only the fibroglandular disc of each mammogram was used to compute the directional components, due to the fact that most of the directional components such as connective tissues and ligaments exist in this specific region of the breast. The fibroglandular disc ROIs of the mammograms selected for illustration are shown in Figures 8.60 (c), 8.60 (d), 8.61 (c), and 8.61 (d).

Figure 8.62 shows the principal components obtained by applying the KLT to the Gabor filter outputs at the orientation of 135° and four scales for the ROI in Figure 8.61 (d). It can be seen that the relevant information is concentrated in the first two principal components. This is evident based upon an evaluation of the eigenvalues, listed in the caption of Figure 8.62. In this example, only the first two principal components were used to represent the oriented components in the 135° orientation, because their eigenvalues add to 99.34% ($> 95\%$) of the total variance. After thresholding the filtered images with Otsu's method in order to eliminate the effects of spectral leakage, the magnitude and phase images were composed by vector summation, as illustrated in Figures 8.63 and 8.64.

Figures 8.63 (a) – (d) show the result of Gabor filtering for the normal case in Figure 8.60. The rose diagrams in Figures 8.63 (c) and (d) show the distribution of the tissues in the fibroglandular discs of both the left and right views. An inspection of the rose diagrams shows that the results obtained are in good agreement with visual analysis of the filtered results in Figures 8.63 (a) and (b), and the corresponding ROIs in Figures 8.60 (c) and (d). The most relevant angular information indicated in the rose diagrams are similar.

The results of the filtering process for the case of architectural distortion (see Figure 8.61), along with the respective rose diagrams, are shown in Figure 8.64. By analyzing the results of filtering, we can notice a modification of the tissue pattern caused by the presence of a high-density region. An important characteristic of the Gabor filters may be seen in the result: the filters do not respond to regions with nearly uniform intensity, that is, to regions without directional information [see Figures 8.61 (c) and 8.64 (a)]. This is an important outcome that could be used to detect asymmetric dense regions and local foci of architecture distortion. The global distortion of the tissue flow pattern is readily seen by comparison of the rose diagrams of the left and right breasts.

**FIGURE 8.60**

Images mdb043 and mdb044 of a normal case [376]. (a) and (b) Original images ($1,024 \times 1,024$ pixels at $200 \mu\text{m}/\text{pixel}$). The breast boundary (white) and pectoral muscle edge (black) detected are also shown. (c) and (d) Fibroglandular discs segmented and enlarged (512×512 pixels). Histogram equalization was applied to enhance the global contrast of each ROI for display purposes only. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, and A.F. Frère, "Analysis of asymmetry in mammograms via directional filtering with Gabor wavelets", *IEEE Transactions on Medical Imaging*, 20(9): 953 – 964, 2001. © IEEE.





FIGURE 8.61

Images mdb119 and mdb120 of a case of architectural distortion [376]. (a) and (b) Original images ($1,024 \times 1,024$ pixels at $200 \mu\text{m}/\text{pixel}$). The breast boundary (white) and pectoral muscle edge (black) detected are also shown. (c) and (d) Fibroglandular discs segmented and enlarged (512×512 pixels). Histogram equalization was applied to enhance the global contrast of each ROI for display purposes only. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, and A.F. Frère, “Analysis of asymmetry in mammograms via directional filtering with Gabor wavelets”, *IEEE Transactions on Medical Imaging*, 20(9): 953 – 964, 2001. © IEEE.

**FIGURE 8.62**

The images (a), (b), (c), and (d) are, respectively, the first, second, third, and fourth components resulting from the KLT applied to the Gabor filter responses with orientation 135° to the ROI of the image mdb120 shown in Figure 8.61 (d). The eigenvalues of the four components above are: $\lambda_1 = 10.80$, $\lambda_2 = 0.89$, $\lambda_3 = 0.09$, and $\lambda_4 = 0.01$. The images were full brightness-contrast corrected for improved visualization. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, and A.F. Frère, “Analysis of asymmetry in mammograms via directional filtering with Gabor wavelets”, *IEEE Transactions on Medical Imaging*, 20(9): 953 – 964, 2001. © IEEE.

The rose diagrams in Figures 8.63 and 8.64 present a strong visual association with the directional components of the corresponding images, and may be used by radiologists as an aid in the interpretation of mammograms.

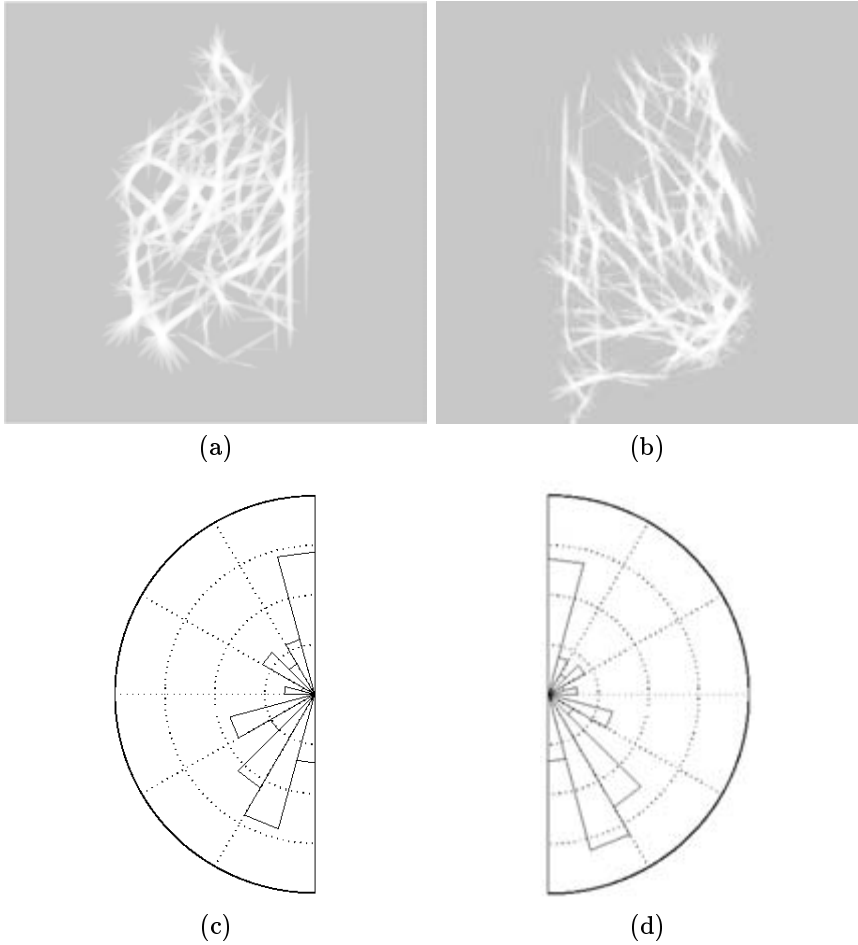
Feature extraction and pattern classification: In order to characterize bilateral asymmetry in an objective manner, three features were derived: the entropy H (Equation 8.10), the first moment M_1 (Equation 8.6), and the second central moment or variance M_2 (Equation 8.7) of the rose diagram given by the difference between the rose diagrams computed for the left and right mammograms of the same individual.

Classification of the normal and asymmetric cases was conducted by using the Bayesian linear classifier [721] (see Section 12.4.2). The Gaussian distribution was assumed in order to model the PDF, and the parameters of the model were estimated by using the training samples. The prior probabilities of the normal and asymmetry classes were assumed to be equal, and the covariance matrix was calculated in a pooled manner by averaging the covariance matrices of the normal and asymmetric classes. The leave-one-out methodology [721] was used to estimate the classification accuracy.

The directional analysis scheme was applied to 80 images (20 normal cases, 14 cases of asymmetry, and six cases of architectural distortion) from the Mini-MIAS database [376]. An exhaustive combination approach was used to select the best set of features. The selection was conducted based upon the classification results obtained by using the leave-one-out method. The best result, by using only one feature in the classification process, was achieved by the first-order angular moment (M_1), with the sensitivity, specificity, and average accuracy values equal to 77.3%, 71.4%, and 74.4%, respectively. When using two features, the best result was achieved with the combination of the first-order angular moment (M_1) and the entropy (H) features, indicating that 80% of the asymmetric and distortion cases, and 65% of the normal cases were correctly classified. The average rate of correct classification in this case was 72.5%. The low rate of specificity may be explained by the fact that even normal cases present natural signs of mild asymmetry; the mammographic imaging procedure may also distort the left and right breasts in different ways.

In a subsequent study, Rangayyan et al. [681] revised the directional analysis procedures as shown in Figure 8.65. The rose diagrams of the left and right mammograms were aligned such that their mean angles corresponded to the straight line perpendicular to the pectoral muscle, and then subtracted to obtain the difference rose diagram. In addition to the features H , M_1 , and M_2 of the difference rose diagram as described above, the dominant orientation θ_R and circular variance s_θ^2 were computed as follows:

$$X_R = \sum_{i=1}^N R_i \cos \theta_i, \quad (8.83)$$

**FIGURE 8.63**

Results obtained for the normal case in Figure 8.60. (a) and (b) Magnitude images. (c) and (d) Rose diagrams. The magnitude images were histogram-equalized for improved visualization. The rose diagrams have been configured to match the mammograms in orientation. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, and A.F. Frère, "Analysis of asymmetry in mammograms via directional filtering with Gabor wavelets", *IEEE Transactions on Medical Imaging*, 20(9): 953 – 964, 2001. © IEEE.




FIGURE 8.64

Results obtained for the case of architectural distortion in Figure 8.61. (a) and (b) Magnitude images. (c) and (d) Rose diagrams. The magnitude images were histogram-equalized for improved visualization. The rose diagrams have been configured to match the mammograms in orientation. Reproduced with permission from R.J. Ferrari, R.M. Rangayyan, J.E.L. Desautels, and A.F. Frère, "Analysis of asymmetry in mammograms via directional filtering with Gabor wavelets", *IEEE Transactions on Medical Imaging*, 20(9): 953 – 964, 2001. © IEEE.

$$Y_R = \sum_{i=1}^N R_i \sin \theta_i, \quad (8.84)$$

$$\theta_R = \text{atan} \left(\frac{Y_R}{X_R} \right), \quad (8.85)$$

and

$$s_\theta^2 = 1 - \sqrt{X_R^2 + Y_R^2}, \quad (8.86)$$

where R_i is the normalized value and θ_i is the central angle of the i^{th} angle band of the difference rose diagram, and N is the number of bins in the rose diagram.

In addition, a set of 11 features including seven of Hu's moments (see Section 6.2.2 and Equation 8.3) and the area, average density, eccentricity η , and stretch ρ were computed to characterize the shape of the segmented fibroglandular discs. Eccentricity was computed as

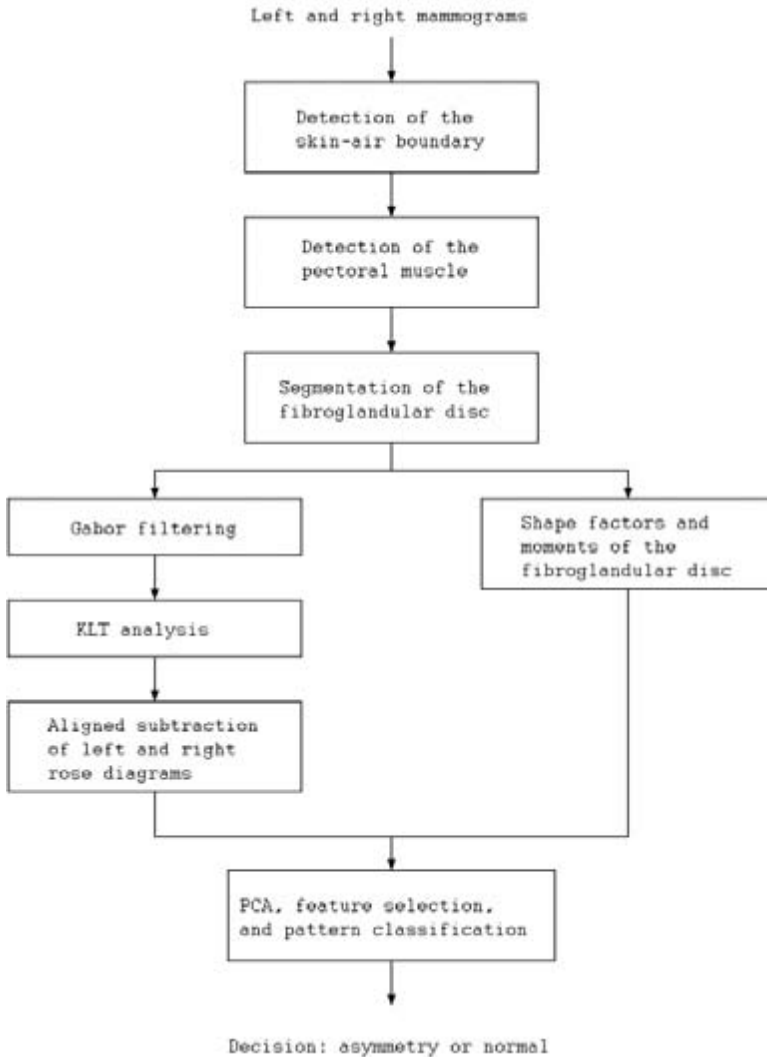
$$\eta = \frac{(m_{20} - m_{02})^2 + 4m_{11}^2}{(m_{20} + m_{02})^2}, \quad (8.87)$$

where m_{pq} are the geometric invariant moments as described in Section 6.2.2. The stretch parameter was computed as

$$\rho = \frac{x_{\max} - x_{\min}}{y_{\max} - y_{\min}}, \quad (8.88)$$

where x_{\max} , x_{\min} , y_{\max} , and y_{\min} are the corner coordinates of the rectangle delimiting the fibroglandular disc. Feature selection was performed by PCA and exhaustive combination techniques. With PCA, only the components associated with 98% of the total variance were used in the classification step. Classification was performed using linear and quadratic Bayesian classifiers with the leave-one-out method.

The revised directional analysis scheme was applied to 88 images (22 normal cases, 14 cases of asymmetry, and eight cases of architectural distortion) from the Mini-MIAS database [376]. The best overall classification accuracy of 84.4% (with a sensitivity of 82.6% and specificity of 86.4%) was obtained using the four features θ_R , M_1 , M_2 , and H computed from the aligned-difference rose diagrams using the quadratic classifier. The morphometric measures and moments, after PCA-based feature selection, resulted in an overall classification accuracy of only 71.1% with the linear classifier. The combination of all of the directional statistics, morphometric measures, and moments, after PCA-based feature selection, resulted in an overall classification accuracy of 82.2%, with a sensitivity of 78.3% and specificity of 86.4% with the linear classifier. The results indicate the importance of directional analysis of the fibroglandular tissue in the detection of bilateral asymmetry.

**FIGURE 8.65**

Block diagram of the revised procedure for the analysis of bilateral asymmetry [681].

8.10 Application: Architectural Distortion in Mammograms

Mammography is the best available tool for detecting early breast cancer; screening programs have been shown to reduce mortality rates by 30% to 70% [55] (Chapter 19), [722]. However, the sensitivity of screening mammography is affected by image quality and the radiologist's level of expertise. Bird et al. [723] estimated the sensitivity of screening mammography to be between 85% and 90%; they also observed that misinterpretation of breast cancer signs accounts for 52% of the errors, and overlooking signs is responsible for 43% of the missed lesions. The extent of errors due to overlooking of lesions reinforces the need for CAD tools in mammography.

CAD techniques and systems have been proposed to enhance the sensitivity of the detection of breast cancer: although these techniques are effective in detecting masses and calcifications, they have been found to fail in the detection of architectural distortion with adequate levels of accuracy [724]. Therefore, new CAD systems are desirable for targeted detection of subtle mammographic abnormalities, such as architectural distortion, which are the most frequent source of screening errors and false-negative findings related to cases of interval cancer [660].

Architectural distortion is defined in BI-RADSTM [403] as follows: "The normal architecture (of the breast) is distorted with no definite mass visible. This includes spiculations radiating from a point and focal retraction or distortion at the edge of the parenchyma.". Focal retraction is considered to be easier to perceive than spiculated distortion within the breast parenchyma [54] (p61). Architectural distortion could be categorized as malignant or benign, the former including cancer, and the latter including scar and soft-tissue damage due to trauma.

According to van Dijck et al. [725], "in nearly half of the screen-detected cancers, minimal signs appeared to be present on the previous screening mammogram two years before the diagnosis". Burrell et al. [660], in a study of screening interval breast cancers, showed that architectural distortion is the most commonly missed abnormality in false-negative cases. Sickles [726] reported that indirect signs of malignancy (such as architectural distortion, bilateral asymmetry, single dilated duct, and developing densities) account for almost 20% of the detected cancers. Broeders et al. [727] suggested that improvement in the detection of architectural distortion could lead to an effective improvement in the prognosis of breast cancer patients.

8.10.1 Detection of spiculated lesions and distortion

The breast contains several piecewise linear structures, such as ligaments, ducts, and blood vessels, that cause oriented texture in mammograms. The

presence of architectural distortion is expected to change the normal oriented texture of the breast.

Figure 8.66 (a) shows a mammogram of a normal breast, where the normal oriented texture is observed. The mammogram in Figure 8.66 (b) exhibits architectural distortion; Figure 8.67 shows an enlarged view of the site of architectural distortion. Observe the appearance of lines radiating from a central point in the ROI. It is also noticeable that the perceived increased density close to the center of the ROI is due to the fibroglandular disk, as it can be observed in the full mammogram view in Figure 8.66 (b).




FIGURE 8.66

(a) Mammogram showing a normal breast; image mdb243 from the Mini-MIAS database [376]. Width of image = 650 pixels = 130 mm. (b) Architectural distortion present in a mammogram from the Mini-MIAS database (mdb115). Width of image = 650 pixels = 130 mm. The square box overlaid on the figure represents the ROI including the site of architectural distortion, shown enlarged in Figure 8.67. Reproduced with permission from F.J. Ayres and R.M. Rangayyan, “Characterization of architectural distortion in mammograms via analysis of oriented texture”, *IEEE Engineering in Medicine and Biology Magazine*, January 2005. © IEEE.




FIGURE 8.67

Detail of mammogram mdb115 showing the site of architectural distortion marked by the box in Figure 8.66 (b). Width of image = 300 pixels = 60 mm. Reproduced with permission from F.J. Ayres and R.M. Rangayyan, “Characterization of architectural distortion in mammograms via analysis of oriented texture”, *IEEE Engineering in Medicine and Biology Magazine*, January 2005. © IEEE.

Several researchers have directed attention to the detection of stellate and spiculated patterns associated with masses in mammograms; see Mudigonda et al. [275] and Section 8.8 for a review on this subject. However, most of such methods have been directed toward and applied to the detection of tumors with spiculations. A common strategy in most of the methods reported in the literature for this application is to assume a stellate appearance for the lesion; detection is performed by fitting a star pattern to the texture of the breast parenchyma.

Qian et al. [623] proposed a directional wavelet transform to detect spicules radiating from masses. Kegelmeyer et al. [728] used local edge orientation histograms and Laws’ texture features to detect stellate lesions. A sensitivity of 100% was obtained, with a specificity of 82%. Karssemeijer and te Brake [641] applied operators sensitive to radial patterns of straight lines to an orientation field to detect stellate patterns, and obtained a sensitivity of 90% with one false positive per image; however, only lesions with radiating spicules (regardless of the presence of a central density) are detected in their method. Mudigonda et al. [275] investigated the detection of masses using density slicing and texture flow-field; a sensitivity of 85% with 2.45 false positives per image was achieved (see Section 8.8).

Zwiggelaar et al. [652] proposed a technique to detect abnormal patterns of linear structures, by detecting the radiating spicules and/or the central mass

expected to occur with spiculated lesions. PCA was applied to a training set of mammograms including normal tissue patterns and spiculated lesions. The results of PCA were used to construct a basis set of oriented texture patterns, which was used to analyze radiating structures. A sensitivity of 80% was obtained in the classification of pixels as belonging to normal tissue or lesions, but with low specificity. Sampat and Bovik [729] employed filtering in the Radon domain to enhance mammograms, followed by the usage of radial spiculation filters to detect spiculated lesions; however, statistical evaluation of the performance of the technique was not conducted.

Mudigonda and Rangayyan [730] proposed the use of texture flow-field to detect architectural distortion, based upon the local coherence of texture orientation; only preliminary results were given, indicating the potential of the technique in the detection of architectural distortion. Ferrari et al. [381] used Gabor filters to detect oriented patterns, with subsequent analysis designed to characterize global bilateral asymmetry; however, the method was not aimed at detecting focal architectural distortion (see Section 8.9). Matsubara et al. [731] used mathematical morphology to detect architectural distortion around the skin line, and a concentration index to detect architectural distortion within the mammary gland; they reported sensitivity rates of 94% with 2.3 false positives per image and 84% with 2.4 false positives per image, respectively, in the two tasks.

Burhenne et al. [732] studied the performance of a commercial CAD system in the detection of masses and calcifications in screening mammography, obtaining a sensitivity of 75% in the detection of masses and architectural distortion. Evans et al. [733] investigated the ability of a commercial CAD system to mark invasive lobular carcinoma of the breast: the system identified correctly 17 of 20 cases of architectural distortion. Birdwell et al. [734] evaluated the performance of a commercial CAD system in marking cancers that were overlooked by radiologists: the software detected five out of six cases of architectural distortion. However, Baker et al. [724] found the sensitivity of two commercial CAD systems to be poor in detecting architectural distortion: fewer than 50% of the cases of architectural distortion were detected. These findings indicate the need for further research in this area, and the development of algorithms designed specifically to characterize architectural distortion.

Ayres and Rangayyan [595, 679, 680] proposed the application of Gabor filters and phase portraits to characterize architectural distortion in ROIs selected from mammograms, as well as to detect sites of architectural distortion in full mammograms; their methods and results are described in the following subsections.

8.10.2 Phase portraits

Phase portraits provide an analytical tool to study systems of first-order differential equations [735]. The method has proved to be useful in characterizing oriented texture [432, 736].

Let $p(t)$ and $q(t)$ denote two differentiable functions of time t , related by a system of first-order differential equations as

$$\begin{aligned}\dot{p}(t) &= F[p(t), q(t)] \\ \dot{q}(t) &= G[p(t), q(t)]\end{aligned}\quad (8.89)$$

where the dot above the variable indicates the first-order derivative of the function with respect to time, and F and G represent functions of p and q . Given initial conditions $p(0)$ and $q(0)$, the solution $[p(t), q(t)]$ to Equation 8.89 can be viewed as a parametric trajectory of a hypothetical particle in the (p, q) plane, placed at $[p(0), q(0)]$ at time $t = 0$, and moving through the (p, q) plane with velocity $[\dot{p}(t), \dot{q}(t)]$. The (p, q) plane is referred to as the *phase plane* of the system of first-order differential equations. The path traced by the hypothetical particle is called a *streamline* of the vector field (\dot{p}, \dot{q}) . The *phase portrait* is a graph of the possible streamlines in the phase plane. A *fixed point* of Equation 8.89 is a point in the phase plane where $\dot{p}(t) = 0$ and $\dot{q}(t) = 0$: a particle left at a fixed point remains stationary.

When the system of first-order differential equations is linear, Equation 8.89 assumes the form

$$\begin{bmatrix} \dot{p}(t) \\ \dot{q}(t) \end{bmatrix} = \mathbf{A} \begin{bmatrix} p(t) \\ q(t) \end{bmatrix} + \mathbf{b}, \quad (8.90)$$

where \mathbf{A} is a 2×2 matrix and \mathbf{b} is a 2×1 column matrix (a vector). In this case, there are only three types of phase portraits: node, saddle, and spiral [735]. The type of phase portrait can be determined from the nature of the eigenvalues of \mathbf{A} , as shown in Table 8.4. The center (p_0, q_0) of the phase portrait is given by the fixed point of Equation 8.90:

$$\begin{bmatrix} \dot{p}(t) \\ \dot{q}(t) \end{bmatrix} = 0 \Rightarrow \begin{bmatrix} p_0 \\ q_0 \end{bmatrix} = -\mathbf{A}^{-1}\mathbf{b}. \quad (8.91)$$

Solving Equation 8.90 yields a linear combination of complex exponentials for $p(t)$ and $q(t)$, whose exponents are given by the eigenvalues of A multiplied by the time variable t . Table 8.4 illustrates the streamlines obtained by solving Equation 8.90 for a node, a saddle, and a spiral phase portrait: the solid lines indicate the movement of the $p(t)$ and the $q(t)$ components of the solution, and the dashed lines indicate the streamlines. The formation of each phase portrait type is explained as follows:

- *Node*: the components $p(t)$ and $q(t)$ are exponentials that either simultaneously converge to, or diverge from, the fixed-point coordinates p_0 and q_0 .

- *Saddle*: the components $p(t)$ and $q(t)$ are exponentials; while one of the components [either $p(t)$ or $q(t)$] converges to the fixed point, the other diverges from the fixed point.
- *Spiral*: the components $p(t)$ and $q(t)$ are exponentially modulated sinusoidal functions — the resulting streamline forms a spiral curve.

Associating the functions $p(t)$ and $q(t)$ with the x and y coordinates of the Cartesian (image) plane, we can define the *orientation field* generated by Equation 8.90 as

$$\phi(x, y | \mathbf{A}, \mathbf{b}) = \arctan \left(\frac{\dot{q}(t)}{\dot{p}(t)} \right), \quad (8.92)$$

which is the angle of the velocity vector $[\dot{p}(t), \dot{q}(t)]$ with the x axis at $(x, y) = [p(t), q(t)]$. Table 8.4 lists the three phase portraits and the corresponding orientation fields generated by a system of linear first-order differential equations.

Using the concepts presented above, the orientation field of a textured image may be described qualitatively by determining the type of the phase portrait that is most similar to the orientation field, along with the center of the phase portrait. This notion was employed by Ayres and Rangayyan [595, 679, 680] to characterize architectural distortion.

8.10.3 Estimating the orientation field

The analysis of oriented texture is an important task in computer vision, and several methods have been proposed for this task, such as weighted angle histograms [653], integral curves [678], and phase portraits [432, 736]. Ayres and Rangayyan [595, 679, 680] analyzed the orientation field in ROIs of mammograms, through the usage of phase portraits, in order to characterize architectural distortion. In order to extract the texture orientation at each pixel of the ROI, the ROI was filtered with a bank of Gabor filters of different orientations [384] (see Section 8.4). The Gabor filter kernel oriented at the angle $\theta = -\pi$ was formulated as







$$g(x, y) = \frac{1}{2\pi\sigma_x\sigma_y} \exp \left[-\frac{1}{2} \left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} \right) \right] \cos(2\pi f_o x). \quad (8.93)$$

Kernels at other angles were obtained by rotating this kernel. A set of 180 kernels was used, with angles spaced evenly over the range $[-\pi/2, \pi/2]$.

Gabor filters may be used as line detectors. In the work of Ayres and Rangayyan [595, 679, 680], the parameters in Equation 8.93, namely σ_x , σ_y , and f_o , were derived from a design rule as follows. Let τ be the thickness of the line detector. This parameter constrains σ_x and f_o as follows:

- The amplitude of the exponential term in Equation 8.93, that is, the Gaussian term, is reduced to one half of its maximum at $x = \tau/2$ and $y = 0$; therefore, $\sigma_x = \tau / (2\sqrt{2 \ln 2})$.

TABLE 8.4
Phase Portraits for a System of Linear First-order Differential Equations [736].

Phase portrait type	Eigenvalues	Streamlines	Appearance of the orientation field
Node	Real eigenvalues of same sign		
Saddle	Real eigenvalues of opposite sign		
Spiral	Complex eigenvalues		

Solid lines indicate the movement of the $p(t)$ and the $q(t)$ components of the solution; dashed lines indicate the streamlines. Reproduced with permission from F.J. Ayres and R.M. Rangayyan, "Characterization of architectural distortion in mammograms via analysis of oriented texture", *IEEE Engineering in Medicine and Biology Magazine*, January 2005. © IEEE.

- The cosine term has a period of τ ; therefore, $f_o = 1/\tau$.

The value of σ_y was defined as $\sigma_y = l \sigma_x$, where l determines the elongation of the Gabor filter in the orientation direction, with respect to its thickness. The values $\tau = 4$ pixels (corresponding to a thickness of 0.8 mm at a pixel size of $200 \mu\text{m}$) and $l = 8$ were determined empirically, by observing the typical spicule width and length in mammograms with architectural distortion in the Mini-MIAS database [376].

The effects of the different design parameters are shown in Figure 8.68, and are as follows:

- Figures 8.68 (a) and (e) show the impulse response of a Gabor filter and its Fourier transform magnitude, respectively.
- In Figure 8.68 (b), the Gabor filter of Figure 8.68 (a) is stretched in the x direction, by increasing the elongation factor l . Observe that the Fourier spectrum of the new Gabor filter, shown in Figure 8.68 (f), is compressed in the horizontal direction.
- The Gabor filter shown in Figure 8.68 (c) was obtained by increasing the parameter τ of the original Gabor filter, thus enlarging the filter in both the x and y directions. Correspondingly, the Fourier spectrum of the enlarged filter, shown in Figure 8.68 (g), has been shrunk in both the vertical and horizontal directions.
- The effect of rotating the Gabor filter by 30° counterclockwise is displayed in Figures 8.68 (d) and (h), that show the rotated Gabor filter's impulse response and the corresponding Fourier spectrum.

The texture orientation at a pixel was estimated as the orientation of the Gabor filter that yielded the highest magnitude response at that pixel. The orientation at every pixel was used to compose the orientation field. The magnitude of the corresponding filter response was used to form the *magnitude image*. The magnitude image was not used in the estimation of the phase portrait, but was found to be useful for illustrative purposes.

Let $\theta(x, y)$ be the texture orientation at (x, y) , and $g_k(x, y)$, $k = 0, 1, \dots, 179$, be the Gabor filter oriented at $\alpha_k = -\pi/2 + \pi k/180$. Let $f(x, y)$ be the ROI of the mammogram being processed, and $f_k(x, y) = (f * g_k)(x, y)$, where the asterisk denotes linear 2D convolution. Then, the orientation field of $f(x, y)$ is given by

$$\theta(x, y) = \alpha_{k_{\max}} \quad \text{where } k_{\max} = \arg\{\max_k [|f_k(x, y)|]\}. \quad (8.94)$$

8.10.4 Characterizing orientation fields with phase portraits

In the work of Ayres and Rangayyan [595, 679, 680], the analysis of oriented texture patterns was performed in a two-step process. First, the orientation




FIGURE 8.68

Effects of the different parameters of the Gabor filter. (a) Example of the impulse response of a Gabor filter. (b) The parameter l is increased: the Gabor filter is elongated in the x direction. (c) The parameter τ is increased: the Gabor filter is enlarged in the x and y directions. (d) The angle of the Gabor filter is modified. Figures (e) – (h) correspond to the magnitude of the Fourier transforms of the Gabor filters in (a) – (d), respectively. The $(0, 0)$ frequency component is at the center of the spectra displayed. Reproduced with permission from F.J. Ayres and R.M. Rangayyan, “Characterization of architectural distortion in mammograms via analysis of oriented texture”, *IEEE Engineering in Medicine and Biology Magazine*, January 2005. © IEEE.

field $\theta(x, y)$ of the ROI was computed in a small analysis window. The sliding analysis window was centered at pixels within the ROI, avoiding window positions with incomplete data at the edges of the ROI for the estimation of \mathbf{A} and \mathbf{b} . Second, the matrix \mathbf{A} and the vector \mathbf{b} in Equation 8.90 were estimated such that the best match was achieved between $\theta(x, y)$ and $\phi(x, y|\mathbf{A}, \mathbf{b})$. The eigenvalues of \mathbf{A} determine the type of the phase portrait present in $\theta(x, y)$; the fixed point of the phase portrait is given by Equation 8.91.

The estimates of \mathbf{A} and \mathbf{b} were obtained as follows. For every point (x, y) , let $\Delta(x, y) = \sin[\theta(x, y) - \phi(x, y|\mathbf{A}, \mathbf{b})]$ represent the error between the orientation of the texture given by Equation 8.94 and the orientation of the model given by Equation 8.92. Then, the estimation problem is that of finding \mathbf{A} and \mathbf{b} that minimize the sum of the squared error

$$\epsilon^2 = \sum_x \sum_y \Delta^2(x, y) = \sum_x \sum_y \{\sin[\theta(x, y) - \phi(x, y|\mathbf{A}, \mathbf{b})]\}^2, \quad (8.95)$$

which may be solved using a nonlinear least-squares algorithm [737].

The ROI was investigated by sliding the analysis window through the orientation field of the ROI, and accumulating the information obtained, that is, the type of the phase portrait and the location of the fixed point, for each window position, as follows:

1. Create three maps, one for each type of phase portrait (hereafter called the *phase portrait maps*), that will be used to accumulate information from the sliding analysis window. The maps are initialized to zero, and are of the same size as the ROI or the image being processed.
2. Slide the analysis window through the orientation field of the ROI. At each position of the sliding window, determine the type of the phase portrait and compute the fixed point of the orientation field.
3. Increment the value at the location of the fixed point in the corresponding phase portrait map.

The size of the sliding analysis window was set at 44×44 pixels (8.8×8.8 mm).

The three maps obtained as above provide the results of a voting procedure, and indicate the possible locations of fixed points corresponding to texture patterns that (approximately) match the node, saddle, and spiral phase portraits. It is possible that, for some positions of the sliding analysis window, the location of the fixed point falls outside the spatial limits of the ROI or image being processed; the votes related to such results were ignored. The value at each location (x, y) in a phase portrait map provides the degree of confidence in determining the existence of the corresponding phase portrait type centered at (x, y) . The three phase portraits were expected to relate to different types of architectural distortion.

8.10.5 Feature extraction for pattern classification

The estimates of the fixed-point location for a given phase portrait pattern can be scattered around the true fixed-point position, due to the limited precision of the estimation procedure, the presence of multiple overlapping patterns, the availability of limited data within the sliding analysis window, and the presence of noise. A local accumulation of the votes is necessary to diminish the effect of fixed-point location errors. Ayres and Rangayyan [595, 679, 680] employed a Gaussian smoothing filter with a standard deviation of 25 pixels (5 mm) for this purpose.

For the purpose of pattern classification, six features were extracted to characterize each ROI: the maximum of each phase portrait map (three features), and the entropy of each phase portrait map (three features). The maximum of each map conveys information about the likelihood of the presence of the corresponding phase portrait type, and the entropy relates to the uncertainty in the location of the fixed point in each map. The entropy H of a map $h(x, y)$ was computed as

$$H[h(x, y)] = - \sum_x \sum_y \frac{h(x, y)}{S_h} \ln \left[\frac{h(x, y)}{S_h} \right], \quad (8.96)$$

where

$$S_h = \sum_x \sum_y h(x, y). \quad (8.97)$$

A map with a dense spatial concentration of votes is expected to have a large maximum value and a low entropy. On the contrary, a map with a wide scatter of votes may be expected to have a low maximum and a large entropy.

8.10.6 Application to segments of mammograms

Ayres and Rangayyan [595, 679] analyzed a set of 106 ROIs, each of size 230×230 pixels (46×46 mm, with a resolution of $200 \mu\text{m}$), selected from the Mini-MIAS database [376]. The set included 17 ROIs with architectural distortion (all the cases of architectural distortion available in the MIAS database), 45 ROIs with normal tissue patterns, eight ROIs with spiculated malignant masses, four ROIs with circumscribed malignant masses, 11 ROIs with spiculated benign masses, 19 ROIs with circumscribed benign masses, and two ROIs with malignant calcifications. The size of the ROIs was chosen to accommodate the largest area of architectural distortion or mass identified in the Mini-MIAS database. ROIs related to all of the masses in the database were included. The normal ROIs included examples of overlapping ducts, ligaments, and other parenchymal patterns. Only the central portion of 150×150 pixels of each ROI was investigated using a sliding analysis window of size 44×44 pixels; the remaining outer ribbon of pixels was not processed in order to discard the effects of the preceding filtering steps.

TABLE 8.5
Results of Linear Discriminant Analysis for ROIs with
Architectural Distortion Using the Leave-one-out Method.

Architectural distortion	#ROIs	Classified as	
		Architectural distortion	Other
Benign	9	7	2
Malignant	8	6	2
Total	17	TP = 13	FN = 4

TP = true positives, FN = false negatives. The results correspond to the prior probability of belonging to the architectural distortion class being 0.465. Sensitivity = 76.5%. Reproduced with permission from F.J. Ayres and R.M. Rangayyan, "Characterization of architectural distortion in mammograms via analysis of oriented texture", *IEEE Engineering in Medicine and Biology Magazine*, January 2005. © IEEE.

Figure 8.69 illustrates the results obtained for an image with architectural distortion (mdb115). The maximum of the node map is larger than the maxima of the other two maps. Also, the scattering of votes in the node map is less than that in the saddle and spiral maps. These results indicate that the degree of scattering of the votes (quantified by the entropy of the corresponding map) and the maximum of each of the three phase portrait maps could be useful features to distinguish between architectural distortion and other patterns.

Linear discriminant analysis was performed using SPSS [738], with stepwise feature selection. Architectural distortion was considered as a positive finding, with all other test patterns (normal tissue, masses, and calcifications) being considered as negative findings. The statistically significant features were the entropy of the node map and the entropy of the spiral map: the other features were deemed to be not significant by the statistical analysis package, and were discarded in all subsequent analysis. With the prior probability of architectural distortion set to 50%, the sensitivity obtained was 82.4%, and the specificity was 71.9%. The fraction of cases correctly classified was 73.6%.

Tables 8.5 and 8.6 present the classification results with the prior probability of architectural distortion being 46.5%. An overall classification accuracy of 76.4% was achieved.

8.10.7 Detection of sites of architectural distortion

Ayres and Rangayyan [595, 679, 680] hypothesized that architectural distortion would appear as an oriented texture pattern that can be locally approximated by a linear phase portrait model; furthermore, it was expected that the




FIGURE 8.69

Analysis of the ROI from the image mdb115, which includes architectural distortion: (a) ROI of size 230×230 pixels (46×46 mm); (b) magnitude image; (c) orientation field superimposed on the original ROI; (d) node map, with intensities mapped from $[0, 123]$ to $[0, 255]$; (e) saddle map, $[0, 22]$ mapped to $[0, 255]$; (f) spiral map, $[0, 71]$ mapped to $[0, 255]$. This image was correctly classified as belonging to the “architectural distortion” category (Table 8.5). Reproduced with permission from F.J. Ayres and R.M. Rangayyan, “Characterization of architectural distortion in mammograms via analysis of oriented texture”, *IEEE Engineering in Medicine and Biology Magazine*, January 2005. © IEEE.

TABLE 8.6

Results of Linear Discriminant Analysis for ROIs Without Architectural Distortion Using the Leave-one-out Method.

Type	#ROIs	Classified as		
		Architectural distortion	Other	
Masses	CB	19	4	15
	SB	11	3	8
	CM	4	1	3
	SM	8	3	5
Calcifications	2	1	1	
Normal	45	9	36	
Total	89	FP = 21	TN = 68	

CB = circumscribed benign mass, CM = circumscribed malignant tumor, SB = spiculated benign mass, SM = spiculated malignant tumor, FP = false positives, TN = true negatives. The results correspond to the prior probability of belonging to the architectural distortion class being 0.465. Specificity = 76.4%. Reproduced with permission from F.J. Ayres and R.M. Rangayyan, "Characterization of architectural distortion in mammograms via analysis of oriented texture", *IEEE Engineering in Medicine and Biology Magazine*, January 2005. © IEEE.

fixed-point location of the phase portrait model would fall within the breast area in the mammogram. Then, the numbers of votes cast at each position of the three phase portrait maps would indicate the likelihood that the position considered is a fixed point of a node, a saddle, or a spiral pattern.

Before searching the maps for sites of distortion, the orientation field was filtered and downsampled as follows. Let $h(x, y)$ be a Gaussian filter of standard deviation σ_h , defined as

$$h(x, y) = \frac{1}{2\pi\sigma_h} \exp \left[-\frac{1}{2} \left(\frac{x^2 + y^2}{\sigma_h^2} \right) \right]. \quad (8.98)$$

Define the images $s(x, y) = \sin[2\theta(x, y)]$ and $c(x, y) = \cos[2\theta(x, y)]$, where $\theta(x, y)$ is the orientation field. Then, the filtered orientation field $\theta_f(x, y)$ is obtained as

$$\theta_f(x, y) = \frac{1}{2} \arctan \left(\frac{h(x, y) * s(x, y)}{h(x, y) * c(x, y)} \right), \quad (8.99)$$

where the asterisk denotes 2D convolution.

The filtered orientation field was downsampled by a factor of four, thus producing the downsampled orientation field θ_d as

$$\theta_d(x, y) = \theta_f(4x, 4y). \quad (8.100)$$

The filtering and downsampling procedures, summarized in Figure 8.70, were applied in order to reduce noise and also to reduce the computational effort required for the processing of full mammograms. The filtering procedure described above is a variant of Rao's dominant local orientation method [432]: a Gaussian filter has been used instead of a box filter.

The following procedure was used to detect and locate sites of architectural distortion, using only the node map:

1. The node map is filtered with a Gaussian filter of standard deviation equal to 1.0 pixel (0.8 mm).
2. The filtered node map is thresholded (with the same threshold value for all images).
3. The thresholded image is subjected to the following series of morphological operations to group positive responses that are close to one another, and to reduce each region of positive response to a single point. The resulting points indicate the detected locations of architectural distortion.
 - (a) A closing operation is performed to group clusters of points that are less than 8 mm apart. The structural element is a disk of radius 10 pixels (8 mm).
 - (b) A "close holes" filter is applied to the image. The resulting image includes only compact regions.

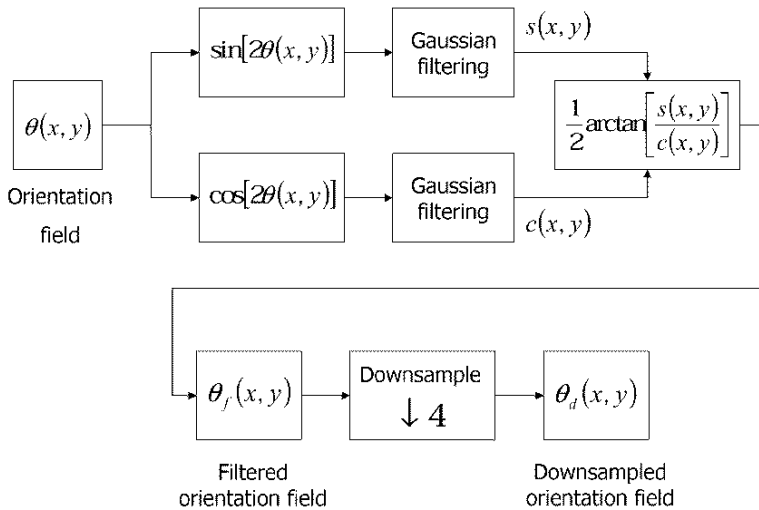


FIGURE 8.70

Filtering and downsampling of the orientation field. Figure courtesy of F.J. Ayres.

- (c) The image is subjected to a “shrink” filter, where each compact region is shrunk to a single pixel.

The threshold value influences the sensitivity of the method and the number of false positives per image. A high threshold value reduces the number of false positives, but also reduces the sensitivity. A low threshold value increases the number of false positives.

The method was applied to 18 mammograms exhibiting architectural distortion, selected from the Mini-MIAS database [376]. The mammograms were MLO views, digitized to $1,024 \times 1,024$ pixels at a resolution of $200 \mu m$ and $8 b/pixel$. Figures 8.71 and 8.72 illustrate the steps of the method, as applied to image mdb115. Observe that the filtered orientation field [Figure 8.71 (d)] is smoother and more coherent as compared to the original orientation field [Figure 8.71 (c)]: the pattern of architectural distortion is displayed better in the filtered orientation field.

The architectural distortion present in the mammogram mdb115 has a stellate or spiculated appearance. As a consequence, a large number of votes have been cast into the node map, at a location close to the center of the distortion, as seen in Figure 8.72 (c). Another point of accumulation of votes in the node map is observed in Figure 8.72 (c), at the location of the nipple. This is not unexpected: the breast has a set of ducts that carry milk to the nipple; the ducts appear in mammograms as linear structures converging to the nipple. Observe that the node map has the strongest response of all maps,




Figure 8.71 (a)

(b)

within the site of architectural distortion given by the Mini-MIAS database. The technique has resulted in the identification of two locations as sites of architectural distortion: one true positive and one false positive, as shown in Figure 8.72 (d).

The free-response receiver operating characteristics (FROC) was derived by varying the threshold level in the detection step; the result is shown in Figure 8.73. (See Section 12.8.1 for details on ROC analysis.) A sensitivity of 88% was obtained at 15 false positives per image. Further work is required in order to reduce the number of false positives and improve the accuracy of detection.

8.11 Remarks

Preferred orientation and directional distributions relate to the functional integrity of several types of tissues and organs; changes in such patterns could indicate structural damage as well as recovery. Directional analysis could,




FIGURE 8.71

(a) Image mdb115 from the Mini-MIAS database [376]. The circle indicates the location and the extent of architectural distortion, as provided in the Mini-MIAS database [376]. (b) Magnitude image after Gabor filtering. (c) Orientation field superimposed on the original image. Needles have been drawn for every fifth pixel. (d) Filtered orientation field superimposed on the original image. Reproduced with permission from F.J. Ayres and R.M. Rangayyan, "Detection of architectural distortion in mammograms using phase portraits", *Proceedings of SPIE Medical Imaging 2004: Image Processing*, Volume 5370, pp 587 – 597, 2004. © SPIE. See also Figure 8.72.





Figure 8.72 (a)



(b)




FIGURE 8.72

Phase portrait maps derived from the orientation field in Figure 8.71 (d), and the detection of architectural distortion. (a) Saddle map: values are scaled from the range $[0, 20]$ to $[0, 255]$. (b) Spiral map: values are scaled from the range $[0, 47]$ to $[0, 255]$. (c) Node map: values are scaled from the range $[0, 84]$ to $[0, 255]$. (d) Detected sites of architectural distortion superimposed on the original image: the solid line indicates the location and spatial extent of architectural distortion as given by the Mini-MIAS database [376]; the dashed lines indicate the detected sites of architectural distortion (one true positive and one false positive). Reproduced with permission from F.J. Ayres and R.M. Rangayyan, “Detection of architectural distortion in mammograms using phase portraits”, *Proceedings of SPIE Medical Imaging 2004: Image Processing*, Volume 5370, pp 587 – 597, 2004. © SPIE.




FIGURE 8.73

Free-response receiver operating characteristics (FROC) curve for the detection of sites of architectural distortion. Reproduced with permission from F.J. Ayres and R.M. Rangayyan, "Detection of architectural distortion in mammograms using phase portraits", *Proceedings of SPIE Medical Imaging 2004: Image Processing*, Volume 5370, pp 587 – 597, 2004. © SPIE.

therefore, be used to study the health and well-being of a tissue or organ, as well as to follow the pathological and physiological processes related to injury, treatment, and healing.

In this chapter, we explored the directional characteristics of several biomedical images. We have seen several examples of the application of fan filters and Gabor wavelets. The importance of multiscale or multiresolution analysis in accounting for variations in the size of the pattern elements of interest has been demonstrated. In spite of theoretical limitations, the methods for directional analysis presented in this chapter have been shown to lead to practically useful results in important applications.

8.12 Study Questions and Problems

Selected data files related to some of the problems and exercises are available at the site

www.enel.ucalgary.ca/People/Ranga/enel697

1. Discuss how entropy can characterize a directional distribution.
2. Discuss the limitations of fan filters. Describe how Gabor functions address these limitations.
3. Using an image with line segments of various widths as an example, discuss the need for multiscale or multiresolution analysis.

8.13 Laboratory Exercises and Projects

1. Prepare a test image with line segments of different directions, lengths, and widths, with overlap. Apply Gabor filters at a few different scales and angles, as appropriate. Evaluate the results in terms of
 - (a) the lengths of the extracted components, and
 - (b) the widths of the extracted components.Discuss the limitations of the methods and the artifacts in the results.
2. Decompose your test image in the preceding problem using eight sector or fan filters spanning the full range of $0^\circ - 180^\circ$.
Apply a thresholding technique to binarize the results.
Compute the area covered by the filtered patterns for each angle band. Compare the results with the known areas of the directional patterns.
Discuss the limitations of the methods and the artifacts in the results.

Image Reconstruction from Projections

Mathematically, the fundamental problem in CT imaging is that of estimating an image (or object) from its projections (or integrals) measured at different angles [11, 43, 80, 82, 739, 740, 741, 742, 743]; see Section 1.6 for an introduction to this topic. A projection of an image is also referred to as the *Radon transform* of the image at the corresponding angle, after the main proponent of the associated mathematical principles [67, 68]. In the continuous space, the projections are ray integrals of the image, measured at different ray positions and angles; in practice, only discrete measurements are available. The solution to the problem of image reconstruction from projections may be formulated variously as completing the corresponding Fourier space, backprojecting and summing the given projection data, or solving a set of simultaneous equations. Each of these methods has its own advantages and disadvantages that determine its suitability to a particular imaging application. In this chapter, we shall study the three basic approaches to image reconstruction from projections mentioned above.

(*Note:* Most of the derivations presented in this chapter closely follow those of Rosenfeld and Kak [11], with permission. For further details, refer to Herman [80, 43] and Kak and Slaney [82]. Parts of this chapter are reproduced, with permission, from R.M. Rangayyan and A. Kantzas, “Image reconstruction”, *Wiley Encyclopedia of Electrical and Electronics Engineering*, Supplement 1, Editor: J. G. Webster, Wiley, New York, NY, pp 249–268, 2000. © This material is used by permission of John Wiley & Sons, Inc.)

9.1 Projection Geometry

Let us consider the problem of reconstructing a 2D image given parallel-ray projections of the image measured at different angles. Referring to Figure 9.1, let $f(x, y)$ represent the density distribution within the image. Although discrete images are used in practice, the initial presentation here will be in continuous-space notations for easier comprehension. Consider the ray AB represented by the equation

$$x \cos \theta + y \sin \theta = t_1. \quad (9.1)$$

The integral of $f(x, y)$ along the ray path AB is given by

$$p_\theta(t_1) = \int_{\text{AB}} f(x, y) ds = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) \delta(x \cos \theta + y \sin \theta - t_1) dx dy, \quad (9.2)$$

where $\delta(\)$ is the Dirac delta function, and $s = -x \sin \theta + y \cos \theta$. The mutually parallel rays within the imaging plane are represented by the coordinates (t, s) that are rotated by angle θ with respect to the (x, y) coordinates as indicated in Figures 1.9, 1.19, and 9.1, with the s axis being parallel to the rays; ds is thus the elemental distance along a ray. When this integral is evaluated for different values of the ray offset t_1 , we obtain the 1D projection $p_\theta(t)$. The function $p_\theta(t)$ is known as the *Radon transform* of $f(x, y)$. [Note: Whereas a single projection $p_\theta(t)$ of a 2D image at a given value of θ is a 1D function, a set of projections for various values of θ could be seen as a 2D function. Observe that t represents the space variable related to ray displacement along a projection, and not time.] Because the various rays within a projection are parallel to one another, this is known as parallel-ray geometry.

Theoretically, we would need an infinite number of projections for all θ to be able to reconstruct the image. Before we consider reconstruction techniques, let us take a look at the projection or Fourier slice theorem.

9.2 The Fourier Slice Theorem

The projection or Fourier slice theorem relates the three spaces we encounter in image reconstruction from projections: the image, Fourier, and projection (Radon) spaces. Considering a 2D image, the theorem states that the 1D Fourier transform of a 1D projection of the 2D image is equal to the radial section (slice or profile) of the 2D Fourier transform of the 2D image at the angle of the projection. This is illustrated graphically in Figure 9.2, and may be derived as follows.

Let $F(u, v)$ represent the 2D Fourier transform of $f(x, y)$, given by

$$F(u, v) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) \exp[-j2\pi(ux + vy)] dx dy. \quad (9.3)$$

Let $P_\theta(w)$ represent the 1D Fourier transform of the projection $p_\theta(t)$, that is,

$$P_\theta(w) = \int_{-\infty}^{\infty} p_\theta(t) \exp(-j2\pi wt) dt, \quad (9.4)$$

where w represents the frequency variable corresponding to t . (Note: If x, y, s , and t are in *mm*, the units for u, v , and w will be *cycles/mm* or *mm⁻¹*.) Let


**FIGURE 9.1**

Illustration of a ray path AB through a sectional plane or image $f(x, y)$. The (t, s) axis system is rotated by angle θ with respect to the (x, y) axis system. ds represents the elemental distance along the ray path AB . $p_{\theta}(t_1)$ is the ray integral of $f(x, y)$ for the ray path AB . $p_{\theta}(t)$ is the parallel-ray projection (Radon transform or integral) of $f(x, y)$ at angle θ . See also Figures 1.9 and 1.19. Adapted, with permission, from A. Rosenfeld and A.C. Kak, *Digital Picture Processing*, 2nd ed., New York, NY, 1982. © Academic Press.




FIGURE 9.2

Illustration of the Fourier slice theorem. $F(u, v)$ is the 2D Fourier transform of $f(x, y)$. $F(w, \theta_1) = P_{\theta_1}(w)$ is the 1D Fourier transform of $p_{\theta_1}(t)$. $F(w, \theta_2) = P_{\theta_2}(w)$ is the 1D Fourier transform of $p_{\theta_2}(t)$. Reproduced, with permission, from R.M. Rangayyan and A. Kantzas, "Image reconstruction", *Wiley Encyclopedia of Electrical and Electronics Engineering*, Supplement 1, Editor: J. G. Webster, Wiley, New York, NY, pp 249–268, 2000. © This material is used by permission of John Wiley & Sons, Inc.

$f(t, s)$ represent the image $f(x, y)$ rotated by angle θ , with the transformation given by

$$\begin{bmatrix} t \\ s \end{bmatrix} = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}. \quad (9.5)$$

Then,

$$p_{\theta}(t) = \int_{-\infty}^{\infty} f(t, s) ds. \quad (9.6)$$

$$\begin{aligned} P_{\theta}(w) &= \int_{-\infty}^{\infty} p_{\theta}(t) \exp(-j2\pi wt) dt \\ &= \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} f(t, s) ds \right] \exp(-j2\pi wt) dt. \end{aligned} \quad (9.7)$$

Transforming from (t, s) to (x, y) , we get

$$\begin{aligned} P_{\theta}(w) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) \exp[-j2\pi w(x \cos \theta + y \sin \theta)] dx dy \\ &= F(u, v) \text{ for } u = w \cos \theta, v = w \sin \theta \\ &= F(w, \theta), \end{aligned} \quad (9.8)$$

which expresses the projection theorem. Observe that $t = x \cos \theta + y \sin \theta$ and $dx dy = ds dt$.

It immediately follows that if we have projections available at all angles from 0° to 180° , we can take their 1D Fourier transforms, fill the 2D Fourier space with the corresponding radial sections or slices, and take an inverse 2D Fourier transform to obtain the image $f(x, y)$. The difficulty lies in the fact that, in practice, only a finite number of projections will be available, measured at discrete angular positions or steps. Thus, some form of interpolation will be essential in the 2D Fourier space [72, 73]. Extrapolation may also be required if the given projections do not span the entire angular range. This method of reconstruction from projections, known as the Fourier method, succinctly relates the image, Fourier, and Radon spaces. The Fourier method is the most commonly used method for the reconstruction of MR images.

A practical limitation of the Fourier method of reconstruction is that interpolation errors are larger for higher frequencies due to the increased spacing between the samples available on a discrete grid. Samples of $P_\theta(w)$ computed from $p_\theta(t)$ will be available on a polar grid, whereas the 2D Fourier transform $F(u, v)$ and/or the inverse-transformed image will be required on a Cartesian (rectangular grid). This limitation could cause poor reconstruction of high-frequency (sharp) details.

9.3 Backprojection

Let us now consider the simplest reconstruction procedure: backprojection (BP). Assuming the rays to be ideal straight lines, rather than strips of finite width, and the image to be made of dimensionless points rather than pixels or voxels of finite size, it can be seen that each point in the image $f(x, y)$ contributes to only one ray integral per parallel-ray projection $p_\theta(t)$, with $t = x \cos \theta + y \sin \theta$. We may obtain an estimate of the density at a point by simply summing (integrating) all rays that pass through it at various angles, that is, by backprojecting the individual rays. In doing so, however, the contributions to the various rays of all of the other points along their paths are also added up, causing smearing or blurring; yet this method produces a reasonable estimate of the image. Mathematically, simple BP can be expressed as [11]

$$f(x, y) \simeq \int_0^\pi p_\theta(t) d\theta, \quad \text{where } t = x \cos \theta + y \sin \theta. \quad (9.9)$$

This is a sinusoidal path of integration in the (θ, t) Radon space. In practice, only a finite number of projections and a finite number of rays per projection will be available, that is, the (θ, t) space will be discretized; hence, interpolation will be required.

Examples of reconstructed images: Figure 9.3 (a) shows a synthetic 2D image (phantom), which we will consider to represent a cross-sectional plane of a 3D object. The objects in the image were defined on a discrete grid, and hence have step and/or jagged edges. Figure 9.4 (a) is a plot of the projection of the phantom image computed at 90° ; observe that the values are all positive.





FIGURE 9.3


(a) A synthetic 2D image (phantom) with 101×101 eight-bit pixels, representing a cross-section of a 3D object. (b) Reconstruction of the phantom in (a) obtained using 90 projections from 2° to 180° in steps of 2° with the simple BP algorithm. Reproduced, with permission, from R.M. Rangayyan and A. Kantzas, “Image reconstruction”, *Wiley Encyclopedia of Electrical and Electronics Engineering*, Supplement 1, Editor: J. G. Webster, Wiley, New York, NY, pp 249–268, 2000. © This material is used by permission of John Wiley & Sons, Inc.

Figure 9.3 (b) shows the reconstruction of the phantom obtained using 90 projections from 2° to 180° in steps of 2° with the simple BP algorithm. While the objects in the image are faintly visible, the smearing effect of the BP algorithm is obvious.

Considering a point source as the image to be reconstructed, it becomes evident that BP produces a spoke-like pattern with straight lines at all projection angles, intersecting at the position of the point source. This may be considered to be the PSF of the reconstruction process, which is responsible for the blurring of details.



(a)



(b)

FIGURE 9.4

(a) Projection of the phantom image in Figure 9.3 (a) computed at 90° . (b) Filtered version of the projection using only the ramp filter inherent to the FBP algorithm. Reproduced, with permission, from R.M. Rangayyan and A. Kantzas, "Image reconstruction", *Wiley Encyclopedia of Electrical and Electronics Engineering*, Supplement 1, Editor: J. G. Webster, Wiley, New York, NY, pp 249–268, 2000. © This material is used by permission of John Wiley & Sons, Inc.

The use of limited projection data in reconstruction results in geometric distortion and streaking artifacts [744, 745, 746, 747, 748]. The distortion may be modeled by the PSF of the reconstruction process if it is linear and shift-invariant; this condition is satisfied by the BP process. The PSFs of the simple BP method are shown as images in Figure 9.5 (a) for the case with 10 projections over 180° , and in Figure 9.5 (b) for the case with 10 projections from 40° to 130° . The reconstructed image is given by the convolution of the original image with the PSF; the images in parts (c) and (d) of Figure 9.5 illustrate the corresponding reconstructed images of the phantom in Figure 9.3 (a). Limited improvement in image quality may be obtained by applying deconvolution filters to the reconstructed image [744, 745, 746, 747, 748, 749, 750, 751]. Deconvolution is implicit in the filtered (convolution) backprojection technique, which is described next.

9.3.1 Filtered backprojection

Consider the inverse Fourier transform relationship

$$f(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} F(u, v) \exp[j2\pi(ux + vy)] du dv. \quad (9.10)$$

Changing from the Cartesian coordinates (u, v) to the polar coordinates (w, θ) , where $w = \sqrt{u^2 + v^2}$ and $\theta = \tan^{-1}(v/u)$, we get

$$\begin{aligned} f(x, y) &= \int_0^{2\pi} \int_0^{\infty} F(w, \theta) \exp[j2\pi w(x \cos \theta + y \sin \theta)] w dw d\theta \\ &= \int_0^{\pi} \int_0^{\infty} F(w, \theta) \exp[j2\pi w(x \cos \theta + y \sin \theta)] w dw d\theta \\ &\quad + \int_0^{\pi} \int_0^{\infty} F(w, \theta + \pi) \\ &\quad \times \exp\{j2\pi w[x \cos(\theta + \pi) + y \sin(\theta + \pi)]\} w dw d\theta. \end{aligned} \quad (9.11)$$

Here, $u = w \cos \theta$, $v = w \sin \theta$, and $du dv = w dw d\theta$. Because $F(w, \theta + \pi) = F(-w, \theta)$, we get


$$\begin{aligned} f(x, y) &= \int_0^{\pi} \left[\int_{-\infty}^{\infty} F(w, \theta) |w| \exp(j2\pi wt) dw \right] d\theta \\ &= \int_0^{\pi} \left[\int_{-\infty}^{\infty} P_{\theta}(w) |w| \exp(j2\pi wt) dw \right] d\theta, \end{aligned} \quad (9.12)$$

with $t = x \cos \theta + y \sin \theta$ as before. If we define

$$q_{\theta}(t) = \int_{-\infty}^{\infty} P_{\theta}(w) |w| \exp(j2\pi wt) dw, \quad (9.13)$$

we get

$$f(x, y) = \int_0^{\pi} q_{\theta}(t) d\theta = \int_0^{\pi} q_{\theta}(x \cos \theta + y \sin \theta) d\theta. \quad (9.14)$$

**FIGURE 9.5**

PSF of the BP procedure using: (a) 10 projections from 18° to 180° in steps of 18° ; (b) 10 projections from 40° to 130° in steps of 10° . The images (a) and (b) have been enhanced with $\gamma = 3$. Reconstruction of the phantom in Figure 9.3 (a) obtained using (c) 10 projections as in (a) with the BP algorithm; (d) 10 projections as in (b) with the BP algorithm. Reproduced, with permission, from R.M. Rangayyan and A. Kantzas, "Image reconstruction", *Wiley Encyclopedia of Electrical and Electronics Engineering*, Supplement 1, Editor: J. G. Webster, Wiley, New York, NY, pp 249–268, 2000. © This material is used by permission of John Wiley & Sons, Inc.

It is now seen that a perfect reconstruction of $f(x, y)$ may be obtained by backprojecting *filtered* projections $q_\theta(t)$ instead of backprojecting the original projections $p_\theta(t)$; hence the name *filtered backprojection* (FBP). The filter is represented by the $|w|$ function, known as the *ramp filter*; see Figure 9.6.

Observe that the limits of integration in Equation 9.12 are $(0, \pi)$ for θ and $(-\infty, \infty)$ for w . In practice, a smoothing window should be applied to reduce the amplification of high-frequency noise by the $|w|$ function. Furthermore, the integrals change to summations in practice due to the finite number of projections available, as well as the discrete nature of the projections themselves and of the Fourier transform computations employed. (Details of the discrete version of FBP are provided in the next section.)

An important feature of the FBP technique is that each projection may be filtered and backprojected while further projection data are being acquired, which was of help in on-line processing with the first-generation CT scanners (see Figure 1.20). Furthermore, the inverse Fourier transform of the filter $|w|$ (with modifications to account for the discrete nature of measurements, smoothing window, etc.; see Figure 9.7) could be used to convolve the projections directly in the t space [74] using fast array processors. FBP is the most widely used procedure for image reconstruction from projections; however, the procedure provides good reconstructed images only when a large number of projections spanning the full angular range of 0° to 180° are available.

9.3.2 Discrete filtered backprojection

The filtering procedure with the $|w|$ function, in theory, must be performed over $-\infty \leq w \leq \infty$. In practice, the signal energy above a certain frequency limit W will be negligible, and $|w|$ filtering beyond the limit will only amplify noise. Thus, we may consider the projections to be bandlimited to $\pm W$. Then, using the sampling theorem, $p_\theta(t)$ can be represented by its samples at the sampling rate $2W$ as


$$p_\theta(t) = \sum_{m=-\infty}^{\infty} p_\theta\left(\frac{m}{2W}\right) \frac{\sin 2\pi W\left(t - \frac{m}{2W}\right)}{2\pi W\left(t - \frac{m}{2W}\right)}. \quad (9.15)$$

Then,

$$P_\theta(w) = \frac{1}{2W} \sum_{m=-\infty}^{\infty} p_\theta\left(\frac{m}{2W}\right) \exp\left[-j2\pi w \frac{m}{2W}\right] b_W(w), \quad (9.16)$$

where

$$b_W(w) = 1, \quad \text{if } |w| \leq W \\ = 0, \quad \text{otherwise.} \quad (9.17)$$

**FIGURE 9.6**

(a) The (bandlimited) ramp filter inherent to the FBP algorithm. (b) The ramp filter weighted with a Butterworth lowpass filter. The filters are shown for both positive and negative frequency along the w axis with $w = 0$ at the center. The corresponding filters are shown in the Radon domain (t space) in Figure 9.7.




FIGURE 9.7

(a) The inverse Fourier transform of the (bandlimited) ramp filter inherent to the FBP algorithm (in the t space). (b) The inverse Fourier transform of the ramp filter weighted with a Butterworth lowpass filter. The functions, which are used for convolution with projections in the CBP method, are shown for both $+t$ and $-t$, with $t = 0$ at the center. The corresponding filters are shown in the frequency domain (w space) in Figure 9.6.

If the projections are of finite order, that is, they can be represented by a finite number of samples $N + 1$, then

$$P_\theta(w) = \frac{1}{2W} \sum_{m=-N/2}^{N/2} p_\theta\left(\frac{m}{2W}\right) \exp\left[-j2\pi w \frac{m}{2W}\right] b_W(w). \quad (9.18)$$

Let us assume that N is even, and let the frequency axis be discretized as

$$w = k \frac{2W}{N} \text{ for } k = -\frac{N}{2}, \dots, 0, \dots, \frac{N}{2}. \quad (9.19)$$

Then,

$$P_\theta\left(k \frac{2W}{N}\right) = \frac{1}{2W} \sum_{m=-N/2}^{N/2} p_\theta\left(\frac{m}{2W}\right) \exp\left(-j \frac{2\pi}{N} mk\right), \quad (9.20)$$

$k = -\frac{N}{2}, \dots, 0, \dots, \frac{N}{2}$. This represents a DFT relationship, and may be evaluated using the FFT algorithm.

The filtered projection $q_\theta(t)$ may then be obtained as

$$q_\theta(t) = \int_{-W}^W P_\theta(w) |w| \exp(j2\pi wt) dw \quad (9.21)$$

$$\simeq \frac{2W}{N} \sum_{k=-N/2}^{N/2} P_\theta\left(k \frac{2W}{N}\right) \left|k \frac{2W}{N}\right| \exp\left[j2\pi k \frac{2W}{N} t\right]. \quad (9.22)$$

If we want to evaluate $q_\theta(t)$ for only those values of t at which $p_\theta(t)$ has been sampled, we get

$$q_\theta\left(\frac{m}{2W}\right) \simeq \frac{2W}{N} \sum_{k=-N/2}^{N/2} P_\theta\left(k \frac{2W}{N}\right) \left|k \frac{2W}{N}\right| \exp\left(j \frac{2\pi}{N} mk\right), \quad (9.23)$$

$$m = -\frac{N}{2}, \dots, -1, 0, 1, \dots, \frac{N}{2}.$$

In order to control noise enhancement by the $|k \frac{2W}{N}|$ filter, it may be beneficial to include a filter window such as the Hamming window; then,

$$\begin{aligned} q_\theta\left(\frac{m}{2W}\right) &\simeq \frac{2W}{N} \sum_{k=-N/2}^{N/2} P_\theta\left(k \frac{2W}{N}\right) \left|k \frac{2W}{N}\right| \\ &\times G\left(k \frac{2W}{N}\right) \exp\left(j \frac{2\pi}{N} mk\right), \end{aligned} \quad (9.24)$$

with

$$G\left(k \frac{2W}{N}\right) = 0.54 + 0.46 \cos\left(k \frac{2W}{N} \frac{\pi}{W}\right), \quad k = -\frac{N}{2}, \dots, 0, \dots, \frac{N}{2}. \quad (9.25)$$

Using the convolution theorem, we get

$$q_\theta \left(\frac{m}{2W} \right) \simeq \frac{2W}{N} p_\theta \left(\frac{m}{2W} \right) * g_1 \left(\frac{m}{2W} \right), \quad (9.26)$$

where $*$ denotes circular (periodic) convolution, and $g_1 \left(\frac{m}{2W} \right)$, $m = -\frac{N}{2}, \dots, 0, \dots, \frac{N}{2}$, is the inverse DFT of $|k \frac{2W}{N}| G \left(k \frac{2W}{N} \right)$. Butterworth or other lowpass filters may also be used instead of the Hamming window.

Observe that the inverse Fourier transform of $|w|$ does not exist due to the fact that $|w|$ is neither absolutely nor square integrable. However, if we consider the inverse Fourier transform of $|w| \exp(-\varepsilon|w|)$ as $\varepsilon \rightarrow 0$, we get the function [11]

$$p_\varepsilon(t) = \frac{\varepsilon^2 - (2\pi t)^2}{[\varepsilon^2 + (2\pi t)^2]^2}; \quad (9.27)$$

for large t , $p_\varepsilon(t) \simeq -\frac{1}{(2\pi t)^2}$.

The reconstructed image may be obtained as

$$\tilde{f}(x, y) = \frac{\pi}{L} \sum_{l=1}^L q_{\theta_l}(x \cos \theta_l + y \sin \theta_l), \quad (9.28)$$

where the L angles θ_l are those at which the projections $p_{\theta_l}(t)$ are available, and the (x, y) coordinates are discretized as appropriate.

For practical implementation of discrete FBP, let us consider the situation where the projections have been sampled with an interval of τ mm with no aliasing error. Each projection $p_{\theta_l}(m\tau)$ is then limited to the frequency band $(-W, W)$, with $W = \frac{1}{2\tau}$ cycles/mm. The continuous versions of the filtered projections are

$$q_{\theta_l}(t) = \int_{-\infty}^{\infty} P_{\theta_l}(w) H(w) \exp(j2\pi wt) dw, \quad (9.29)$$

where the filter $H(w) = |w| b_W(w)$, with $b_W(w)$ as defined in Equation 9.17. The impulse response of the filter $H(w)$ is [11]

$$h(t) = \frac{1}{2\tau^2} \frac{\sin(2\pi t/2\tau)}{2\pi t/2\tau} - \frac{1}{4\tau^2} \left(\frac{\sin(\pi t/2\tau)}{\pi t/2\tau} \right)^2. \quad (9.30)$$

Because we require $h(t)$ only at integral multiples of the sampling interval τ , we have

$$h(n\tau) = \begin{cases} \frac{1}{4\tau^2}, & n = 0 \\ 0, & n \text{ even} \\ -\frac{1}{n^2 \pi^2 \tau^2}, & n \text{ odd} \end{cases}. \quad (9.31)$$

The filtered projections $q_{\theta_l}(m\tau)$ may be obtained as

$$q_{\theta_l}(m\tau) = \tau \sum_{n=0}^{N-1} p_{\theta_l}(n\tau) h(m\tau - n\tau), \quad m = 0, 1, \dots, N-1, \quad (9.32)$$

where N is the finite number of samples in the projection $p_{\theta_l}(m\tau)$. Observe that $h(n\tau)$ is required for $n = -(N-1), \dots, 0, \dots, N-1$. When the filter is implemented as a convolution, the FBP method is also referred to as *convolution backprojection* (CBP).

The procedure for FBP may be expressed in algorithmic form as follows:

1. Measure the projection $p_{\theta_l}(m\tau)$.
2. Compute the filtered projection $q_{\theta_l}(m\tau)$.
3. Backproject the filtered projection $q_{\theta_l}(m\tau)$.
4. Repeat Steps 1 – 3 for all projection angles θ_l , $l = 1, 2, \dots, L$.

The FBP algorithm is suitable for on-line implementation in a *translate-rotate* CT scanner because each parallel-ray projection may be filtered and backprojected as soon as it is acquired, while the scanner is acquiring the next projection. The reconstructed image is ready as soon as the last projection is acquired, filtered, and backprojected. If the projections are acquired using fan-beam geometry (see Figure 1.20), one could either rebin the fan-beam data to compose parallel-ray projections, or use reconstruction algorithms specifically tailored to fan-beam geometry [11, 82].

Examples of reconstructed images: Figure 9.4 (b) shows a plot of the filtered version of the projection in Figure 9.4 (a) using only the ramp filter ($|w|$) inherent to the FBP algorithm. Observe that the filtered projection has negative values.

Figure 9.8 (a) shows the reconstruction of the phantom in Figure 9.3 (a) obtained using 90 projections with the FBP algorithm; only the ramp filter that is implicit in the FBP process was used, with no other smoothing or low-pass filter function. The contrast and visibility of the objects are better than those in the case of the simple BP result in Figure 9.3 (b); however, the image is noisy due to the increasing gain of the ramp filter at higher frequencies. The reconstructed image also exhibits artifacts related to the computation of the projections on a discrete grid; refer to Herman [43], Herman et al. [752], and Kak and Slaney [82] for discussions on this topic. The use of additional filters could reduce the noise and artifacts: Figure 9.8 (b) shows the result of reconstruction with the FBP algorithm including a fourth-order Butterworth filter with the -3 dB cutoff at 0.4 times the maximum frequency present in the data to filter the projections. The Butterworth filter has suppressed the noise and artifacts at the expense of blurring the edges of the objects in the image.




FIGURE 9.8


(a) Reconstruction of the phantom in Figure 9.3 (a) obtained using 90 projections from 2° to 180° in steps of 2° with the FBP algorithm; only the ramp filter that is implicit in the FBP process was used. (b) Reconstruction of the phantom with the FBP algorithm as in (a), but with the additional use of a Butterworth lowpass filter. See also Figure 9.5. Reproduced, with permission, from R.M. Rangayyan and A. Kantzas, "Image reconstruction", *Wiley Encyclopedia of Electrical and Electronics Engineering*, Supplement 1, Editor: J. G. Webster, Wiley, New York, NY, pp 249–268, 2000. © This material is used by permission of John Wiley & Sons, Inc.

The Radon transform may be interpreted as a transformation of the given image from the (x, y) space to the (θ, t) space. In practical CT scanning, the projection or ray-integral data are obtained as samples at discrete intervals in t and θ . Just as we encounter the (Nyquist or Shannon) sampling theorem in the representation of a 1D signal in terms of its samples in time, we now encounter the requirement to sample adequately along both the t and θ axes. A major distinction lies in the fact that the measurements made in CT scanners are discrete to begin with, and the signal (the body or object being imaged) cannot be prefiltered to prevent aliasing. Undersampling in either axis will lead to aliasing errors and poor reconstructed images.

Figure 9.9 (a) shows the reconstructed version of the phantom in Figure 9.3 (a) obtained using only 10 projections spanning the $0^\circ - 180^\circ$ range in sampling steps of 18° and using the FBP algorithm. Although the edges of the objects in the image are sharper than those in the reconstruction obtained using the BP algorithm with the same parameters [see Figure 9.5 (c)], the image is affected by severe streaking artifacts [753] due to the limited number of projections used. Figure 9.9 (b) shows the reconstructed image of the phantom obtained using 10 projections but spanning only the angular range of $40^\circ - 130^\circ$ in steps of 10° . The limited angular coverage provided by the projections has clearly affected the quality of the image, and has introduced geometric distortion [753, 746, 748, 744, 747].

9.4 Algebraic Reconstruction Techniques

The algebraic reconstruction technique (ART) [11, 742, 743] is related to the Kaczmarz method [754] of projections for solving simultaneous equations. The Kaczmarz method takes an approach that is completely different from that of the Fourier or FBP methods: the available projections — treated as individual ray sums in a discrete representation — are seen as a set of simultaneous equations, with the unknown quantities being the discrete pixels of the image. The large size of images encountered in practice precludes the use of the usual methods for solving simultaneous equations. Furthermore, in many practical applications, the number of available equations may be far less than the number of pixels in the image to be reconstructed; the set of simultaneous equations is then under-determined. The Kaczmarz method of projections is an elegant iterative method that may be implemented easily. (*Note:* The Kaczmarz method uses the term “projection” in the vectorial or geometric sense, and individual ray sums are processed one at a time. Observe that a set of ray sums or integrals is also known as a projection. The distinction should be clear from the context.)

**FIGURE 9.9**

(a) Reconstruction of the phantom in Figure 9.3 (a) obtained using 10 projections from 18° to 180° in steps of 18° with the FBP algorithm. (b) Reconstruction of the phantom obtained using 10 projections from 40° to 130° in steps of 10° with the FBP algorithm. The ramp filter that is implicit in the FBP process was combined with a Butterworth lowpass filter in both cases. See also Figure 9.5. Reproduced, with permission, from R.M. Rangayyan and A. Kantzas, "Image reconstruction", *Wiley Encyclopedia of Electrical and Electronics Engineering*, Supplement 1, Editor: J. G. Webster, Wiley, New York, NY, pp 249–268, 2000. © This material is used by permission of John Wiley & Sons, Inc.

Let the image to be reconstructed be divided into N cells, f_n denoting the value in the n^{th} cell. The image density or intensity is assumed to be constant within each cell. Let M be the number of ray sums available, expressed as

$$p_m = \sum_{n=1}^N w_{mn} f_n, \quad m = 1, 2, \dots, M, \quad (9.33)$$

where w_{mn} is the contribution factor of the n^{th} image element to the m^{th} ray sum, equal to the fractional area of the n^{th} cell crossed by the m^{th} ray path, as illustrated in Figure 9.10. (*Note:* An image and its Radon transform are each represented using only one index in this formulation of ART.) Observe that for a given ray m , most of w_{mn} will be zero, because only a few elements of the image contribute to the corresponding ray sum. Equation 9.33 may also be expressed as

$$\begin{aligned} w_{11}f_1 + w_{12}f_2 + \dots + w_{1N}f_N &= p_1, \\ w_{21}f_1 + w_{22}f_2 + \dots + w_{2N}f_N &= p_2, \\ \vdots & \\ w_{M1}f_1 + w_{M2}f_2 + \dots + w_{MN}f_N &= p_M. \end{aligned} \quad (9.34)$$

A grid representation with N cells gives the image N degrees of freedom. Thus, an image represented by $\mathbf{f} = [f_1, f_2, \dots, f_N]^T$ may be considered to be a single point in an N -dimensional hyperspace. Then, each of the above ray-sum equations will represent a hyperplane in this hyperspace. If a unique solution exists, it is given by the intersection of all the hyperplanes at a single point. To arrive at the solution, the Kaczmarz method takes the approach of successively and iteratively projecting an initial guess and its successors from one hyperplane to the next.


Let us, for simplicity, consider a 2D version of the situation (with $N = M = 2$), as illustrated in Figures 9.11 and 9.12. Let $\mathbf{f}^{(0)}$ represent vectorially the initial guess to the solution, and let $\mathbf{w}_1 = [w_{11}, w_{12}]^T$ represent vectorially the series of weights (coefficients) in the first ray equation. The first ray sum may then be written as

$$\mathbf{w}_1 \cdot \mathbf{f} = p_1. \quad (9.35)$$

The hyperplane represented by this equation is orthogonal to \mathbf{w}_1 . (Consider two images or points \mathbf{f}_1 and \mathbf{f}_2 belonging to the hyperplane. We have $\mathbf{w}_1 \cdot \mathbf{f}_1 = p_1$ and $\mathbf{w}_1 \cdot \mathbf{f}_2 = p_1$. Hence, $\mathbf{w}_1 \cdot [\mathbf{f}_1 - \mathbf{f}_2] = 0$. Therefore, \mathbf{w}_1 is orthogonal to the hyperplane.)

With reference to Figure 9.12, Equation 9.35 indicates that, for the vector \mathbf{OC} corresponding to any point C on the hyperplane, its projection on to the vector \mathbf{w}_1 is of a constant length. The unit vector \mathbf{OU} along \mathbf{w}_1 is given by

$$\mathbf{OU} = \frac{\mathbf{w}_1}{\sqrt{\mathbf{w}_1 \cdot \mathbf{w}_1}}. \quad (9.36)$$

**FIGURE 9.10**

ART treats the image as a matrix of discrete pixels of finite size $(\Delta x, \Delta y)$. Each ray has a finite width. The fraction of the area of the n^{th} pixel crossed by the m^{th} ray is represented by the weighting factor $w_{mn} = \text{area of ABCD}/(\Delta x \Delta y)$ for the n^{th} pixel f_n in the figure. Adapted, with permission, from A. Rosenfeld and A.C. Kak, *Digital Picture Processing*, 2nd ed., New York, NY, 1982. © Academic Press.

The perpendicular distance of the hyperplane from the origin is

$$\|\mathbf{OA}\| = \mathbf{OU} \cdot \mathbf{OC} = \frac{\mathbf{w}_1 \cdot \mathbf{f}}{\sqrt{\mathbf{w}_1 \cdot \mathbf{w}_1}} = \frac{p_1}{\sqrt{\mathbf{w}_1 \cdot \mathbf{w}_1}}. \quad (9.37)$$

Now,

$$\mathbf{f}^{(1)} = \mathbf{f}^{(0)} - \mathbf{GH}, \quad (9.38)$$

and

$$\begin{aligned} \|\mathbf{GH}\| &= \|\mathbf{OF}\| - \|\mathbf{OA}\| = \mathbf{f}^{(0)} \cdot \mathbf{OU} - \|\mathbf{OA}\| \\ &= \frac{\mathbf{f}^{(0)} \cdot \mathbf{w}_1}{\sqrt{\mathbf{w}_1 \cdot \mathbf{w}_1}} - \frac{p_1}{\sqrt{\mathbf{w}_1 \cdot \mathbf{w}_1}} = \frac{\mathbf{f}^{(0)} \cdot \mathbf{w}_1 - p_1}{\sqrt{\mathbf{w}_1 \cdot \mathbf{w}_1}}. \end{aligned} \quad (9.39)$$

Because the directions of \mathbf{GH} and \mathbf{OU} are the same, $\mathbf{GH} = \|\mathbf{GH}\| \mathbf{OU}$. Thus,

$$\mathbf{GH} = \left(\frac{\mathbf{f}^{(0)} \cdot \mathbf{w}_1 - p_1}{\sqrt{\mathbf{w}_1 \cdot \mathbf{w}_1}} \right) \frac{\mathbf{w}_1}{\sqrt{\mathbf{w}_1 \cdot \mathbf{w}_1}} = \left(\frac{\mathbf{f}^{(0)} \cdot \mathbf{w}_1 - p_1}{\mathbf{w}_1 \cdot \mathbf{w}_1} \right) \mathbf{w}_1. \quad (9.40)$$

Therefore,

$$\mathbf{f}^{(1)} = \mathbf{f}^{(0)} - \left(\frac{\mathbf{f}^{(0)} \cdot \mathbf{w}_1 - p_1}{\mathbf{w}_1 \cdot \mathbf{w}_1} \right) \mathbf{w}_1. \quad (9.41)$$

In general, the m^{th} estimate is obtained from the $(m-1)^{\text{th}}$ estimate as

$$\mathbf{f}^{(m)} = \mathbf{f}^{(m-1)} - \left(\frac{\mathbf{f}^{(m-1)} \cdot \mathbf{w}_m - p_m}{\mathbf{w}_m \cdot \mathbf{w}_m} \right) \mathbf{w}_m. \quad (9.42)$$

That is, the $(m-1)^{\text{th}}$ estimate on hand is projected on to the hyperplane of the m^{th} ray sum, and the deviation from the true ray sum p_m is obtained. This deviation is normalized and applied as a correction to all the pixels according to the weighting factors in \mathbf{w}_m . When this process is applied to all the M ray-sum hyperplanes given, one cycle or iteration is completed. (*Note:* Because the image is updated by altering the pixels along each individual ray sum, the index of the updated estimate or of the iteration is equal to the index of the latest ray sum used. However, as the entire process is iterated, the index of the estimate is reset at the beginning of each iteration.)

Depending upon the initial guess and the organization of the hyperplanes, a number of iterations may have to be completed in order to obtain the solution (if it exists). The following important characteristics of ART are worth observing:

- ART proceeds ray by ray, and is iterative.




FIGURE 9.11

Illustration of the Kaczmarz method of solving a pair of simultaneous equations in two unknowns. The solution is $\mathbf{f} = [3, 4]^T$. The weight vectors for the two ray sums (straight lines) are $\mathbf{w}_1 = [2, -1]^T$ and $\mathbf{w}_2 = [1, 1]^T$. The equations of the straight lines are $\mathbf{w}_1 \cdot \mathbf{f} = 2f_1 - f_2 = 2 = p_1$ and $\mathbf{w}_2 \cdot \mathbf{f} = f_1 + f_2 = 7 = p_2$. The initial estimate is $\mathbf{f}^{(0)} = [4, 1]^T$. The first updated estimate is $\mathbf{f}^{(1)} = [2, 2]^T$; the second updated estimate is $\mathbf{f}^{(2)} = [3.5, 3.5]^T$. Because two ray sums are given, two corrections constitute one cycle (or iteration) of ART. The path of the second cycle of ART is also illustrated in the figure. Reproduced, with permission, from R.M. Rangayyan and A. Kantzas, "Image reconstruction", *Wiley Encyclopedia of Electrical and Electronics Engineering*, Supplement 1, Editor: J. G. Webster, Wiley, New York, NY, pp 249–268, 2000. © This material is used by permission of John Wiley & Sons, Inc.

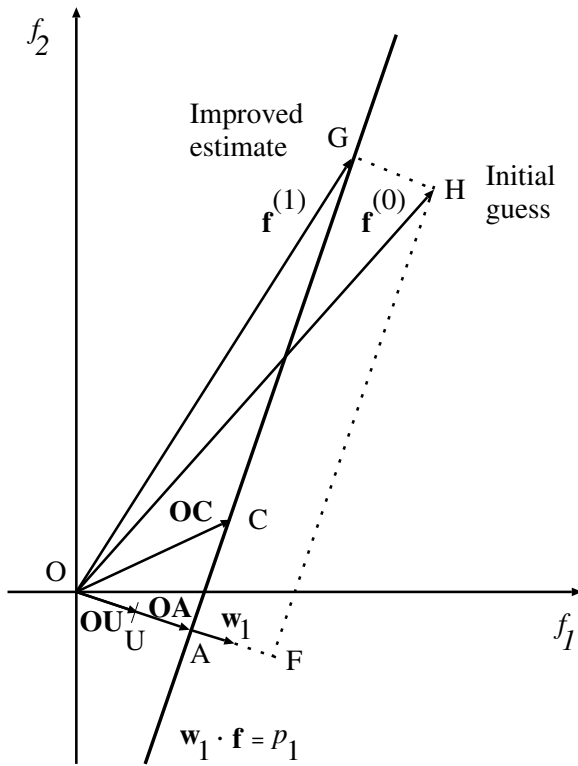


FIGURE 9.12

Illustration of the algebraic reconstruction technique. $\mathbf{f}^{(1)}$ is an improved estimate computed by projecting the initial guess $\mathbf{f}^{(0)}$ on to the hyperplane (the straight line AG in the illustration) corresponding to the first ray sum given by the equation $\mathbf{w}_1 \cdot \mathbf{f} = p_1$. Adapted, with permission, from A. Rosenfeld and A.C. Kak, *Digital Picture Processing*, 2nd ed., New York, NY, 1982.

© Academic Press.

- If the hyperplanes of all the given ray sums are mutually orthogonal, we may start with any initial guess and reach the solution in only one cycle.

On the other hand, if the hyperplanes subtend small angles with one another, a large number of iterations will be required. The number of iterations may be reduced by using optimized ray-access schemes [755].

- If the number of ray sums is greater than the number of pixels, that is, $M \geq N$, but the measurements are noisy, no unique solution exists — the procedure will oscillate in the neighborhood of the intersections of the hyperplanes.
- If $M < N$, the system is under-determined and an indefinite or infinite number of partial solutions exist. It has been shown that unconstrained ART converges to the minimum-variance estimate [752].
- The major advantage of ART is that any *a priori* information available about the image may be introduced easily into the iterative procedure (for example, upper and/or lower limits on pixel values, and the spatial boundaries of the image). This may help in obtaining a useful “solution” even if the system is under-determined.

9.4.1 Approximations to the Kaczmarz method

We could rewrite the reconstruction step in Equation 9.42 at the n^{th} pixel level as

$$f_n^{(m)} = f_n^{(m-1)} + \left[\frac{p_m - q_m}{\sum_{k=1}^N w_{mk}^2} \right] w_{mn}, \quad (9.43)$$

where $q_m = \mathbf{f}^{(m-1)} \cdot \mathbf{w}_m = \sum_{k=1}^N f_k^{(m-1)} w_{mk}$. This equation indicates that when we project the $(m-1)^{\text{th}}$ estimate on to the m^{th} hyperplane, the correction factor for the n^{th} cell is

$$\Delta f_n^{(m)} = f_n^{(m)} - f_n^{(m-1)} = \left[\frac{p_m - q_m}{\sum_{k=1}^N w_{mk}^2} \right] w_{mn}. \quad (9.44)$$

Here, p_m is the given (true) ray sum for the m^{th} ray, and q_m is the computed ray sum for the same ray for the estimated image on hand. $(p_m - q_m)$ is the error in the estimate, which is normalized and applied as a correction to all the pixels with appropriate weighting. Because the correction factor is added to the current image, this version of ART is known as additive ART.

In one of the approximations to Equation 9.43, the weights w_{mn} are simply replaced by zeros or ones depending upon whether the center of the n^{th} image cell is within the m^{th} ray (of finite width) or not [742, 743]. Then, the coefficients need not be computed and stored — we may instead determine the

pixels to be corrected for the ray considered during the reconstruction procedure. Furthermore, it follows that $\sum_{k=1}^N w_{mk}^2 = N_m$, the number of pixels crossed by the m^{th} ray. The correction applicable to *all of the pixels* along the m^{th} ray is $(p_m - q_m)/N_m$. Then,

$$f_n^{(m)} = f_n^{(m-1)} + \frac{p_m - q_m}{N_m}. \quad (9.45)$$

Because the corrections could be negative, negative pixel values may be encountered. Because negative values are not meaningful in most imaging applications, the constrained (and thereby nonlinear) version of ART is defined as

$$f_n^{(m)} = \max \left[0, f_n^{(m-1)} + \frac{p_m - q_m}{N_m} \right]. \quad (9.46)$$

The corrections could also be multiplicative [75]:

$$f_n^{(m)} = f_n^{(m-1)} \frac{p_m}{q_m}. \quad (9.47)$$

This version of ART is known as multiplicative ART. In this case, no positivity constraint is required. Furthermore, the convex hull of the image is almost guaranteed (subject to approximation related to the number of ray sums available and their angular coverage), because a pixel once set to zero will remain so during subsequent iterations. It has been shown that the multiplicative version of ART converges to the maximum-entropy estimate of the image [43, 756].

A generic ART procedure may be expressed in the following algorithmic form:

1. Prepare an initial estimate $\mathbf{f}^{(0)}$ of the image. All of the pixels in the initial image could be zero for additive ART; however, for multiplicative ART, pixels within at least the convex hull of the object in the image must have values other than zero.
2. Compute the ray sum q_m for the first ray path ($m = 1$) for the estimate of the image on hand.
3. Obtain the difference between the true ray sum p_m and the computed ray sum q_m , and apply the correction to all the pixels belonging to the ray according to one of the ART equations (for example, Equation 9.43, 9.45, 9.46, or 9.47). Apply constraints, if any, based upon the *a priori* information available.
4. Perform Steps 2 and 3 for all rays available, $m = 1, 2, \dots, M$.
5. Steps 2 – 4 constitute one cycle or iteration (over all available ray sums). Repeat Steps 2 – 4 as many times as required. If desired, compute a

measure of convergence, such as

$$E_1 = \sum_{m=1}^M (p_m - q_m)^2 \quad (9.48)$$

or


$$E_2 = \sum_{n=1}^N \left(f_n^{(m)} - f_n^{(m-1)} \right)^2. \quad (9.49)$$

Stop if the error or difference is less than a prespecified limit; else, go back to Step 2.

For improved convergence, a simultaneous correction procedure (Simultaneous Iterative Reconstruction Technique — SIRT [757]) has been proposed, where the corrections to all the pixels from all the rays are first computed, and the averaged corrections are applied at the same time to all the pixels (that is, only one correction is applied per pixel per iteration). Guan and Gordon [755] proposed different ray-access schemes to improve convergence, including the consecutive use of rays in mutually orthogonal directions. In the absence of complete projection data spanning the full angular range of 0° to 180° , ART typically yields better results than the FBP or the Fourier methods.

Examples of reconstructed images: Figure 9.13 (b) shows the reconstruction of the phantom shown in part (a) of the same figure, obtained using 90 parallel-ray projections with three iterations of constrained additive ART, as in Equation 9.46. Ray sums for use with ART were computed from the phantom image data using angle-dependent ray width, given by $\max(|\sin(\theta)|, |\cos(\theta)|)$ [742, 753]. The ART reconstruction is better than that given by the BP or the FBP algorithm. The advantages of constrained ART due to the use of the positivity constraint — that is, the *a priori* knowledge imposed — and of the ability to iterate, are seen in the improved quality of the result.

Figure 9.14 shows reconstructed images of the phantom obtained using only 10 projections spanning the angular ranges of (a) 18° to 180° in steps of 18° , and (b) 40° – 130° in steps of 10° , respectively. The limited number of projections used and the limited angular coverage of the projections in the second case have affected the quality of the reconstructed images and introduced geometric distortion. However, when compared with the results of BP and FBP with similar parameters (see Figures 9.5 and 9.9), ART has provided better results.

**FIGURE 9.13**

(a) A synthetic 2D image (phantom) with 101×101 eight-bit pixels, representing a cross-section of a 3D object [the same image as in Figure 9.3 (a)]. (b) Reconstruction of the phantom obtained using 90 projections from 2° to 180° in steps of 2° with three iterations of constrained additive ART. See also Figure 9.8. Reproduced, with permission, from R.M. Rangayyan and A. Kantzas, "Image reconstruction", *Wiley Encyclopedia of Electrical and Electronics Engineering*, Supplement 1, Editor: J. G. Webster, Wiley, New York, NY, pp 249–268, 2000. © This material is used by permission of John Wiley & Sons, Inc.




FIGURE 9.14

Reconstruction of the phantom in Figure 9.13 (a) obtained using: (a) 10 projections from 18° to 180° in steps of 18° with three iterations of constrained additive ART; (b) 10 projections from 40° to 130° in steps of 10° with three iterations of constrained additive ART. See also Figures 9.5 and 9.9. Reproduced, with permission, from R.M. Rangayyan and A. Kantzas, "Image reconstruction", *Wiley Encyclopedia of Electrical and Electronics Engineering*, Supplement 1, Editor: J. G. Webster, Wiley, New York, NY, pp 249–268, 2000. © This material is used by permission of John Wiley & Sons, Inc.

9.5 Imaging with Diffracting Sources

In some applications of CT imaging, such as imaging pregnant women, X-ray imaging may not be advisable. Imaging with nonionizing forms of radiation, such as acoustic (ultrasonic) [82, 83] and electromagnetic (optical or thermal) imaging [85], is then a valuable alternative. X-ray imaging is also not suitable when the object to be imaged has poor contrast in density or atomic number distribution. An important point to observe in acoustic or electromagnetic imaging is that these forms of energy do not propagate along straight-line ray paths through a body due to refraction and diffraction. When the dimensions of the inhomogeneities in the object being imaged are comparable to or smaller than the wavelength of the radiation used, geometric propagation concepts cannot be applied; it becomes necessary to consider wave propagation and diffraction-based methods.

When the body being imaged may be treated as a weakly scattering object in the 2D sectional plane and invariant in the axial direction, the Fourier diffraction theorem is applicable [82]. This theorem states that the 1D Fourier transform of a projection including the effects of diffraction gives values of the 2D Fourier transform of the image along a semicircular arc. Interpolation methods may be developed in the Fourier space taking this property into account for reconstruction of images from projections obtained with diffracting sources. Backpropagation and algebraic techniques have also been proposed for the case of imaging with diffracting sources [82].

9.6 Display of CT Images

X-ray CT is capable of producing images with high density resolution, on the order of one part in 1,000. For display purposes, the attenuation coefficients are normalized with respect to that of water and expressed as

$$HU = K \left(\frac{\mu}{\mu_w} - 1 \right), \quad (9.50)$$

where μ is the measured attenuation coefficient, and μ_w is the attenuation coefficient of water. The K parameter used to be set at 500 in early models of the CT scanner. It is now common to use $K = 1,000$ to obtain the CT number in Hounsfield units (HU) [758], named after the inventor of the first commercial medical CT scanner [77]. This scale results in values of about +1,000 for bone, 0 for water, about -1,000 for air, -80 to 20 for soft tissue, and about -800 for lung tissue [38]. Table 9.1 shows the mean and SD of the CT values in HU for several types of tissue in the abdomen.

TABLE 9.1
 Mean and SD of CT Values in
 Hounsfield Units (*HU*) for a Few
 Types of Abdominal Tissue.

Tissue	Mean <i>HU</i>	SD
Air ²	-1,006	2
Fat ¹	-90	18
Bile ¹	+16	8
Kidney ¹	+32	10
Pancreas ¹	+40	14
Blood (aorta) ¹	+42	18
Muscle ¹	+44	14
Necrosis ²	+45	15
Spleen ¹	+46	12
Liver ¹	+60	14
Viable tumor ²	+91	25
Marrow ¹	+142	48
Calcification ²	+345	155
Bone ²	+1,005	103

¹Based upon Mategrano et al. [759]. ²Estimated from CT exams with contrast (see Section 9.9), based upon 1,000 – 4,000 pixels in each category. The contrast medium is expected to increase the CT values of vascularized tissues by 30 – 40 *HU*. The CT number for air should be –1000; the estimated value is slightly different due to noise in the images. Reproduced with permission from F.J. Ayres, M.K. Zuffo, R.M. Rangayyan, G.S. Boag, V. Odone Filho, and M. Valente, “Estimation of the tissue composition of the tumor mass in neuroblastoma using segmented CT images”, *Medical and Biological Engineering and Computing*, 42:366 – 377, 2004. © IFMBE.

The dynamic range of CT values is much wider than those of common display devices and that of the human visual system at a given level of adaptation. Furthermore, clinical diagnosis requires detailed visualization and analysis of small density differences. For these reasons, the presentation of the entire range of values available in a CT image in a single display is neither practically feasible nor desirable. In practice, small “windows” of the CT number scale are selected and linearly expanded to occupy the capacity of the display device. The window width and level (center) values may be chosen interactively to display different density ranges with improved perceptibility of details within the chosen density window. Values above or below the window limits are displayed as totally white or black, respectively. This technique, known as windowing or density slicing, may be expressed as

$$g(x, y) = \begin{cases} 0 & \text{if } f(x, y) \leq m \\ \frac{N}{(M-m)} [f(x, y) - m] & \text{if } m < f(x, y) < M \\ N & \text{if } f(x, y) \geq M \end{cases}, \quad (9.51)$$

where $f(x, y)$ is the original image in CT numbers, $g(x, y)$ is the windowed image to be displayed, $[m, M]$ is the range of CT values in the window to be displayed, and $[0, N]$ is the range of the display values. The window width is $[M - m]$ and the window level (or center) is $(M + m)/2$; the display range is typically $[0, 255]$ with 8-bit display systems.

Example: Figure 9.15 shows a set of two CT images of a patient with head injury, with each image displayed using two sets of window level and width. The effects of the density window chosen on the features of the image displayed are clearly seen in the figure: either the fractured bone or the brain matter are seen in detail in the windowed images, but not both in the same image. See Figures 1.24 and 4.4 for more examples of density windowing in the display of CT images. See also Figures 1.22, 1.23, and 2.15, as well as Section 9.9 for more examples of X-ray CT images. Examples of images reconstructed from projection data from other modalities of medical imaging such as MRI, SPECT, and PET are provided in Sections 1.7 and 1.9.

A dramatic visualization of details may be achieved with pseudo-color techniques. Arbitrary or structured color scales could be assigned to CT values by LUTs or gray-scale-to-color transformations. Some of the popular color transforms are the rainbow (VIBGYOR: violet – indigo – blue – green – yellow – orange – red) and the heated metal color (black – red – yellow – white) sequences. Difficulties may arise, however, in associating density values with different colors if the transformation is arbitrary and not monotonic in intensity or total brightness. Furthermore, small changes in CT values could cause abrupt changes in the corresponding colors displayed, especially with a mapping such as VIBGYOR. An LUT linking the displayed colors to CT numbers or other pixel attributes may assist in improved visual analysis of image features in engineering and scientific applications.




FIGURE 9.15

Two CT images of a patient with head injury, with each image displayed with two sets of window level and width. Images (a) and (b) are of the same section; images (c) and (d) are of another section. The window levels used to obtain images (a) – (d) are 18, 138, 22, and 138 *HU*, respectively; the window widths used are 75, 400, 75, and 400 *HU*, respectively. The windows in (b) and (d) display the skull, the fracture, and the bone segments, but the brain matter is not visible; the windows (a) and (c) display the brain matter in detail, but the fracture area is saturated. Images courtesy of W. Gordon, Health Sciences Centre, Winnipeg, MB, Canada.

9.7 Agricultural and Forestry Applications

Cruvinel et al. [760] and Vaz et al. [761] developed a portable X-ray and gamma-ray minitomograph for application in soil science, and used the scanner to measure water content and bulk density of soil samples. Soil-related studies address the identification of features such as fractures, wormholes, and roots, and assist in studies of flow of various contaminants in soil.

Forestry applications of CT have appeared in the literature in the form of scanning of live trees to measure growth rings and detect decay using a portable X-ray CT scanner [762], and monitoring tree trunks or logs in the timber and lumber industry. Figures 9.16 and 9.17 show a portable CT scanner in operation and images of a utility pole.




FIGURE 9.16

A portable CT scanner to image live trees. Reproduced with permission from A.M. Onoe, J.W. Tsao, H. Yamada, H. Nakamura, J. Kogure, H. Kawamura, and M. Yoshimatsu, "Computed tomography for measuring annual rings of a live tree", *Proceedings of the IEEE*, 71(7):907–908, 1983. © IEEE.




FIGURE 9.17

CT images of a Douglas fir utility pole and photographs of the corresponding physical sections. The sections in (a) demonstrate normal annual growth rings. The sections in (b) indicate severe decay in the heartwood region. Reproduced with permission from A.M. Onoe, J.W. Tsao, H. Yamada, H. Nakamura, J. Kogure, H. Kawamura, and M. Yoshimatsu, "Computed tomography for measuring annual rings of a live tree", *Proceedings of the IEEE*, 71(7):907-908, 1983. © IEEE.

9.8 Microtomography

The resolution of common CT devices used in medical and other applications varies from the common figure of $1 \times 1 \text{ mm}$ to about $200 \times 200 \mu\text{m}$ in cross-section, and $1\text{--}5 \text{ mm}$ between slices. Special systems have been built to image small samples of the order of 1 cm^3 in volume with resolution of the order of $5\text{--}10 \mu\text{m}$ in cross-section [134, 763, 764, 765, 766, 767, 768, 769, 770, 771, 772]. (See Section 2.12 for illustrations of the LSF and MTF of a μCT system.) Such an imaging procedure is called microtomography, microCT, or μCT , being a hybrid of tomography and microscopy. Most μCT studies are performed with finely focused and nearly monochromatic X-ray beams produced by a particle accelerator (such as a synchrotron). Stock [763] provides a review of the basic principles, techniques, and applications of μCT . Whereas most μCT studies have been limited to small, excised samples, Sasov [768] discusses the design of a μCT to image whole, small animals with resolution of the order of $10 \mu\text{m}$.

Umetani et al. [767] present the application of synchrotron-based μCT in a microangiography mode to the study of microcirculation within the heart of small animals. Shimizu et al. [766] studied the effects of alveolar hemorrhage and alveolitis (pneumonia) on the microstructure of the lung using synchrotron-based μCT . Johnson et al. [764] developed a microfocal X-ray imaging system to image the arterial structure in the rat lung, and studied the decrease in distensibility due to pulmonary hypertension.

Shaler et al. [769] applied μCT to study the fiber orientation and void structure in wood, paper, and wood composites. Illman and Dowd [765] studied the xylem tissue structure of wood samples, and analyzed the loss of structural integrity due to fungal degradation.

Example of application to the analysis of bone structure: Injury to the anterior cruciate ligament is expected to lead to a decrease in bone mineral density. Post-traumatic osteoarthritis is known to cause joint space narrowing and full-thickness cartilage erosion. It has been established that the cancellous architecture of bone is related to its mechanical properties and strength [134, 771, 772].

Conventional studies on bone structure have been performed through histological analysis of thin slices of bone samples. Spatial resolution of the order of $0.3 \mu\text{m}$ can be realized with optical microscopy or microradiography. However, in addition to being destructive, this procedure could cause artifacts due to the slicing operation. Furthermore, limitations exist in the 3D information derived from a few slices. Boyd et al. [771] applied μCT techniques to analyze the morphometric and anisotropic changes in periarticular cancellous bone due to ligament injury and osteoarthritis.

In the work of Boyd et al., anterior cruciate ligaments of dogs were transected by arthrotomy. After a specific recovery period, at euthanasia, the femora and tibia were extracted. Cylindrical bone cores of diameter 6 mm

and length 12 – 14 *mm* were obtained from the weight-bearing regions of the medial femoral condyles and the medial tibial plateaus. Contralateral bone core samples were also extracted and processed to serve as internal controls. The cores were scanned at a resolution of 34 μm , and 3D images with diameter of 165 voxels and length of up to 353 voxels were reconstructed. Morphometric parameters such as bone volume ratio, relative surface density, and trabecular thickness were estimated from the images. It was postulated that the anisotropy of the bone fabric or trabecular orientation is related to mechanical anisotropy and loading conditions.

Figure 9.18 shows two sample bone-core sectional images each of a case of transected anterior cruciate ligament of a dog after 12 weeks of recovery, and the corresponding contralateral control sample; Figure 9.19 shows 3D renditions of the same samples. The bone core related to ligament transection demonstrates increased trabecular spacing, and hence lower bone density, than the contralateral sample. Boyd et al. observed that significant periarticular bone changes occur as early as three weeks after ligament transection; the changes were more pronounced 12 weeks after transection.




FIGURE 9.18

(a) Two sample sectional images of a bone core sample in a case of anterior cruciate ligament transection after 12 weeks of recovery. (b) Two sample sectional images of the corresponding contralateral bone core sample. The diameter of each sample is 6 *mm*. See also Figure 9.19. Images courtesy of S.K. Boyd, University of Calgary [772].

Example of application to the study of microcirculation in the heart: Umetani et al. [767] developed a μCT system using monochromatized synchrotron radiation for use as a microangiography tool to study circulatory disorders and early-stage malignant tumors. Two types of detection systems were used: an indirect system including a fluorescent screen, optical coupling, and a CCD camera; and a direct system with a beryllium faceplate, a photoconductive layer, and an electron-beam scanner.




FIGURE 9.19

3D renditions of μ CT reconstructions of a bone core sample in a case of anterior cruciate ligament transection after 12 weeks of recovery (left), and the corresponding contralateral bone core sample (right). The diameter of each sample is 6 mm. See also Figure 9.18. Images courtesy of S.K. Boyd, University of Calgary [772].

In one of the experiments of Umetani et al., the left side of the heart of a rat was fixed in formalin after barium sulphate was injected into the coronary artery. Figure 9.20 (a) shows a projection image (a microradiograph) of the specimen obtained at a resolution of $24 \mu\text{m}$. The image clearly shows the left-anterior-descending coronary artery. Figure 9.20 (b) shows a 3D visualization of a part of the specimen (of diameter 3.5 mm and height 5 mm) reconstructed at a resolution of $6 \mu\text{m}$. The 3D structure of small blood vessels with diameter of the order of $30 - 40 \mu\text{m}$ was visualized by this method. Visualization of tumor-induced small vessels that feed lesions was found to be useful in the diagnosis of malignant tumors.


9.9 Application: Analysis of the Tumor in Neuroblastoma

9.9.1 Neuroblastoma

Neuroblastoma is a malignant tumor of neural-crest origin that may arise anywhere along the sympathetic ganglia or within the adrenal medulla [773, 774]. There are three types of ganglion cell lesions that form a spectrum of neoplastic disease. Neuroblastoma is the most immature and malignant form of the three, usually presenting before the age of five years. Ganglioneuroblastoma is a more mature form that retains some malignant characteristics, with peak incidence between five and 10 years of age. Ganglioneuroma is well differentiated and benign, typically presenting after 10 years of age [775].

Neuroblastoma is the most common extra-cranial solid malignant tumor in children [776]; it is the third most common malignancy of childhood. It accounts for $8 - 10\%$ of all childhood cancers [777], and 15% of all deaths related to cancer in the pediatric age group. The median age at diagnosis is two years, and 90% of the diagnosed cases are in children under the age of five years [776].

In the US, about 650 children and adolescents younger than 20 years of age are diagnosed with neuroblastoma every year [778]. Neuroblastoma is the most common cancer of infancy, with an incidence rate that is almost double that of leukemia, the next most common malignancy occurring during the first year of life [778]. The rate of incidence of neuroblastoma among infants in the US has increased from 53 per million in the period 1976 – 84 to 74 per million in the period 1986 – 94 [778]. Gurney et al. [779] estimated an annual increase rate of 3.4% for extracranial neuroblastoma. Although some countries instituted screening programs to detect neuroblastoma in infants, studies have indicated that the possible benefit of screening on mortality is small and has not yet been demonstrated in reliable data [780, 781].



(a)



(b)

FIGURE 9.20

(a) Projection image of a rat heart specimen obtained using synchrotron radiation. (b) 3D μ CT image of the rat heart specimen. Reproduced with permission from K. Umetani, N. Yagi, Y. Suzuki, Y. Ogasawara, F. Kajiya, T. Matsumoto, H. Tachibana, M. Goto, T. Yamashita, S. Imai, and Y. Kajihara, "Observation and analysis of microcirculation using high-spatial-resolution image detectors and synchrotron radiation", *Proceedings SPIE 3977: Medical Imaging 2000 - Physics of Medical Imaging*, pp 522–533, 2000, © SPIE.

Sixty-five percent of the tumors related to neuroblastoma are located in the abdomen; approximately two-thirds of these arise in the adrenal gland. Fifteen percent of neuroblastoma are thoracic, usually located in the sympathetic ganglia of the posterior mediastinum. Ten to twelve percent of neuroblastoma are disseminated without a known site of origin [782].

Staging and prognosis: The most recent staging system for neuroblastoma is the International Neuroblastoma Staging System, which takes into account radiologic findings, surgical resectability, lymph-node involvement, and bone-marrow involvement [783]. The staging ranges from Stage 1 for a localized tumor with no lymph-node involvement, to Stage 4 with the disease spread to distant lymph nodes, bone, liver, and other organs [783].

The main determinant factors for prognosis are the patient's age and the stage of the disease [776, 782]. The survival rate of patients diagnosed with neuroblastoma under the age of one year is 74%, whereas it is only 14% for patients over the age of three years [776]. Whereas the survival rate of patients diagnosed with Stage 1 neuroblastoma is in the range 95 – 100%, it is only 10 – 30% for patients diagnosed with Stage 4 of the disease [776].

The site of the primary tumor is also said to be of relevance in the overall prognosis. Tumors arising in the abdomen and pelvis have the worst prognosis, with adrenal tumors having the highest mortality. Thoracic neuroblastoma has a better overall survival rate (61%) than abdominal tumors (20%) [782].

Surgical resection of the primary tumor is recommended whenever possible. However, the primary tumor of a patient with advanced neuroblastoma (Stages 3 and 4) can be unresectable, if there is risk of damaging vital structures in the procedure, notably when the mass encases the aorta; see Figure 9.21. Treatment in these cases requires chemotherapy or radiotherapy for initial shrinkage of the mass, after which (delayed) surgical resection may be performed [773].

Radiological analysis: The radiology of neuroblastoma has been studied extensively over the past three decades, and several review articles concerning the radiological aspects of the disease have been published [776, 782, 784, 785, 786]. Radiological exams can be useful in the initial diagnosis, assessment of extension, staging, presurgical evaluation, treatment, and follow-up.

Several imaging modalities have been investigated in the context of neuroblastoma, including ultrasound, CT, MRI, bone and marrow scintigraphy, excretory urography, and chest radiography. CT and MRI are regarded to be the best modalities for the evaluation of tumor stage [776], resectability [787], and prognosis and follow-up [776]. CT and MRI exams are mandatory for the analysis of the primary tumor, and some investigators have reported that MRI could be more useful than CT in evaluating metastatic disease [788, 789, 790].

Comparing CT and MRI, the latter has a higher diagnostic accuracy (the highest among all imaging modalities) and is said to be more suitable for the demonstration of tumor spread and visualization of the tumor in relationship to neighboring blood vessels and vessels within the tumor [776], which are important factors in therapy and assessment of resectability. MRI also provides




FIGURE 9.21

CT image of a patient with Stage 3 neuroblastoma, with the mass outline drawn by a radiologist. The mass encases the aorta, which is the small, circular object just above the spine. Reproduced with permission from F.J. Ayres, M.K. Zuffo, R.M. Rangayyan, G.S. Boag, V. Odone Filho, and M. Valente, “Estimation of the tissue composition of the tumor mass in neuroblastoma using segmented CT images”, *Medical and Biological Engineering and Computing*, 42:366 – 377, 2004. © IFMBE.

better soft-tissue contrast than CT, and is more promising for the estimation of tissue composition [791, 792]. Another advantageous feature is that MRI uses nonionizing radiation. Nevertheless, CT is considered to be the modality of choice in several investigations, because it is comparable in usefulness to MRI in evaluating local disease [776], and is more cost-effective. CT is known to be effective in detecting calcifications (a finding in favor of neuroblastoma [784] and against Wilms tumor), detection and assessment of the extension of local disease, evaluation of lymph-node involvement, recurrent disease, and nonskeletal metastasis [782].

On CT exams, abdominal neuroblastoma is seen as a mass of soft tissue, commonly suprarenal or paravertebral, irregularly shaped, lobulated, extending from the flank toward the midline, and lacking a capsule. The mass tends to be inhomogeneous due to tumor necrosis intermixed with viable tumor, and contains calcifications in 85% of patients. Calcifications are usually dense, amorphous, and mottled in appearance. Sometimes, neuroblastoma presents areas of central necrosis, shown as low-attenuation areas, that are more apparent after contrast enhancement [773, 776, 782]. Figure 9.21 shows a CT image of a patient with Stage 3 neuroblastoma, with the mass outline drawn by a radiologist. The mass encases the aorta, which makes it unresectable.

Despite the proven usefulness of imaging techniques in the detection, delineation, and staging of the primary tumor, there is a need for improvement in the usage of these techniques for more accurate assessment of the local

disease that could lead to better treatment planning and follow-up. Foglia et al. [793] argued that the primary tumor status in advanced neuroblastoma cannot be assessed definitively by diagnostic imaging, due to errors in sensitivity and specificity as high as 38%, when assessing tumor viability by imaging methods and comparing it to findings in delayed surgery. They also reported that CT exams could not differentiate viable tumor from fibrotic tissue or nonviable tumor destroyed by previous chemotherapy.

Computer-aided image analysis could improve radiological analysis of neuroblastoma by offering more sophisticated, quantitative, accurate, and reproducible measures of information in the image data [251, 794, 795]. Nevertheless, few researchers have investigated the potential of CAD of neuroblastoma using diagnostic imaging; related published works are limited to tumor volume measurement using manually segmented CT slices (planimetry) [796, 797]. Ayres et al. [798, 799, 800] proposed methods for computer-aided quantitative analysis of the primary tumor mass, in patients with advanced neuroblastoma, using X-ray CT exams. Specifically, they proposed a methodology for the estimation of the tissue content of the mass via statistical modeling and analysis of manually segmented CT images. The results of the method were compared with the results of histological analysis of surgically resected masses.

9.9.2 Tissue characterization using CT

The linear attenuation coefficient μ of tissue is the physical entity that is measured in CT (see Equations 1.1 and 1.2). The linear attenuation coefficient varies with two material properties: density and elemental composition [758]. The value of μ has been measured and tabulated for several materials [121, 801], including human and animal tissues, at different X-ray energies (including measurements with multiple energies, such as dual-energy imaging [758, 802, 803]). However, it is not common to display CT images in terms of the linear attenuation coefficient (which is dependent on the energy used [121, 803]). Instead, normalized CT units that are more convenient and independent, to a certain extent, of the X-ray energy are used [38]; see Equation 9.50. Table 9.1 shows the mean and SD of the CT values in *HU* for several types of tissue in the abdomen.

Soon after the invention of the CT scanner, several researchers focused their attention on the problem of estimating the composition of body tissues from CT images. Alter [804] evaluated the clinical usefulness of the CT scanner and reported on the appearance on CT images of several organs, tissues, and diseases, and used their Hounsfield value in tissue characterization. Phelps et al. [121] and Rao and Gregg [801] described the linear attenuation coefficient of several tissues, and related it to the ideal CT unit as reported by Wilson [802] and Brooks [803]. Mategrano et al. [759] presented measures of the attenuation coefficient for tissues in the abdominal region; see Table 9.1.

A discussion of potential sources of error in the measurement of the attenuation coefficient is found in the work of Williams et al. [805]. Duerinckx and Macovski [806] discussed the nature of noise in CT images. Pullan et al. [807] worked on the characterization of regions in a CT image using statistical central moments (such as mean, SD, and skewness) obtained from histograms of CT values and a gradient measure (taken as a simplified measure of texture) within delimited regions, with application to brain tumors and lesions of the spleen and liver. Kramer et al. [808] reported on a work similar to that of Pullan et al. Latchaw et al. [809] opined that CT is nonspecific in the separation of solid tumors and cystic lesions in the brain.

Intravascular contrast is usually employed in abdominal CT studies. The contrast agent is injected rapidly into the venous system, with scanning commencing shortly after completion of the injection. The objective of this technique is to image the patient while the intravascular concentration of contrast is at its peak, and before redistribution of contrast into soft tissues occurs, thus maximizing the density difference between vascular structures and other body organs, and allowing assessment of regional blood flow. The intravascular concentration of contrast decreases initially due to dilution in the blood volume, then by redistribution throughout perfused tissues, and lastly by renal excretion. The effect of contrast on *HU* measurements is dependent on many factors, including: blood volume, body weight, contrast volume and injection rate, time elapsed since injection, and vascularity of the structure of interest. Some of the *HU* values listed in Table 9.1 were estimated using CT data with contrast.

Some authors have reported relative success in using CT units to differentiate between tissues that are visually similar in CT images [759, 805, 807, 808], while others have reported failure in the same task [793, 809]. Although the measurement of the linear attenuation coefficient of tissues *in vitro* can be performed with good precision [121, 801], several sources of error can degrade the performance of *in vivo* CT sensitometry [810], some of which are motion artifacts, noise, partial-volume effect, and spectral spread of the energy of the X ray.

9.9.3 Estimation of tissue composition from CT images

Ayres et al. [798, 799, 800] investigated quantitative analysis of the primary tumor mass, in patients with advanced neuroblastoma, using CT exams. Their methodology includes a statistical parametric model for the tissue composition of the tumor, and a method to estimate the parameters of the model. Segmentation of the tumor was performed manually by a radiologist. The histogram of the tumor mass was computed from the segmented regions over all applicable CT image slices (see Section 2.7). The statistical model employed is the Gaussian mixture model [700], and the algorithm for parameter estimation is the EM algorithm [811]; see also Section 8.9.2.

The complete methodology is shown schematically in Figure 9.22. The upper path shows the sequence of algorithmic procedures that lead to the estimation of the statistical model. The lower path consists of obtaining the histological information regarding the tumor from biopsy and delayed surgery [793].

Estimation of tissue composition: The tumor mass in neuroblastoma is inhomogeneous, due to intermixed necrosis and viable tumoral tissue, and sometimes presents central areas of necrosis, shown as low-attenuation regions inside the mass. Therefore, it is appropriate to develop a global description of the mass that could lead to the estimation of the fractional volume corresponding to each tissue type, rather than attempting to separate the mass into distinct regions.

Assume that the CT value for a voxel that arises from a given tissue inside the mass (benign mass, necrosis, malignant tumor, fibrotic tissue, etc.) is a Gaussian random variable. Then, the whole tumor mass may be modeled statistically as a mixture of Gaussian variables, known as the Gaussian mixture model [812]. Let x denote the CT attenuation value for a given voxel, and $\theta_i = (\mu_i, \sigma_i)$ be the set of parameters that describes the Gaussian PDF of the CT values of the i^{th} type of tissue. The PDF for x , given that x belongs to the i^{th} tissue type, is $p_i(x|\theta_i)$, and is represented as

$$p_i(x|\theta_i) = \frac{1}{\sqrt{2\pi} \sigma_i} \exp \left[-\frac{(x - \mu_i)^2}{2\sigma_i^2} \right]. \quad (9.52)$$

Let M be the true number of the different types of tissue in the given tumor mass (assumed to be known for the moment; estimation of the value of M will be described later.) Let α_i be the probability that a given voxel came from the i^{th} type of tissue in the tumor mass. By definition, $\sum_{i=1}^M \alpha_i = 1$. The value of α_i can be seen as the fraction of the tumor volume that is composed of the i^{th} type of tissue. Then, the PDF for the entire mass is a mixture of Gaussians, specified by the parameter vector $\Theta = (\alpha_1, \dots, \alpha_M, \theta_1, \dots, \theta_M)^T$, and described by $p(x|\Theta) = \sum_{i=1}^M \alpha_i p_i(x|\theta_i)$.

Let N be the number of voxels in the tumor mass, and let $\mathbf{x} = (x_1, x_2, \dots, x_N)^T$ be a vector composed of the values of the voxels (the observed data). Under the condition of the observed data, the *posterior* probability of the parameters is obtained by Bayes rule [700] (see Section 12.4.1) as

$$p(\Theta|\mathbf{x}) = \frac{p(\Theta) p(\mathbf{x}|\Theta)}{p(\mathbf{x})}. \quad (9.53)$$

Here, $p(\mathbf{x})$ is the probability of the data, regardless of the parameters; because the estimation problem is conditional upon the observed data, $p(\mathbf{x})$ is a constant term. The term $p(\Theta)$ is the *prior* probability of the vector of parameters Θ . The term $p(\mathbf{x}|\Theta)$ is called the *likelihood* of Θ , denoted by $L(\Theta|\mathbf{x})$, and given by

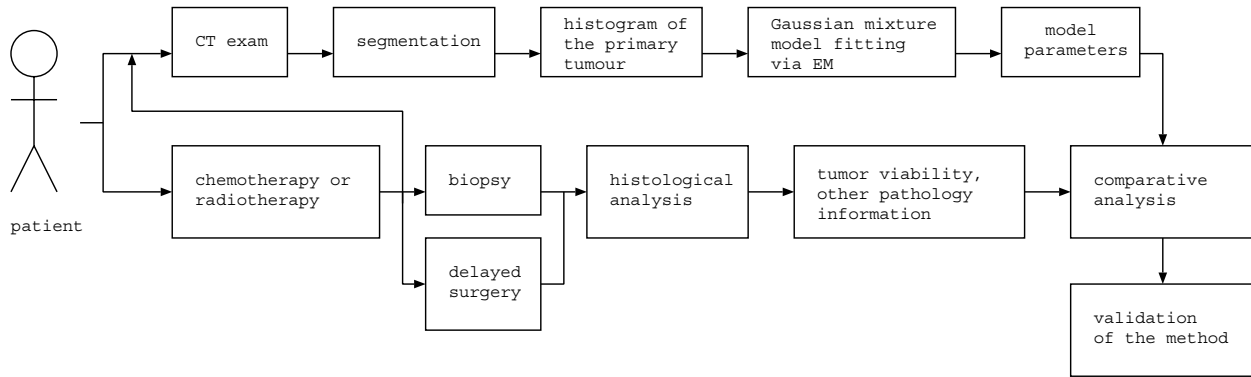


FIGURE 9.22

Schematic representation of the proposed method for the analysis of neuroblastoma. EM: Expectation-maximization. Reproduced with permission from F.J. Ayres, M.K. Zuffo, R.M. Rangayyan, G.S. Boag, V. Odone Filho, and M. Valente, “Estimation of the tissue composition of the tumor mass in neuroblastoma using segmented CT images”, *Medical and Biological Engineering and Computing*, 42:366 – 377, 2004. © IFMBE.

$$L(\Theta|\mathbf{x}) \triangleq p(\mathbf{x}|\Theta) = \prod_{j=1}^N p(x_j|\Theta) . \quad (9.54)$$

In the case that there is no prior belief about Θ , that is, nothing is known about its prior probability, the situation is known as a *flat prior* [812]. In this case, Equation 9.53 becomes $p(\Theta|\mathbf{x}) = c p(\mathbf{x}|\Theta)$, where c is a normalizing constant. Thus, finding the most probable value of Θ given the data, without any prior knowledge about the PDF of the parameters, is the same as finding the value of Θ that maximizes the likelihood, or the log-likelihood defined as $\log[L(\Theta|\mathbf{x})]$: this is the *maximum likelihood* (ML) principle [700, 812]. The adoption of this principle leads to simplified calculations with reasonable results [813]. Fully Bayesian approaches for classification and parameter estimation can provide better performance, at the expense of greater computational requirements and increased complexity of implementation [814, 815].

In order to maximize the likelihood, Ayres et al. used the EM algorithm [700, 811, 812, 813, 816]. The EM algorithm is an iterative procedure that starts with an initial guess Θ_g of the parameters, and iteratively improves the estimate toward the local maximum of the likelihood. The generic EM algorithm is comprised of two steps: the expectation step (or E-step) and the maximization step (or M-step). In the E-step, one computes the parametric probability model given the current estimate of the parameter vector. In the M-step, one finds the parameter vector that maximizes the newly calculated model, which is then treated as the new best estimate of the parameters. The iterative procedure continues until some stopping condition is met, for example, the difference $\log[L(\Theta_{n+1}|\mathbf{x})] - \log[L(\Theta_n|\mathbf{x})]$ or the modulus $|\Theta_{n+1} - \Theta_n|$ of the difference vector between successive iterations n and $n + 1$ is smaller than a predefined value.

For each tissue type i , let $p(i|x_j, \Theta)$ represent the probability that the j^{th} voxel, with the value x_j , belongs to the i^{th} tissue type. This can be calculated using Bayes rule as

$$p(i|x_j, \Theta) = \frac{p(i|\Theta) p(x_j|i, \Theta)}{p(x_j|\Theta)} = \frac{\alpha_i p_i(x_j|\theta_i)}{p(x_j|\Theta)} . \quad (9.55)$$

The derivation of the EM algorithm for the Gaussian mixture model leads to a set of iterative equations that perform the E-step and the M-step simultaneously. For the i^{th} tissue type, the update equations are:

$$\alpha_i^{\text{new}} = \frac{1}{N} \sum_{j=1}^N p(i|x_j, \Theta^{\text{old}}) , \quad (9.56)$$

$$\mu_i^{\text{new}} = \frac{\sum_{j=1}^N x_j p(i|x_j, \Theta^{\text{old}})}{\sum_{j=1}^N p(i|x_j, \Theta^{\text{old}})} , \quad (9.57)$$

$$\sigma_i^{\text{new}} = \sqrt{\frac{\sum_{j=1}^N (x_j - \mu_i^{\text{new}})^2 p(i|x_j, \Theta^{\text{old}})}{\sum_{j=1}^N p(i|x_j, \Theta^{\text{old}})}}. \quad (9.58)$$

In order to estimate the value of M , that is, the number of types of tissue in the mass, one cannot model M as a random variable and directly apply the ML principle, because the maximum likelihood of Θ is a nondecreasing function of M [811]. The estimated value of M should be the value that minimizes a cost function that penalizes higher values of M . The common choice for such a cost function is one that follows the MDL criterion [811]; however, other criteria exist to find the value of M [811]. Ferrari et al. [375, 381, 817] successfully used the MDL criterion to find the number of Gaussian kernels in a Gaussian mixture model, in the context of detecting the fibroglandular disc in mammograms; see Section 8.9.2. However, Ayres et al. found that it is not appropriate to use the MDL criterion in the application to neuroblastoma because the Gaussian kernels to be identified overlap significantly in the HU domain.

Finite mixture models are regarded as powerful tools in unsupervised classification tasks [811]. Gaussian mixture models are the most common type of mixture models [811], and the EM algorithm is the common method of estimation of the parameters in a Gaussian mixture model [811, 818]. Mixture models have been employed with success in image processing for unsupervised classification [811, 818], automatic segmentation of brain MR images [818] and mammograms [375, 381, 817], automatic target recognition [814], correction of intensity nonuniformity in MRI [813], tissue characterization [813, 819], and partial volume segmentation in MRI [819].

Jain et al. [811] point out that current problems and research topics in using the EM algorithm are: dealing with its local nature, which causes the algorithm to be critically dependent on the initial value of Θ ; and the unbounded nature of the parameters (because, in the ML principle, no prior probability is assigned to Θ) that could cause Θ to converge to undesired points in the feature space, such as having α_i and σ_i approach zero simultaneously for the i^{th} Gaussian kernel. Although the latter problem was not encountered by Ayres et al., they did face the former problem of finding a good initial estimate of Θ .

Parameter selection and initialization: The tumor bulk in neuroblastoma commonly contains up to three different tissue components: low-attenuation necrotic tissue, intermediate-attenuation viable tumor, and high-attenuation calcified tissue. The relative quantity of each of these tissue types varies from tumor to tumor. Although the typical mean HU and standard deviation values of these types of tissue are known (as shown in Table 9.1), the statistics of the tissue types could vary from one imaging system to another, depend upon the imaging protocol (including the use of contrast agents), and be influenced by the partial-volume effect. It should also be noted that the ranges of HU values of necrotic tissue, viable tumor, and several abdominal organs overlap. For these reasons, it would be inappropriate to use fixed

bands of HU values to analyze the density distribution of a given tumor mass. The same reasons make it inappropriate to use fixed initial values for the EM algorithm. In the work of Ayres et al., the EM algorithm was initialized with three mean values ($M = 3$) computed as the mean of the histogram of the tumor, and the mean \pm one-half of the standard deviation of the histogram. The variance of all three Gaussians was initialized to the variance of the histogram of the tumor.

9.9.4 Results of application to clinical cases

Ayres et al. analyzed ten CT exams of four patients with Stage 3 neuroblastoma from the Alberta Children's Hospital, Calgary, Alberta, Canada. Tumor outlines were manually drawn on the images by a radiologist. Each patient had had an initial CT scan to assess the state of the disease prior to chemotherapy. Two patients had follow-up CT exams during treatment. All patients had a presurgical CT exam. After surgical resection, the tumor masses were analyzed by a pathologist. The following paragraphs describe the results obtained with two of the cases.

Case 1: The two-year-old male patient had an initial diagnostic CT scan in April 2001 [labeled as Exam 1a, see Figure 9.23 (a)]. The patient had a follow-up CT scan in June 2001 [labeled as Exam 1b, see Figure 9.23 (b)], and a presurgical CT scan in September 2001 [labeled as Exam 1c, see Figure 9.23 (c)]. Surgical resection of the tumor was performed in September 2001. Pathologic analysis showed extensive necrosis and dystrophic calcification.

Figure 9.24 shows the results of decomposition of the histogram of Exam 1a with different numbers of Gaussian components. (*Note:* Although only one CT slice is shown for each exam in Figure 9.23, all applicable slices of each exam were processed to obtain the corresponding histograms.) The results of estimation of the tissue composition for all exams of Case 1, assuming the existence of three tissue types, are shown in Figure 9.25 and Figure 9.26, along with the tumor volume in each CT scan in the latter figure.

The initial diagnostic scan of the patient [Exam 1a, see Figure 9.23 (a)] showed a large mass with several components. Radiological analysis indicated the existence of a calcified mass with a size of about $4.5 \times 4.4 \times 5.9$ cm, located in the right suprarenal region. The predominant components in this case are low-density necrotic tissue, intermediate-density tumor, and high-density areas of calcification, probably representing dystrophic calcification in necrotic tumor. These three components are well-demonstrated in the histogram corresponding to Exam 1a in Figure 9.26.

Exam 1b [see Figure 9.23 (b)] represents an intermediate scan performed part way through the presurgical chemotherapy regimen. The scan demonstrated an overall decrease in tumor volume together with an increasing amount of calcification. The corresponding histogram in Figure 9.26 is of interest in that it indicates a disproportionate increase in the intermediate

values. Observe, however, that the mean CT value for the central component is significantly higher than that for the initial diagnostic scan. This probably represents areas of early faint calcification within necrotic tissue; this component has likely been emphasized by partial-volume averaging, which has resulted in a higher value for the intermediate density.

Exam 1c [Figure 9.23 (c)] shows a smaller, but largely and densely calcified tumor, with very little remaining of the lower-density component. The corresponding histogram in Figure 9.26 correlates with this increasing overall density. Observe that the mean density of all three components now is high, with the emphasis particularly on the calcification. This suggests that previous necrotic tumor has progressed to dystrophic calcification with little in the way of potentially viable residual tumor.

Case 2: The two-year-old female patient had the initial diagnostic CT scan in March 2000 [labeled Exam 2a, see Figure 9.27 (a)]. The patient had the presurgical CT scan in July 2000 [labeled Exam 2b, see Figure 9.27 (b)]. Pathologic analysis of the resected mass indicated residual tumor consistent with differentiating neuroblastoma. Sections from the tumor showed extensive necrosis (consistent with previous chemotherapy), and fibrosis. The results of estimation of the tissue composition, assuming the existence of three tissue types, are shown in Figure 9.28, along with the tumor volume in each CT scan.

The initial diagnostic images of this patient demonstrated a large mass, predominantly of soft tissue (viable tumor) composition. There were significant areas of lower-attenuation necrotic tissue, but very little calcification. Radiological analysis of the presurgical scan indicated the existence of a mixed-density mass in the left adrenal region, with a size of about $5 \times 5 \times 3.6$ cm, showing peripheral calcification and central low density. The histogram for Exam 2a in Figure 9.28 correlates well with these findings.

Exam 2b [Figure 9.27 (a)] shows a post-chemotherapy, presurgical CT scan of the patient. This scan demonstrated a significant overall decrease in tumor volume. However, the composition had changed relatively little. The tumor was still composed largely of soft-tissue, low-density material with significant areas of necrosis and relatively little calcification. The histogram of Exam 2b in Figure 9.28 shows a similar composition, although there is considerable overlap between the components; observe that the mean densities of the components differ little. The lack of progression to calcification suggests that there is still considerable viable tumor remaining, with less evidence of necrosis and subsequent dystrophic calcification. These findings were confirmed by pathologic analysis, which showed residual viable tumor.

9.9.5 Discussion

With treatment, all of the four cases in the study of Ayres et al. demonstrated a significant response with an overall reduction in tumor bulk. Frequently, the tumor undergoes necrosis, seen as an increase in the relative


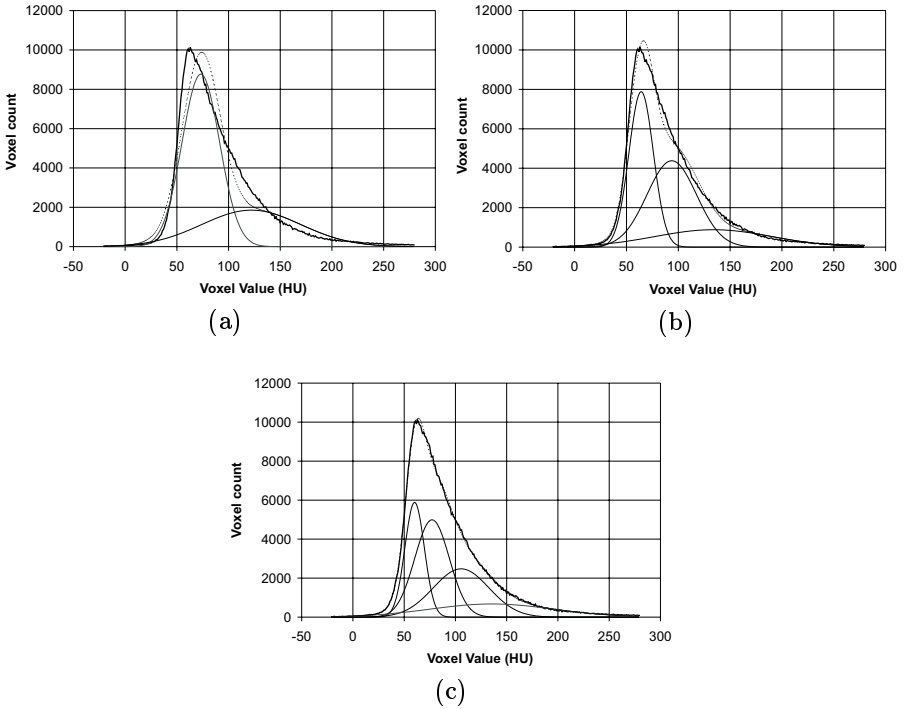



FIGURE 9.23

(a) Initial diagnostic CT image of Case 1, Exam 1a (April 2001). (b) Intermediate follow-up CT image, Exam 1b (June 2001). (c) Presurgical CT image, Exam 1c (September 2001). The contours of the tumor mass drawn by a radiologist are also shown. Reproduced with permission from F.J. Ayres, M.K. Zuffo, R.M. Rangayyan, G.S. Boag, V. Odone Filho, and M. Valente, "Estimation of the tissue composition of the tumor mass in neuroblastoma using segmented CT images", *Medical and Biological Engineering and Computing*, 42:366 – 377, 2004. © IFMBE.

**FIGURE 9.24**

Results of decomposition of the histogram of Exam 1a. Plots (a), (b), and (c) show two, three, and four estimated Gaussian kernels (thin lines), respectively, and the original histogram (thick line) for comparison. The sum of the Gaussian components is indicated in each case by the dotted curve; however, this curve is not clearly visible in (c) because it overlaps the original histogram. Figure courtesy of F.J. Ayres.



(a)





Figure 9.25 (b)

**FIGURE 9.25**

Results of decomposition of the histograms of the three CT exams of Case 1 (Figure 9.23) with three estimated Gaussian kernels (thin lines) for each histogram. The original histograms (thick line) are also shown for comparison. In each case, the sum of the Gaussian components is indicated by the dotted curve; however, this curve may not be clearly visible due to close matching with the original histogram. Reproduced with permission from F.J. Ayres, M.K. Zuffo, R.M. Rangayyan, G.S. Boag, V. Odone Filho, and M. Valente, “Estimation of the tissue composition of the tumor mass in neuroblastoma using segmented CT images”, *Medical and Biological Engineering and Computing*, 42:366 – 377, 2004. © IFMBE.





FIGURE 9.26

Results of estimation of the tumor volume and tissue composition of each CT Exam of Case 1. Reproduced with permission from F.J. Ayres, M.K. Zuffo, R.M. Rangayyan, G.S. Boag, V. Odone Filho, and M. Valente, "Estimation of the tissue composition of the tumor mass in neuroblastoma using segmented CT images", *Medical and Biological Engineering and Computing*, 42:366 – 377, 2004. © IFMBE.

**FIGURE 9.27**

(a) Initial diagnostic CT image of Case 2, Exam 2a (March 2000). (b) Presurgical CT image, Exam 2b (July 2000). The contours of the tumor mass drawn by a radiologist are also shown. Reproduced with permission from F.J. Ayres, M.K. Zuffo, R.M. Rangayyan, G.S. Boag, V. Odone Filho, and M. Valente, "Estimation of the tissue composition of the tumor mass in neuroblastoma using segmented CT images", *Medical and Biological Engineering and Computing*, 42:366 – 377, 2004. © IFMBE.




FIGURE 9.28

Results of estimation of the tumor volume and tissue composition of each CT exam of Case 2. Reproduced with permission from F.J. Ayres, M.K. Zuffo, R.M. Rangayyan, G.S. Boag, V. Odone Filho, and M. Valente, "Estimation of the tissue composition of the tumor mass in neuroblastoma using segmented CT images", *Medical and Biological Engineering and Computing*, 42:366 – 377, 2004. © IFMBE.

volume of tissue with lower attenuation values. The necrotic tissue may subsequently undergo calcification, and therefore, ultimately result in an increase in the high-attenuation calcified component. One may hypothesize, therefore, that a progression in the pattern of the histograms from predominantly intermediate-density tissues to predominantly low-attenuation necrotic tissue and ultimately to predominantly high-attenuation calcified tissue represents a good response to therapy, with the tumor progressing through necrosis to ultimate dystrophic calcification. On the contrary, the absence of this progression from necrosis to calcification, and the persistence of significant proportions of intermediate-attenuation soft tissue may be a predictor of residual viable tumor. As such, the technique proposed by Ayres et al. may be of considerable value in assessing response to therapy in patients with neuroblastoma.

Objective demonstration of the progression of a tumor through various stages, as described above, requires the use of Gaussians of variable mean values. In order to allow for the three possible tissue types mentioned above, it is necessary to allow the use of at least three Gaussians in the mixture model. However, when a tumor lacks a certain type of tissue, two (or more) of the Gaussians derived could possibly be associated with the same tissue type. This is evident, for example, in Exam 1c (Figure 9.26) where the two Gaussians with mean values of 215 and 398 *HU* correspond to calcified tissue. (Varying degrees of calcification of tissues and the partial-volume effect could have contributed to a wide range of *HU* values for calcified tissue in Exam 1c.) Furthermore, the results for Exam 1c indicate the clear absence of viable tumor, and those for Exam 2a the clear absence of calcification. It may be desirable to apply some heuristics to combine similar Gaussians (of comparable mean and variance).

Although some initial work in tissue characterization of this type was performed using CT, many investigators have shifted their interest away from CT toward MRI for the purpose of tissue characterization. Although MRI shows more long-term promise in this field due to its inherently superior definition of soft tissues, the CT technique may still be of considerable value. Specifically, MRI scanners remain an expensive and difficult-to-access specialty in many areas, whereas CT scanners have become much more economical and widespread. With regard to the clinical problem of neuroblastoma presenting in young children, the current standards of medical care for such patients include assessment by CT in almost all cases. On the other hand, MRI is used only in a minority of cases, due to the lower level of accessibility, the need for anesthesia or sedation in young children, expense, and difficulties with artifact due to bowel peristalsis. As such, CT methods for tissue characterization and assessment of tumor bulk, tissue composition, and response to therapy may be of considerable value in neuroblastoma.

It is clear from the study of Ayres et al., as well as past clinical experience, that the CT number by itself is not sufficient to define tumor versus normal tissues. Tumor definition and diagnosis require an analysis of the spatial distribution of the various CT densities coupled with a knowledge of

normal anatomy. Some work has been conducted in attempts to define automatically the boundaries of normal anatomical structures, and subsequently identify focal or diffuse abnormalities within those organs [820]. Ayres et al. made no attempt to automatically define normal versus abnormal structures, but rather attempted an analysis of the tissues in a manually identified abnormality. However, this process may ultimately prove of value for the analysis of abnormalities identified automatically by future image analysis techniques [365, 366].

9.10 Remarks

The Radon transform offers a method to convert a 2D image to a series of 1D functions (projections). This facilitates improved or convenient implementation of some image processing tasks in the Radon domain in 1D instead of in the original 2D image plane; some examples of this approach include edge detection [821] and the removal of repeated versions of a basic pattern [444, 505] (see Section 10.3).

The 1980s and 1990s brought out many new developments in CT imaging. Continuing development of versatile imaging equipment and image processing algorithms has been opening up newer applications of CT imaging. 3D imaging of moving organs such as the heart is now feasible. 3D display systems and algorithms have been developed to provide new and intriguing displays of the interior of the human body. 3D images obtained by CT are being used in planning surgery and radiation therapy, thereby creating the new fields of image-guided surgery and treatment. The practical realization of portable scanners has also made possible field applications in agricultural sciences and other biological applications. CT is a truly revolutionary investigative imaging technique — a remarkable synthesis of many scientific principles.

9.11 Study Questions and Problems

Selected data files related to some of the problems and exercises are available at the site

www.enel.ucalgary.ca/People/Ranga/enel697

1. A 2×2 image has the pixel values

$$\begin{bmatrix} 7 & 2 \\ 4 & 0 \end{bmatrix}. \quad (9.59)$$

Compute parallel-ray projections of the image at 0° and 90° . Compute a reconstruction of the image using the simple backprojection method.

- State and explain the Fourier slice theorem. Given the notations $f(x, y)$ for a function in the image domain, $p_\theta(t)$ for a function in the projection or Radon domain, and $F(u, v)$ as well as $P_\theta(w)$ for functions in the frequency or Fourier domain, explain the relationships between these functions.

With reference to the notations provided above, what do the variables $x, y, \theta, t, u, v,$ and w stand for? What are their units?

- A researcher has obtained parallel-ray projections of an image at the angles $30^\circ, 50^\circ, 70^\circ, 90^\circ, 110^\circ, 130^\circ,$ and 150° . The only algorithm available for reconstruction of the image is the Fourier method.

Draw a schematic representation of the information available in the Fourier domain. Propose methods to help the researcher obtain the best possible reconstruction of the image.

Under what conditions can a perfect reconstruction be obtained?

- Give a step-by-step description of the Fourier method for reconstructing an image from its projections. Explain the limitations of the method.
- A 2×2 image has the pixel values

$$\begin{bmatrix} 2 & 3 \\ 4 & 5 \end{bmatrix}. \quad (9.60)$$

Compute parallel-ray projections of the image at 0° and 90° . Starting with an initial estimate with all pixels equal to unity, compute reconstructions of the image over one iteration of (a) additive ART, and (b) multiplicative ART.

- One of the properties of ART is that if the hyperplanes of all the given ray sums are mutually orthogonal, we may start with any initial guess and reach the solution in only one cycle (or iteration) of projections (or corrections).

Prepare a set of two simultaneous equations in two unknowns such that the corresponding straight lines in the 2D plane are mutually orthogonal. Show graphically that, starting from any initial guess, the solution may be reached in just one iteration (two projections).

9.12 Laboratory Exercises and Projects

- Create a numerical phantom image by placing circles, ellipses, rectangles, and triangles of different intensity values within an ellipse. Compute parallel-ray projections of the image at a few different angles.

Compute reconstructions of the image using the simple backprojection and the filtered backprojection methods using various numbers of projections with different angular sampling and coverage. Compare the quality of the results obtained.

- Repeat the preceding exercise with additive and multiplicative ART.

Deconvolution, Deblurring, and Restoration

Image enhancement techniques are typically designed to yield “better looking” images satisfying some subjective criteria. In comparison with the given image, the processed image may not be closer to the true image in any sense. On the other hand, image restoration [8, 9, 11, 589, 822, 823, 824, 825, 826, 827, 828] is defined as image quality improvement under objective evaluation criteria to find the *best possible estimate of the original unknown image from the given degraded image*. The commonly used criteria are LMS, MMSE, and distance measures of several types. Additional constraints based upon prior and independent knowledge about the original image may also be imposed to limit the scope of the solution. Image restoration may then be posed as a constrained optimization problem.

Image restoration requires precise information about the degrading phenomenon, and analysis of the system that produced the degraded image. Typical items of information required are estimates or models of the impulse response of the degrading filter (the PSF, or equivalently, the MTF); the PSD (or ACF) of the original image; and the PSD of the noise. If the degrading system is shift-variant, then a model of the variation of its impulse response across the field of imaging would be required. The success of a procedure for image restoration depends upon the accuracy of the model of degradation used, and the accuracy of the functions used to represent the image degrading phenomena. In this chapter, we shall explore several techniques for image restoration under varying conditions of degradation, available information, and optimization.

10.1 Linear Space-invariant Restoration Filters

Assuming the degrading phenomenon to be linear and shift-invariant, the simplest model of image degradation is

$$\begin{aligned}g(x, y) &= h(x, y) * f(x, y) + \eta(x, y), \\G(u, v) &= H(u, v) F(u, v) + \eta(u, v),\end{aligned}\tag{10.1}$$

where f is the original image, h is the impulse response of the degrading LSI system, g is the observed (degraded) image, and η is additive random noise that is statistically independent of the image-generating process. The functions represented by upper-case letters represent the Fourier transforms of the image-domain functions represented by the corresponding lower-case letters. A block diagram of the image degradation system as above is given in Figure 10.1.




FIGURE 10.1

Image degradation model involving an LSI system and additive noise.

The image restoration problem is defined as follows: Given g and some knowledge of h , f , and η , find the best possible estimate of f . When the degrading phenomenon can be represented by an LSI system, it is possible to design LSI filters to restore the image, within certain limits. A few well-known LSI filters for image restoration are described in the following subsections.

10.1.1 Inverse filtering

Let us consider the degradation model expressed in matrix form (see Section 3.5) as

$$\mathbf{g} = \mathbf{h} \mathbf{f}, \quad (10.2)$$

with no noise being present. The restoration problem may be stated as follows: Given \mathbf{g} and \mathbf{h} , estimate \mathbf{f} .

In order to develop a mathematical statement of the problem, let us consider an approximation $\tilde{\mathbf{f}}$ to \mathbf{f} . In the least-squares approach [9], the criterion for obtaining the optimal solution is stated as follows: Minimize the squared error between the observed response \mathbf{g} , and the response $\tilde{\mathbf{g}}$ had the input been $\tilde{\mathbf{f}}$.

The error between \mathbf{g} and $\tilde{\mathbf{g}}$ is given by

$$\boldsymbol{\epsilon} = \mathbf{g} - \tilde{\mathbf{g}} = \mathbf{g} - \mathbf{h} \tilde{\mathbf{f}}. \quad (10.3)$$

The squared error is given as

$$\begin{aligned} \epsilon^2 &= \boldsymbol{\epsilon}^T \boldsymbol{\epsilon} \\ &= (\mathbf{g} - \mathbf{h} \tilde{\mathbf{f}})^T (\mathbf{g} - \mathbf{h} \tilde{\mathbf{f}}) \\ &= \mathbf{g}^T \mathbf{g} - \tilde{\mathbf{f}}^T \mathbf{h}^T \mathbf{g} - \mathbf{g}^T \mathbf{h} \tilde{\mathbf{f}} + \tilde{\mathbf{f}}^T \mathbf{h}^T \mathbf{h} \tilde{\mathbf{f}}. \end{aligned} \quad (10.4)$$

Now, we can state the image restoration problem as an optimization problem: Find $\tilde{\mathbf{f}}$ that minimizes ϵ^2 . Taking the derivative of the squared error ϵ^2 in Equation 10.4 with respect to $\tilde{\mathbf{f}}$, we get

$$\frac{\partial \epsilon^2}{\partial \tilde{\mathbf{f}}} = -2 \mathbf{h}^T \mathbf{g} + 2 \mathbf{h}^T \mathbf{h} \tilde{\mathbf{f}}. \quad (10.5)$$

Setting this expression to zero, we get

$$\tilde{\mathbf{f}} = (\mathbf{h}^T \mathbf{h})^{-1} \mathbf{h}^T \mathbf{g}. \quad (10.6)$$

This is the least-squares or pseudo-inverse solution. If \mathbf{h} is square and non-singular, we get

$$\tilde{\mathbf{f}} = \mathbf{h}^{-1} \mathbf{g}. \quad (10.7)$$

If \mathbf{h} is circulant or block-circulant, we have $\mathbf{h}^{-1} = \mathbf{W} \mathbf{D}_h^{-1} \mathbf{W}^{-1}$ (see Section 3.5.5). Then,

$$\tilde{\mathbf{f}} = \mathbf{W} \mathbf{D}_h^{-1} \mathbf{W}^{-1} \mathbf{g}, \quad (10.8)$$

which leads to

$$\tilde{F}(u, v) = \frac{G(u, v)}{H(u, v)}. \quad (10.9)$$

This operation represents the inverse filter, which may be expressed as

$$L_I(u, v) = \frac{1}{H(u, v)}. \quad (10.10)$$

It is evident that the inverse filter requires knowledge of the MTF of the degradation process; see Sections 2.9, 2.12, and 10.1.6 for discussions on methods to derive this information.

The major drawback of the inverse filter is that it fails if $H(u, v)$ has zeros, or if \mathbf{h} is singular. Furthermore, if noise is present (as in Equation 10.1), we get

$$\tilde{F}(u, v) = F(u, v) + \frac{\eta(u, v)}{H(u, v)}. \quad (10.11)$$

Problems arise because $H(u, v)$ is usually a lowpass function, whereas $\eta(u, v)$ is uniformly distributed over the entire spectrum; then, the amplified noise at higher frequencies (the second component in the equation above) overshadows the restored image.

An approach to address the singularity problem associated with the inverse filter is the use of the singular value decomposition (SVD) method [825]. A widely used implementation of this approach is an iterative algorithm based on Bialy's theorem to solve the normal equation [829]; the algorithm is also known as the Landweber iterative method [830]. McGlamery [831] demonstrated the application of the inverse filter to restore images blurred by atmospheric turbulence.


In an interesting extension of the inverse filter to compensate for distortions or aberrations caused by abnormalities in the human eye, Alonso and Barreto [832] applied a predistortion or precompensation inverse filter to test images prior to being displayed on a computer monitor. The PSF of the affected eye was estimated using the wavefront aberration function measured using a wavefront analyzer. In order to overcome the limitations of the inverse filter, a weighting function similar to the parametric Wiener filter (see Section 10.1.3) was applied. The subjects participating in the study indicated improved visual acuity in reading predistorted images of test-chart letters than in reading directly displayed test images.

Example: The original “Shapes” test image (of size 128×128 pixels) is shown in Figure 10.2 (a), along with its log-magnitude spectrum in part (b) of the figure. The image was blurred via convolution with an isotropic Gaussian PSF having a radial standard deviation of two pixels. The PSF and the related MTF are shown in parts (c) and (d) of the figure, respectively. The blurred image is shown in part (e) of the figure. Gaussian-distributed noise of variance 0.01 was added to the blurred image after normalizing the image to the range $[0, 1]$; the degraded, noisy image is shown in part (f) of the figure.


The results of application of the inverse filter to the noise-free and noisy blurred versions of the “Shapes” image are shown in Figure 10.3 in both the space and frequency domains. The result of inverse filtering of the noise-free blurred image for radial frequencies up to the maximum frequency in (u, v) , shown in part (a) of the figure, demonstrates effective deblurring. A small amount of ringing artifact may be observed upon close inspection, due to the removal of frequency components beyond a circular region [see the spectrum in part (b) of the figure]. Inverse filtering of the noisy degraded image, even when limited to radial frequencies less than 0.4 times the maximum frequency in (u, v) , resulted in significant amplification of noise that led to the complete loss of the restored image information, as shown in part (c) of the figure. Limiting the inverse filter to radial frequencies less than 0.2 times the maximum frequency in (u, v) prevented noise amplification, but also severely curtailed the restoration process, as shown by the result in part (e) of the figure. The results illustrate a severe practical limitation of the inverse filter. (See also Figures 10.5 and 10.6.)

10.1.2 Power spectrum equalization

Considering the degradation model in Equation 10.1, the method of power spectrum equalization (PSE) [825] takes the following approach: Find a linear transform \mathcal{L} so as to obtain an estimate $\tilde{f}(x, y) = \mathcal{L}[g(x, y)]$, subject to the constraint $\Phi_{\tilde{f}}(u, v) = \Phi_f(u, v)$, that is, the PSD of the restored image be equal to the PSD of the original image. Applying the linear transform \mathcal{L} to

**FIGURE 10.2**

(a) “Shapes” test image; size 128×128 pixels. (b) Log-magnitude spectrum of the test image. (c) PSF with Gaussian shape; radial standard deviation = 2 pixels. (d) MTF related to the PSF in (c). (e) Test image blurred with the PSF in (c). (f) Blurred image in (e) after normalization to $[0, 1]$ and the addition of Gaussian noise with variance = 0.01.

**FIGURE 10.3**

(a) Result of inverse filtering the blurred “Shapes” image in Figure 10.2 (e). Result of inverse filtering the noisy blurred “Shapes” image in Figure 10.2 (f) using the inverse filter up to the radial frequency equal to (c) 0.4 times and (e) 0.2 times the maximum frequency in (u, v) . The log-magnitude spectra of the images in (a), (c), and (e) are shown in (b), (d), and (f), respectively.

Equation 10.1 as well as the constraint mentioned above to the result, we get

$$\begin{aligned}\Phi_{\hat{f}}(u, v) &= |L(u, v)|^2 [|H(u, v)|^2 \Phi_f(u, v) + \Phi_{\eta}(u, v)] \\ &= \Phi_f(u, v),\end{aligned}\quad (10.12)$$

where $L(u, v)$ represents the MTF of the filter \mathcal{L} . Rearranging the expression above, we get

$$L_{\text{PSE}}(u, v) = |L(u, v)| = \left[\frac{\Phi_f(u, v)}{|H(u, v)|^2 \Phi_f(u, v) + \Phi_{\eta}(u, v)} \right]^{\frac{1}{2}} \quad (10.13)$$

$$= \left[\frac{1}{|H(u, v)|^2 + \frac{\Phi_{\eta}(u, v)}{\Phi_f(u, v)}} \right]^{\frac{1}{2}}. \quad (10.14)$$

A detailed inspection of the equation above indicates the following properties of the PSE filter:

- The PSE filter requires knowledge of the PSDs of the original image and noise processes (or models thereof).
- The PSE filter tends toward the inverse filter in magnitude when the noise PSD tends toward zero. This property may be viewed in terms of the entire noise PSD or at individual frequency samples.
- The PSE filter performs restoration in spectral magnitude only. Phase correction, if required, may be applied in a separate step. In most practical cases, the degrading PSF and MTF are isotropic, and $H(u, v)$ has no phase.
- The gain of the PSE filter is not affected by zeros in $H(u, v)$, as long as $\Phi_{\eta}(u, v)$ is also not zero at the same frequencies. (In most cases, the noise PSD is nonzero at all frequencies.)
- The gain of the PSE filter reduces to zero wherever the original image PSD is zero. The noise-to-signal PSD ratio in the denominator of Equation 10.14 controls the gain of the filter in the presence of noise.

Models of the PSDs of the original image and noise processes may be estimated from practical measurements or experiments (see Section 10.1.6). See Section 10.2 for a discussion on extending the PSE filter to blind deblurring. Examples of application of the PSE filter are provided in Sections 10.1.3 and 10.5.

10.1.3 The Wiener filter

Wiener filter theory provides for *optimal* filtering by taking into account the statistical characteristics of the image and noise processes [9, 198, 589, 833].

The filter characteristics are optimized with reference to a performance criterion. The output is guaranteed to be the best achievable result under the conditions imposed and the information provided. The Wiener filter is a powerful conceptual tool that changed traditional approaches to signal processing.

The Wiener filter performs probabilistic (stochastic) restoration with the least-squares error criterion [9, 589]. The basic degradation model used is

$$\mathbf{g} = \mathbf{h} \mathbf{f} + \boldsymbol{\eta}, \quad (10.15)$$

where \mathbf{f} and $\boldsymbol{\eta}$ are real-valued, second-order-stationary random processes that are statistically independent, with known first-order and second-order moments. Observe that this equation is the matrix form of Equation 10.1. The approach taken to estimate the original image is to determine a linear estimate $\tilde{\mathbf{f}} = \mathbf{L} \mathbf{g}$ to \mathbf{f} from the given image \mathbf{g} , where \mathbf{L} is the filter to be derived. The criterion used is to minimize the MSE

$$\epsilon^2 = E \left[\left\| \mathbf{f} - \tilde{\mathbf{f}} \right\|^2 \right]. \quad (10.16)$$

Expressing the MSE as the trace of the outer product matrix of the error vector, we have

$$\epsilon^2 = E \left\{ \text{Tr} \left[(\mathbf{f} - \tilde{\mathbf{f}}) (\mathbf{f} - \tilde{\mathbf{f}})^T \right] \right\}. \quad (10.17)$$

In expanding the expression above, we could make use of the following relationships:

$$(\mathbf{f} - \tilde{\mathbf{f}}) (\mathbf{f} - \tilde{\mathbf{f}})^T = \mathbf{f} \mathbf{f}^T - \mathbf{f} \tilde{\mathbf{f}}^T - \tilde{\mathbf{f}} \mathbf{f}^T + \tilde{\mathbf{f}} \tilde{\mathbf{f}}^T. \quad (10.18)$$

$$\tilde{\mathbf{f}}^T = \mathbf{g}^T \mathbf{L}^T = (\mathbf{f}^T \mathbf{h}^T + \boldsymbol{\eta}^T) \mathbf{L}^T. \quad (10.19)$$

$$\mathbf{f} \tilde{\mathbf{f}}^T = \mathbf{f} \mathbf{f}^T \mathbf{h}^T \mathbf{L}^T + \mathbf{f} \boldsymbol{\eta}^T \mathbf{L}^T. \quad (10.20)$$

$$\tilde{\mathbf{f}} \mathbf{f}^T = \mathbf{L} \mathbf{h} \mathbf{f} \mathbf{f}^T + \mathbf{L} \boldsymbol{\eta} \mathbf{f}^T. \quad (10.21)$$

$$\tilde{\mathbf{f}} \tilde{\mathbf{f}}^T = \mathbf{L} (\mathbf{h} \mathbf{f} \mathbf{f}^T \mathbf{h}^T + \mathbf{h} \mathbf{f} \boldsymbol{\eta}^T + \boldsymbol{\eta} \mathbf{f}^T \mathbf{h}^T + \boldsymbol{\eta} \boldsymbol{\eta}^T) \mathbf{L}^T. \quad (10.22)$$

Because the trace of a sum of matrices is equal to sum of their traces, the E and Tr operators may be interchanged in order. We then obtain the following expressions and relationships:

$$E [\mathbf{f} \mathbf{f}^T] = \boldsymbol{\phi}_f, \quad (10.23)$$

the autocorrelation matrix of the original image.

$$E [\mathbf{f} \tilde{\mathbf{f}}^T] = \boldsymbol{\phi}_f \mathbf{h}^T \mathbf{L}^T, \quad (10.24)$$

with the observation that

$$E [\mathbf{f} \boldsymbol{\eta}^T] = 0 \quad (10.25)$$

because \mathbf{f} and $\boldsymbol{\eta}$ are statistically independent processes and $\mu_{\boldsymbol{\eta}} = 0$.

$$E [\tilde{\mathbf{f}} \mathbf{f}^T] = \mathbf{L} \mathbf{h} \boldsymbol{\phi}_f. \quad (10.26)$$

$$E \left[\tilde{\mathbf{f}} \tilde{\mathbf{f}}^T \right] = \mathbf{L} \mathbf{h} \boldsymbol{\phi}_f \mathbf{h}^T \mathbf{L}^T + \mathbf{L} \boldsymbol{\phi}_\eta \mathbf{L}^T. \quad (10.27)$$

$$\boldsymbol{\phi}_\eta = E \left[\boldsymbol{\eta} \boldsymbol{\eta}^T \right] \quad (10.28)$$

is the autocorrelation matrix of the noise process.

Now, the MSE may be written as:

$$\begin{aligned} \epsilon^2 &= \text{Tr} \left(\boldsymbol{\phi}_f - \boldsymbol{\phi}_f \mathbf{h}^T \mathbf{L}^T - \mathbf{L} \mathbf{h} \boldsymbol{\phi}_f + \mathbf{L} \mathbf{h} \boldsymbol{\phi}_f \mathbf{h}^T \mathbf{L}^T + \mathbf{L} \boldsymbol{\phi}_\eta \mathbf{L}^T \right) \\ &= \text{Tr} \left(\overline{\boldsymbol{\phi}_f} - 2 \boldsymbol{\phi}_f \mathbf{h}^T \mathbf{L}^T + \mathbf{L} \mathbf{h} \boldsymbol{\phi}_f \mathbf{h}^T \mathbf{L}^T + \mathbf{L} \boldsymbol{\phi}_\eta \mathbf{L}^T \right). \end{aligned} \quad (10.29)$$

(Note: $\boldsymbol{\phi}_f^T = \boldsymbol{\phi}_f$ and $\boldsymbol{\phi}_\eta^T = \boldsymbol{\phi}_\eta$ because the autocorrelation matrices are symmetric, and the trace of a matrix is equal to the trace of its transpose.)

At this point, the MSE is no longer a function of \mathbf{f} , \mathbf{g} , or $\boldsymbol{\eta}$, but depends only on the statistical characteristics of \mathbf{f} and $\boldsymbol{\eta}$, as well as on \mathbf{h} and \mathbf{L} .

In order to derive the optimal filter \mathbf{L} , we could set the derivative of ϵ^2 with respect to \mathbf{L} to zero:

$$\frac{\partial \epsilon^2}{\partial \mathbf{L}} = -2 \boldsymbol{\phi}_f \mathbf{h}^T + 2 \mathbf{L} \mathbf{h} \boldsymbol{\phi}_f \mathbf{h}^T + 2 \mathbf{L} \boldsymbol{\phi}_\eta = 0, \quad (10.30)$$

which leads to the optimal Wiener filter function

$$\mathbf{L}_W = \boldsymbol{\phi}_f \mathbf{h}^T \left(\mathbf{h} \boldsymbol{\phi}_f \mathbf{h}^T + \boldsymbol{\phi}_\eta \right)^{-1}. \quad (10.31)$$

Note that this solution does not depend upon the inverses of the individual ACF matrices or \mathbf{h} , but upon the inverse of their combination. The combined matrix could be expected to be invertible even if the individual matrices are singular.

Considerations in implementation of the Wiener filter: Consider the matrix to be inverted in Equation 10.31:

$$\mathbf{Z} = \mathbf{h} \boldsymbol{\phi}_f \mathbf{h}^T + \boldsymbol{\phi}_\eta. \quad (10.32)$$

This matrix would be of size $N^2 \times N^2$ for $N \times N$ images, making inversion practically impossible. Inversion becomes easier if the matrix can be written as a product of diagonal and unitary matrices. A condition that reduces the complexity of the problem is that \mathbf{h} , $\boldsymbol{\phi}_f$, and $\boldsymbol{\phi}_\eta$ each be circulant or block-circulant. We can now make the following observations:

- We know, from Section 3.5, that \mathbf{h} is block-circulant for 2D LSI operations expressed using circular convolution.
- $\boldsymbol{\phi}_\eta$ is a diagonal matrix if η is white (uncorrelated) noise.
- In most real images, the correlation between pixels reduces as the distance (spatial separation) increases: $\boldsymbol{\phi}_f$ is then banded and may be approximated by a block-circulant matrix.

Based upon the observations listed above, we can write (see Section 3.5)

$$\mathbf{h} = \mathbf{W} \mathbf{D}_h \mathbf{W}^{-1}, \quad (10.33)$$

$$\phi_f = \mathbf{W} \mathbf{D}_f \mathbf{W}^{-1}, \quad (10.34)$$

and

$$\phi_\eta = \mathbf{W} \mathbf{D}_\eta \mathbf{W}^{-1}, \quad (10.35)$$

where the matrices \mathbf{D} are diagonal matrices resulting from the application of the DFT to the corresponding block-circulant matrices, with the subscripts indicating the related entity (f : original image, h : degrading system, or η : noise). Then, we have

$$\begin{aligned} \mathbf{Z} &= \mathbf{W} \mathbf{D}_h \mathbf{D}_f \mathbf{D}_h^* \mathbf{W}^{-1} + \mathbf{W} \mathbf{D}_\eta \mathbf{W}^{-1} \\ &= \mathbf{W} (\mathbf{D}_h \mathbf{D}_f \mathbf{D}_h^* + \mathbf{D}_\eta) \mathbf{W}^{-1}. \end{aligned} \quad (10.36)$$

The Wiener filter is then given by

$$\mathbf{L}_W = \mathbf{W} \mathbf{D}_f \mathbf{D}_h^* (\mathbf{D}_h \mathbf{D}_f \mathbf{D}_h^* + \mathbf{D}_\eta)^{-1} \mathbf{W}^{-1}. \quad (10.37)$$

The optimal MMSE estimate is given by

$$\begin{aligned} \tilde{\mathbf{f}} &= \mathbf{L}_W \mathbf{g} \\ &= \mathbf{W} \mathbf{D}_f \mathbf{D}_h^* (\mathbf{D}_h \mathbf{D}_f \mathbf{D}_h^* + \mathbf{D}_\eta)^{-1} \mathbf{W}^{-1} \mathbf{g}. \end{aligned} \quad (10.38)$$

Interpretation of the Wiener filter: With reference to Equation 10.38, we can make the following observations that help in interpreting the nature of the Wiener filter.

- $\mathbf{W}^{-1} \mathbf{g}$ is related to $G(u, v)$, the Fourier transform of the given degraded image $g(x, y)$.
- \mathbf{D}_f is related to the PSD $\Phi_f(u, v)$ of the original image $f(x, y)$.
- \mathbf{D}_η is related to the PSD $\Phi_\eta(u, v)$ of the noise process.
- \mathbf{D}_h is related to the transfer function $H(u, v)$ of the degrading system.

Then, the output of the Wiener filter before the final inverse Fourier transform is given by

$$\begin{aligned} \tilde{F}(u, v) &= \frac{\Phi_f(u, v) H^*(u, v) G(u, v)}{H(u, v) \Phi_f(u, v) H^*(u, v) + \Phi_\eta(u, v)} \\ &= \left[\frac{H^*(u, v)}{|H(u, v)|^2 + \frac{\Phi_\eta(u, v)}{\Phi_f(u, v)}} \right] G(u, v) \\ &= \left[\frac{|H(u, v)|^2}{|H(u, v)|^2 + \frac{\Phi_\eta(u, v)}{\Phi_f(u, v)}} \right] \frac{G(u, v)}{H(u, v)}. \end{aligned} \quad (10.39)$$

The Wiener filter itself is given by the expression

$$L_W(u, v) = \left[\frac{H^*(u, v)}{|H(u, v)|^2 + \frac{\Phi_\eta(u, v)}{\Phi_f(u, v)}} \right]. \quad (10.40)$$

We can now note the following characteristics of the Wiener filter [9, 589]:

- In the absence of noise, we have $\Phi_\eta(u, v) = 0$, and the Wiener filter reduces to the inverse filter. This is also applicable at any frequency (u, v) where $\Phi_\eta(u, v) = 0$.
- The gain of the Wiener filter is modulated by the noise-to-signal (spectral) ratio $\frac{\Phi_\eta(u, v)}{\Phi_f(u, v)}$. If the SNR is high, the Wiener filter is close to the inverse filter.
- If the SNR is poor and both the signal and noise PSDs are “white”, the ratio $\frac{\Phi_\eta}{\Phi_f}$ is large and could be assumed to be a constant. Then, the Wiener filter is close to the matched filter $H^*(u, v)$.
- If the noise-to-signal PSD ratio is not available as a function of (u, v) , it may be set equal to a constant $K = \frac{1}{\overline{\text{SNR}}}$. The filter is then known as the parametric Wiener filter.
- The Wiener filter is not often singular or ill-conditioned. Wherever $H(u, v) = 0$, the output $\tilde{F}(u, v) = 0$.
- If $h(x, y) = \delta(x, y)$, then $H(u, v) = 1$; that is, there is no blurring and the degradation is due to additive noise only. Then

$$\tilde{F}(u, v) = \left[\frac{\Phi_f(u, v)}{\Phi_f(u, v) + \Phi_\eta(u, v)} \right] G(u, v), \quad (10.41)$$

which is the original Wiener filter for the degradation model $\mathbf{g} = \mathbf{f} + \boldsymbol{\eta}$ (see Section 3.6.1 and Rangayyan [31]). The Wiener filter as in Equation 10.38 was proposed by Helstrom [197].

Comparative analysis of the inverse, PSE, and Wiener filters:

- When the noise PSD is zero, or, at all frequencies where the noise PSD is equal to zero, the PSE filter is equivalent to the inverse filter in magnitude.
- When the noise PSD is zero, or, at all frequencies where the noise PSD is equal to zero, the Wiener filter is equivalent to the inverse filter.
- The gains of the inverse, PSE, and Wiener filter are related as

$$|L_I(u, v)| > |L_{\text{PSE}}(u, v)| > |L_W(u, v)|. \quad (10.42)$$

- The PSE filter is the geometric mean of the inverse and Wiener filters, that is,


$$L_{\text{PSE}}(u, v) = [L_{\text{I}}(u, v) L_{\text{W}}(u, v)]^{\frac{1}{2}}. \quad (10.43)$$

- Because $|L_{\text{PSE}}(u, v)| > |L_{\text{W}}(u, v)|$, the PSE filter admits more high-frequency components with larger gain than the Wiener filter. Therefore, the result will be a sharper, but more noisy, image.
- The PSE filter does not have a phase component. Phase correction, if required, may be applied in a separate step.


Examples: The original “Shapes” test image and a degraded version of the image, with blurring via convolution using an isotropic Gaussian PSF having a radial standard deviation of two pixels as well as the addition of Gaussian-distributed noise of variance 0.01 to the blurred image after normalizing the image to the range $[0, 1]$, are shown in Figure 10.4 (a) and (b), respectively. (See also Figures 10.2 and 10.3.) In order to design the appropriate Wiener and PSE filters, a model PSD of the image was prepared by computing the PSD of an image of a square object of size 11×11 pixels; values of the PSD less than a limit equal to 10% of its maximum were replaced by the limit. The noise PSD was prepared as an array of constant value equal to the known variance of the noise times N^2 , where $N \times N$ is the size the PSD array. The results of the Wiener and PSE filters are shown in Figure 10.4 (c) and (d), respectively. It is evident that both filters have removed the blur to some extent; however, as expected, the result of the PSE filter is noisier than that of the Wiener filter.

Figures 10.5 and 10.6 illustrate the application of the inverse, Wiener, and PSE filters to degraded versions of the Lenna and Cameraman images, respectively. The Lenna image was blurred using a Gaussian-shaped blur function with a radial standard deviation of three pixels and additive noise to $SNR = 35 \text{ dB}$. The Cameraman image was blurred with a straight line of length nine pixels, simulating blurring due to motion in the horizontal direction, and then degraded further with additive noise to $SNR = 35 \text{ dB}$. Inverse filtering even to limited ranges of frequency has resulted in distorted images due to the amplification of noise as well as ringing artifacts. The Wiener and PSE filters have removed the blur to a good extent, with control over the noise. Further examples of application of the Wiener filter are provided in Section 10.4.2.


Aghdasi et al. [202] applied the parametric Wiener filter to deblur mammograms after removing noise using the local LMMSE filter (see Section 3.7.1) or a Bayesian filter. King et al. [834] and Honda et al. [835] applied the Wiener filter to the restoration of nuclear medicine images. A comparative analysis of the performance of the PSE, Wiener, and Metz filters, in the restoration of nuclear medicine images, is provided in Section 10.5.

**FIGURE 10.4**

(a) “Shapes” test image; size 128×128 pixels. (b) Test image blurred with a Gaussian PSF of radial standard deviation $\sigma = 2$ pixels and degraded with additive Gaussian noise with variance = 0.01 after normalization to $[0, 1]$. (c) Result of Wiener restoration. (d) Result of restoration using the PSE filter.

**FIGURE 10.5**

(a) Lenna test image of size 128×128 pixels. (b) Blurred image with a Gaussian-shaped blur function and noise to $SNR = 35$ dB. Deblurred images: (c) Inverse filtering up to the radial frequency 0.3 times the maximum frequency in (u, v) ; (d) Inverse filtering up to the radial frequency 0.2 times the maximum frequency in (u, v) ; (e) Wiener filtering; (f) Power spectrum equalization. Reproduced with permission from T.F. Rabie, R.M. Rangayyan, and R.B. Paranjape, "Adaptive-neighborhood image deblurring", *Journal of Electronic Imaging*, 3(4): 368 – 378, 1994. © SPIE.

**FIGURE 10.6**

(a) Cameraman test image of size 128×128 pixels and gray-level range 0–255. (b) Image blurred by 9-pixel horizontal motion and degraded by additive Gaussian noise to $SNR = 35$ dB. Deblurred images: (c) Inverse filtering up to the radial frequency 0.8 times the maximum frequency in (u, v) ; (d) Inverse filtering up to the radial frequency 0.4 times the maximum frequency in (u, v) ; (e) Wiener filtering; (f) Power spectrum equalization. Reproduced with permission from T.F. Rabie, R.M. Rangayyan, and R.B. Paranjape, “Adaptive-neighborhood image deblurring”, *Journal of Electronic Imaging*, 3(4): 368 – 378, 1994. © SPIE.

10.1.4 Constrained least-squares restoration

The Wiener filter is derived using a statistical procedure based on the correlation matrices of the image and noise processes. The filter is optimal, in an average sense, for the class of images represented by the statistical entities used. It is possible that the result provided by the Wiener filter for a specific image on hand is less than satisfactory. Gonzalez and Wintz [589] describe a restoration procedure that is optimal for the specific image given, under particular constraints that are imposed, as follows.

Let \mathbf{L} be a linear filter operator. Using the degradation model as in Equation 10.15, the restoration problem may be posed as the following *constrained optimization problem*: Minimize $\|\mathbf{L}\tilde{\mathbf{f}}\|^2$ subject to $\|\mathbf{g} - \mathbf{h}\tilde{\mathbf{f}}\|^2 = \|\boldsymbol{\eta}\|^2$, where $\tilde{\mathbf{f}}$ is the estimate of \mathbf{f} being derived. Using the method of Lagrange multipliers, we now seek $\tilde{\mathbf{f}}$ that minimizes the function

$$J(\tilde{\mathbf{f}}) = \|\mathbf{L}\tilde{\mathbf{f}}\|^2 + \alpha \left[\|\mathbf{g} - \mathbf{h}\tilde{\mathbf{f}}\|^2 - \|\boldsymbol{\eta}\|^2 \right], \quad (10.44)$$

where α is the Lagrange multiplier. Taking the derivative of $J(\tilde{\mathbf{f}})$ with respect to $\tilde{\mathbf{f}}$ and setting it equal to zero, we get

$$\frac{\partial J(\tilde{\mathbf{f}})}{\partial \tilde{\mathbf{f}}} = 0 = 2\mathbf{L}^T \mathbf{L}\tilde{\mathbf{f}} - 2\alpha \mathbf{h}^T (\mathbf{g} - \mathbf{h}\tilde{\mathbf{f}}), \quad (10.45)$$

which leads to

$$\tilde{\mathbf{f}} = (\mathbf{h}^T \mathbf{h} + \gamma \mathbf{L}^T \mathbf{L})^{-1} \mathbf{h}^T \mathbf{g}, \quad (10.46)$$

where $\gamma = \frac{1}{\alpha}$.

Due to the ill-conditioned nature of restoration procedures, the results are often obscured by noise and high-frequency artifacts. Artifacts and noise may be minimized by formulating a criterion of optimality based on smoothness, such as minimizing the Laplacian of the output. In order to agree with the matrix-vector formulation as above, let us construct a block-circulant matrix \mathbf{L} using the Laplacian operator in Equation 2.83. Now, \mathbf{L} is diagonalized by the 2D DFT as $\mathbf{D}_L = \mathbf{W}^{-1} \mathbf{L} \mathbf{W}$, where \mathbf{D}_L is a diagonal matrix. Then, we have

$$\begin{aligned} \tilde{\mathbf{f}} &= (\mathbf{h}^T \mathbf{h} + \gamma \mathbf{L}^T \mathbf{L})^{-1} \mathbf{h}^T \mathbf{g} \\ &= (\mathbf{W} \mathbf{D}_h^* \mathbf{D}_h \mathbf{W}^{-1} + \gamma \mathbf{W} \mathbf{D}_L^* \mathbf{D}_L \mathbf{W}^{-1})^{-1} \mathbf{W} \mathbf{D}_h^* \mathbf{W}^{-1} \mathbf{g}. \end{aligned} \quad (10.47)$$

The estimate before the final inverse Fourier transform is given by

$$\mathbf{W}^{-1} \tilde{\mathbf{f}} = (\mathbf{D}_h^* \mathbf{D}_h + \gamma \mathbf{D}_L^* \mathbf{D}_L)^{-1} \mathbf{D}_h^* \mathbf{W}^{-1} \mathbf{g}, \quad (10.48)$$

which is equivalent to

$$\tilde{F}(u, v) = \left[\frac{H^*(u, v)}{|H(u, v)|^2 + \gamma |L(u, v)|^2} \right] G(u, v), \quad (10.49)$$

where $L(u, v)$ is the transfer function related to the constraint operator \mathbf{L} .

The expression in Equation 10.49 resembles the result of the Wiener filter in Equation 10.39, but does not require the PSDs of the image and noise processes. However, estimates of the mean and variance of the noise process are required to determine the optimal value for γ , which must be adjusted to satisfy the constraint $\|\mathbf{g} - \mathbf{h}\tilde{\mathbf{f}}\|^2 = \|\boldsymbol{\eta}\|^2$. It is also worth observing that, if $\gamma = 0$, the filter reduces to the inverse filter.

Gonzalez and Wintz [589] give an iterative procedure to estimate γ as follows: Define a residual vector as

$$\begin{aligned}\mathbf{r} &= \mathbf{g} - \mathbf{h}\tilde{\mathbf{f}} \\ &= \mathbf{g} - \mathbf{h}(\mathbf{h}^T\mathbf{h} + \gamma\mathbf{L}^T\mathbf{L})^{-1}\mathbf{h}^T\mathbf{g}.\end{aligned}\quad (10.50)$$

We wish to adjust γ such that

$$\|\mathbf{r}\|^2 = \|\boldsymbol{\eta}\|^2 \pm \epsilon, \quad (10.51)$$

where ϵ is a factor of accuracy. Iterative trial-and-error methods or the Newton-Raphson procedure may be used to determine the optimal value for γ . A simple iterative procedure to find the optimal value for γ is as follows [589]:

1. Choose an initial value for γ .
2. Compute $\tilde{F}(u, v)$ and $\tilde{\mathbf{f}}$.
3. Form the residual vector \mathbf{r} and compute $\|\mathbf{r}\|^2$.
4. Increment γ if $\|\mathbf{r}\|^2 < \|\boldsymbol{\eta}\|^2 - \epsilon$,
or decrement γ if $\|\mathbf{r}\|^2 > \|\boldsymbol{\eta}\|^2 + \epsilon$, and return to Step 2.
Stop if $\|\mathbf{r}\|^2 = \|\boldsymbol{\eta}\|^2 \pm \epsilon$.

Note: $\|\boldsymbol{\eta}\|^2$ is the total squared value or total energy of the noise process, and is given by $(\sigma_\eta^2 + \mu_\eta^2)$ multiplied with the image area.

Other approaches to constrained restoration: Several constrained optimization methods have been developed, such as the constrained least squares method, the maximum-entropy method, and the Leahy-Goutis method [829]. The constrained least squares method requires only information about the mean and variance of the noise function, and the estimate of the image is defined as that which minimizes the output of a linear operator applied to the image; the result may be further subjected to an a priori constraint or an a posteriori measurable feature, such as the smoothness of the image [196, 836]. The maximum-entropy approach maximizes the entropy of the image data subject to the constraint that the solution should reproduce the ideal image exactly or within a tolerable uncertainty defined by the observation noise statistics [11, 837, 838]. The Leahy-Goutis method optimizes a convex criterion function, such as minimum energy or cross entropy, over the intersection of two convex constraint sets [839].

In constrained approaches, the estimate needs to satisfy a priori constraints on the ideal image. The a priori constraints may be obtained from information about the blur, noise, or the image [829]. An example of constrained approaches is the PSE filter, which estimates the image using a linear transform of the image based on the constraint that the PSD of the filtered image be equal to that of the ideal image: see Section 10.1.2.

The maximum-a-posteriori probability (MAP) technique is designed to find the estimate that maximizes the probability of the actual image conditioned on the observation (posterior probability). In the case of linear systems and under Gaussian assumptions, the MAP estimate reduces to the LMMSE estimate [11].

Several other constraints may be imposed upon restoration filters and their results. A simple yet effective constraint is that of nonnegativity of the image values, which is relevant in several practical applications where the pixels represent physical parameters that cannot take negative values. Biraud [840] described techniques to increase the resolving power of signals by imposing nonnegativity. Boas and Kac [841] derived inequalities for the Fourier transforms of positive functions that may be used as limits on spectral components of restored signals or images. Webb et al. [842] developed a constrained deconvolution filter for application to liver SPECT images; several parameters of the filter could be varied to render the filter equivalent to the inverse, Wiener, PSE, and other well-known filters.

10.1.5 The Metz filter

Metz and Beck [843] proposed a modification to the inverse filter for application to nuclear medicine images, including noise suppression at high frequencies. The Metz filter is defined as

$$L_M(u, v) = \frac{1 - [1 - H^2(u, v)]^\chi}{H(u, v)}, \quad (10.52)$$

where χ is a factor that controls the extent in frequency up to which the inverse filter is predominant, after which the noise-suppression feature becomes stronger. It is apparent that, if $\chi = 0$, the gain of the Metz filter reduces to zero. Given that most degradation phenomena have a blurring or smoothing effect, $H(u, v)$ is a lowpass filter. As a consequence, $[1 - H^2(u, v)]$ is a highpass filter, and the numerator in Equation 10.52 is a lowpass filter, whose response is controlled by the factor χ . The factor χ may be selected so as to minimize the MSE between the filtered and the ideal images.

King et al. [834, 844, 845, 846, 847, 848, 849], Gilland et al. [850], Honda et al. [835], Hon et al. [132], and Boulfelfel et al. [86, 749, 751, 851] applied the Wiener and Metz filters for the restoration of nuclear medicine images. A comparative analysis of the performance of the Metz filter with other filters, in the restoration of nuclear medicine images, is provided in Section 10.5.

10.1.6 Information required for image restoration

It is evident from Equations 10.10, 10.14, and 10.40 that image restoration, using techniques such as the inverse, PSE, and Wiener filters, requires specific items of information, including the MTF of the degradation phenomenon $H(u, v)$, the PSD of the noise process $\Phi_n(u, v)$, and the PSD of the original image process $\Phi_f(u, v)$. Although this requirement may appear to be onerous, methods and approaches may be devised to obtain each item based upon prior or additional knowledge, measurements, and derivations.

Several methods to obtain the PSF and MTF are described in Sections 2.9 and 2.12. The PSF or related functions may be measured if the imaging system is accessible and phantoms may be constructed for imaging as described in Sections 2.9 and 2.12. The PSF may also be modeled mathematically if the blurring phenomenon is known precisely, and can be represented as a mathematical process; an example of this possibility is described in Section 10.1.7.

The PSD of an image-generating stochastic process may be estimated as the average of the PSDs of several observations or realizations of the process of interest. However, care needs to be exercised in selecting the samples such that they characterize the same process or type of images. For example, the PSDs of a large collection of high-quality CT images of the brain of similar characteristics may be averaged to derive a PSD model to represent such a population of images. It would be inappropriate to combine CT images of the brain with CT images of the chest or the abdomen, or to combine CT images with MR images. In order to overcome the limitations imposed by noise as well as a specific and small collection of sample images, a practical approach would be to estimate the average ACF of the images, fit a model such as a Laplacian, and apply the Fourier transform to obtain the PSD model.

The PSD of noise may be estimated by using a collection of images taken of phantoms with uniform internal characteristics, or a collection of parts of images outside the patient or object imaged (representing the background).

When the required functions are unavailable, it will not be possible to apply the inverse, PSE, or Wiener filters. However, using a few approximations, it becomes possible to overcome the requirement of explicit and distinct knowledge of some of the functions; image restoration may then be performed while remaining “blind” to these entities. Methods for blind deblurring are described in Section 10.2.

10.1.7 Motion deblurring

Images acquired of living organisms are often affected by blurring due to motion during the period of exposure (imaging). In some cases, it may be appropriate to assume that the motion is restricted to within the plane of the image. Furthermore, it may also be assumed that the velocity is constant over the (usually short) period of exposure. Under such conditions, it becomes

possible to derive an analytical expression to represent the blurring function (PSF).

Consider the imaging of an object $f(x, y)$, over the exposure interval $[0, T]$. Let the relative motion between the object and the camera (or the imaging system) be represented as a displacement of $[\alpha(t), \beta(t)]$ at the instant of time t . Then, the recorded image $g(x, y)$ is given by [589]

$$g(x, y) = \int_0^T f[x - \alpha(t), y - \beta(t)] dt. \quad (10.53)$$

Observe that, due to the integration over the period of exposure, the resulting image $g(x, y)$ is not a function of time t . In order to derive the MTF of the imaging (blurring) system, we may analyze the Fourier transform of $g(x, y)$, as follows.

$$\begin{aligned} G(u, v) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} g(x, y) \exp[-j 2 \pi (ux + vy)] dx dy \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \left\{ \int_0^T f[x - \alpha(t), y - \beta(t)] dt \right\} \exp[-j 2 \pi (ux + vy)] \\ &\quad dx dy. \end{aligned} \quad (10.54)$$

Reversing the order of integration with respect to t and (x, y) , we get

$$G(u, v) = \int_0^T \left\{ \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f[x - \alpha(t), y - \beta(t)] \exp[-j 2 \pi (ux + vy)] dx dy \right\} dt. \quad (10.55)$$

The expression within the braces above represents the Fourier transform of $f[x - \alpha(t), y - \beta(t)]$, which is $\exp\{-j 2 \pi [u \alpha(t) + v \beta(t)]\} F(u, v)$. Then, we have

$$G(u, v) = F(u, v) \int_0^T \exp\{-j 2 \pi [u \alpha(t) + v \beta(t)]\} dt. \quad (10.56)$$

Therefore, the MTF of the system is given by

$$H(u, v) = \int_0^T \exp\{-j 2 \pi [u \alpha(t) + v \beta(t)]\} dt. \quad (10.57)$$

Thus, if the exact nature of the motion during the period of exposure is known, the MTF may be derived in an analytical form.

Let us consider the simple situation of linear motion within the plane of the image and with constant velocity, such that the total displacement during the period of exposure T is given by (a, b) in the (x, y) directions, respectively.

Then, we have $\alpha(t) = \frac{at}{T}$ and $\beta(t) = \frac{bt}{T}$, and

$$\begin{aligned} H(u, v) &= \int_0^T \exp \left[-j 2 \pi \left(u \frac{at}{T} + v \frac{bt}{T} \right) \right] dt \\ &= T \frac{\sin[\pi(ua + vb)]}{\pi(ua + vb)} \exp[-j \pi (ua + vb)]. \end{aligned} \quad (10.58)$$

In the case of linear motion in one direction only, we have $a = 0$ or $b = 0$, and the 2D function above reduces to a 1D function. It follows that the corresponding PSF is a rectangle or box function that reduces to a straight line in the case of motion in one direction only.

Given full knowledge of the MTF of the degradation phenomenon, it becomes possible to design an appropriate restoration filter, such as the inverse, PSE, or Wiener filter. However, it should be observed that the sinc function in Equation 10.58 has several zeros in the (u, v) plane, at which points the inverse filter would pose problems.

An example of simulated motion blurring is given in Section 10.4.2, along with its restoration; see also Figure 10.6. See Sondhi [826], Gonzalez and Wintz [589], and Gonzalez and Woods [8] for several examples of motion deblurring.

10.2 Blind Deblurring

In some cases, it may not be possible to obtain distinct models of the degradation phenomena. An inspection of Equation 10.1 reveals the fact that the given degraded image contains information regarding the blurring system's MTF and the noise spectrum, albeit in a combined form; in particular, we have

$$\Phi_g(u, v) = |H(u, v)|^2 \Phi_f(u, v) + \Phi_\eta(u, v). \quad (10.59)$$

Several methods have been proposed to exploit this feature for "blind" deblurring, with the adjective representing the point that the procedures do not require explicit models of the degrading phenomena. Two procedures for blind deblurring are described in the following paragraphs.

Extension of PSE to blind deblurring: The blurred image itself may be used to derive the parameters required for PSE as follows [9, 589, 825]. Suppose that the given $N \times N$ image $g(m, n)$ is broken into $M \times M$ segments, $g_l(m, n)$, $l = 1, 2, \dots, Q^2$, where $Q = \frac{N}{M}$, and M is larger than the dimensions of the blurring PSF. Then,

$$g_l(m, n) = h(m, n) * f_l(m, n) + \eta_l(m, n), \quad (10.60)$$

and

$$\Phi_{g_i}(u, v) = |H(u, v)|^2 \Phi_{f_i}(u, v) + \Phi_{\eta}(u, v). \quad (10.61)$$

In the expressions above, the combined effect of blurring across the boundaries of adjacent subimages is ignored. If we now average the PSDs Φ_{g_i} over all of the Q^2 available image segments, and make the further assumption that the averaged PSDs tend toward the true signal and noise PSDs, we have

$$\frac{1}{Q^2} \sum_{l=1}^{Q^2} \Phi_{g_l}(u, v) = |H(u, v)|^2 \tilde{\Phi}_f(u, v) + \tilde{\Phi}_\eta(u, v), \quad (10.62)$$

where $\tilde{}$ indicates that the corresponding PSDs are estimates. The expression in Equation 10.62 gives the denominator of the PSE filter as in Equation 10.13; the numerator $\tilde{\Phi}_f$ is required to be estimated separately. The MTF $H(u, v)$ and the noise PSD $\tilde{\Phi}_\eta(u, v)$ need not be estimated individually; the procedure is thus “blind” to these entities.

10.2.1 Iterative blind deblurring

Rabie et al. [852] presented an iterative technique for blind deconvolution under the assumptions that the MTF of the LSI system causing the degradation has zero phase, and that its magnitude is a smoothly varying function of frequency. These assumptions are valid for several types of degradation, such as motion blur and out-of-focus blur [853, 854]. It has been demonstrated in several works that the spectral magnitude in the Fourier representation of a signal is affected by the blur function, while many of the important features of the signal, such as edge locations, are preserved in the phase [854, 855, 856, 857, 858, 859, 860, 861]. For example, it has been shown that the intelligibility of a sentence is retained if the phase of the Fourier transform of a long segment of the speech signal is combined with unit magnitude [854].

The method of Rabie et al. [852] makes use of the image characteristics (edge information) preserved in the phase of the blurred image, and attempts to recover the original magnitude spectrum that was altered by the blur function. Their blind deconvolution algorithm, described in the following paragraphs, differs from earlier related work [853, 862, 863] in that the averaging of the PSD is achieved by smoothing in the frequency domain; hence, the method is based upon the use of the entire image rather than sections of the image (see Section 10.4.1). This feature relaxes the assumption on the region of support (ROS) of the PSF. Another key feature of the algorithm is that further enhancement of the initial estimate of the image is achieved through an iterative approach.

Using the degradation model expressed in Equation 10.1 and neglecting the presence of noise, we have

$$M_g(u, v) = M_f(u, v) M_h(u, v), \quad (10.63)$$

and

$$\theta_g(u, v) = \theta_f(u, v), \quad (10.64)$$

where $M_f(u, v)$ is the spectral magnitude of the original image, $M_g(u, v)$ is the spectral magnitude of the degraded image, $M_h(u, v)$ is the degradation MTF with the property that it is a smooth magnitude function with zero phase, $\theta_f(u, v)$ is the spectral phase of the original image, and $\theta_g(u, v)$ is the spectral phase of the degraded image. The blur model is thus defined in terms of magnitude and phase. The spectral magnitude $M_f(u, v)$ of the original image may be recovered from the spectral magnitude $M_g(u, v)$ of the degraded image as follows. The initial estimate of $M_f(u, v)$ is based on smoothing the spectral magnitude of the blurred image, $M_g(u, v)$, and using the assumption that $M_h(u, v)$ is smooth. If we let $\mathcal{S}[\]$ denote a 2D linear, separable, smoothing operator, then a smoothed $M_g(u, v)$ is given by

$$\mathcal{S}[M_g(u, v)] = \mathcal{S}[M_f(u, v) M_h(u, v)]. \quad (10.65)$$

Because $M_h(u, v)$ is a smooth function and $\mathcal{S}[\]$ is separable, $\mathcal{S}[M_f(u, v) M_h(u, v)]$ may be approximated by $\mathcal{S}[M_f(u, v)] M_h(u, v)$ [854, 864]. Then, we have

$$\mathcal{S}[M_g(u, v)] \simeq \mathcal{S}[M_f(u, v)] M_h(u, v). \quad (10.66)$$

Combining Equation 10.63 and Equation 10.66, we obtain

$$M_f(u, v) \simeq M_g(u, v) \frac{\mathcal{S}[M_f(u, v)]}{\mathcal{S}[M_g(u, v)]}. \quad (10.67)$$

Equation 10.67 suggests that if we can obtain an initial approximation of $M_f(u, v)$, we can rewrite Equation 10.67 in an iterative form, and use it to refine the initial magnitude estimate. Equation 10.67 can thus be written in an iterative form as

$$\tilde{M}_f^{l+1} = M_g \frac{\mathcal{S}[\tilde{M}_f^l]}{\mathcal{S}[M_g]}, \quad (10.68)$$

where l is the iteration number, $\tilde{\ }$ indicates an estimate of the variable under the symbol, and the argument (u, v) has been dropped for compact notation.

The initial estimate \tilde{M}_f^0 may be derived from the phase of the blurred image, that is $\exp[j\theta_g(u, v)]$, which retains most of the high-frequency information (spatial edges) in the image. A simple initial estimate of the original image, in the frequency domain, may be defined as the sum of the phase of the blurred image and the Fourier transform of the blurred image itself:

$$\tilde{M}_f^0 \exp[j\theta_g] = M_g \exp[j\theta_g] + \exp[j\theta_g]; \quad (10.69)$$

in terms of magnitudes, we have

$$\tilde{M}_f^0 = M_g + 1. \quad (10.70)$$

From Equation 10.70, it is evident that a unit constant is being added to the spectral magnitude at all frequencies, which would only raise the entire spectral magnitude response by the corresponding amount. This has an effect comparable to that of adding a high-pass filtered version of the image to the blurred image, which would be of amplifying the high-frequency components in the image. This step would, however, produce a noisy initial approximation of M_f . Rather than simply adding a unit magnitude to recover high frequencies, it would be more appropriate to add those high-frequency components in M_f that were affected by blurring. Although we do not have explicit knowledge of M_f , we do have a ratio between M_f and its smoothed version derived from Equation 10.67, giving

$$\frac{M_f}{\mathcal{S}[M_f]} \simeq \frac{M_g}{\mathcal{S}[M_g]}. \quad (10.71)$$

Adding this ratio to the blurred magnitude spectrum gives

$$\tilde{M}_f^0 = M_g + \frac{M_f}{\mathcal{S}[M_f]}. \quad (10.72)$$

The expression above may be expected to be a more accurate approximation of M_f than that in Equation 10.69, because the amount being added is the ratio $\frac{M_f}{\mathcal{S}[M_f]}$, which contains more information about the original spectral magnitude than a simple constant (unity in Equation 10.70).

The advantages of adding $\frac{M_f}{\mathcal{S}[M_f]}$ to the blurred magnitude spectrum are the following:

- At low frequencies, $M_f \simeq \mathcal{S}[M_f]$. Therefore, $\tilde{M}_f^0 \simeq M_g + 1$, which would not significantly affect the magnitude spectrum.
- At higher frequencies, $\mathcal{S}[M_f] < M_f$ at high amplitudes of M_f , because variations in M_f would be averaged out in $\mathcal{S}[M_f]$. Thus, $\frac{M_f}{\mathcal{S}[M_f]} > 1$ at high-frequency components that are likely to represent spatial edges. Adding this ratio to the blurred magnitude spectrum would have the effect of amplifying high-frequency edges more than the high-frequency noise components. Therefore, the high-frequency components of the phase spectrum $\exp[j\theta_g]$ are emphasized in $\frac{M_f}{\mathcal{S}[M_f]} \exp[j\theta_g]$; for this reason, Rabie et al. referred to the corresponding image as the *enhanced phase* image. Thus, the operation in Equation 10.72 tends to add to M_g the (normalized) high-frequency components that were affected by blurring.

The iterative blind restoration procedure of Rabie et al. may be summarized as follows:

1. Use Equation 10.72 to obtain an initial estimate of M_f .

2. Update the estimate iteratively using Equation 10.68.
3. Stop when the MSE between the l^{th} estimate and the $(l + 1)^{\text{th}}$ estimate is less than a certain limit.
4. Using Equation 10.64, combine the best estimate of $M_f(u, v)$ with the phase function $\theta_f(u, v) = \theta_g(u, v)$. The Fourier transform of the restored image is given as

$$\tilde{F}(u, v) = \tilde{M}_f(u, v) \exp [j\theta_f(u, v)]. \quad (10.73)$$


5. Take the inverse Fourier transform of Equation 10.73 to obtain the deblurred image $\tilde{f}(x, y)$.

Although the method described above neglects noise, it can still be used for deblurring images corrupted by blur and additive noise after first reducing the noise in the blurred image [865, 864].

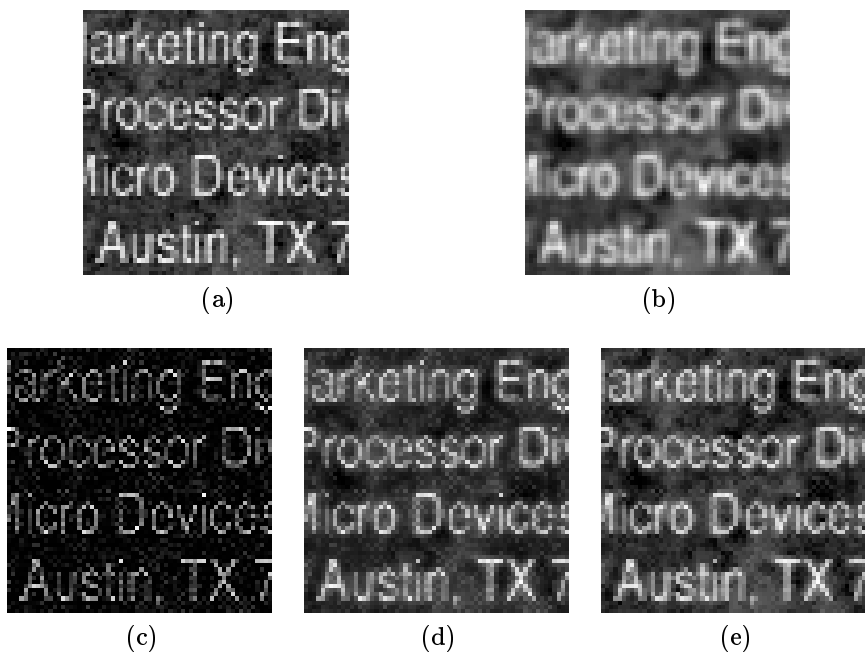
Examples: Figure 10.7 (a) shows the original test image “Lenna”. Part (b) of the figure shows the test image blurred by a Gaussian-shaped PSF with the radial standard deviation $\sigma_r = 3$ pixels. The magnitude of the inverse Fourier transform of the enhanced phase function $\frac{M_f}{S[M_f]} \exp [j\theta_g]$ is shown in part (c) of the figure. It is evident that the enhanced phase image retains most of the edges in the original image that were made weak or imperceptible by blurring. Part (d) of the figure shows the initial estimate of $\tilde{f}(x, y)$, formed by the addition of the image in part (b) of the figure to the image in part (c), as described in Equation 10.72. The addition has emphasized the edges of the blurred image, thus giving a sharper image (albeit with a larger MSE than that of the blurred image, as listed in the caption of Figure 10.7). The dynamic range of the image in part (d) is different from that of the original image in part (a) of the figure.

The image generated by combining the first estimate \tilde{M}_f^1 obtained by using Equation 10.68 with the phase of the blurred image $\exp [j\theta_g]$ is shown in part (e) of the figure. Even after only one iteration, the deblurred image closely resembles the original image. The restored image after four iterations, shown in part (f) of the figure, demonstrates sharp features with minimal artifacts. The smoothing operator \mathcal{S} was defined as the 5×5 mean filter, applied to the 128×128 Fourier spectrum of the image [852, 854, 864].

Figures 10.8 and 10.9 illustrate the restoration of an image with text, blurred to two different levels. The enhanced phase images clearly depict the edge information retained in the phase of the blurred image. The restored images demonstrate significantly improved quality in terms of sharpened edges and improved readability of the text, as well as reduced MSE.

**FIGURE 10.7**

Iterative blind deconvolution with the Lenna test image of size 128×128 pixels and 256 gray levels: (a) original image; (b) blurred image with Gaussian-shaped blur function of radial standard deviation $\sigma_r = 3$ pixels, $\text{MSE} = 606$; (c) enhanced phase image, with the range $[0, 100]$ mapped to the display range $[0, 256]$; (d) initial estimate image, $\text{MSE} = 877$; (e) deblurred image after the first iteration, $\text{MSE} = 128$; (f) deblurred image after four iterations, $\text{MSE} = 110$. Reproduced with permission from T.F. Rabie, R.B. Paranjape, and R.M. Rangayyan, "Iterative method for blind deconvolution", *Journal of Electronic Imaging*, 3(3):245–250, 1994. © SPIE.

**FIGURE 10.8**

Iterative blind deconvolution of a slightly blurred text image, digitized to 64×64 pixels and 256 gray levels: (a) original image; (b) blurred image with Gaussian-shaped blur function of radial standard deviation $\sigma_r = 3$ pixels, $\text{MSE} = 1,041$; (c) enhanced phase image; (d) initial estimate image, $\text{MSE} = 385$; (e) final restored image after eight iterations, $\text{MSE} = 156$. Reproduced with permission from T.F. Rabie, R.B. Paranjape, and R.M. Rangayyan, "Iterative method for blind deconvolution", *Journal of Electronic Imaging*, 3(3):245–250, 1994. © SPIE.

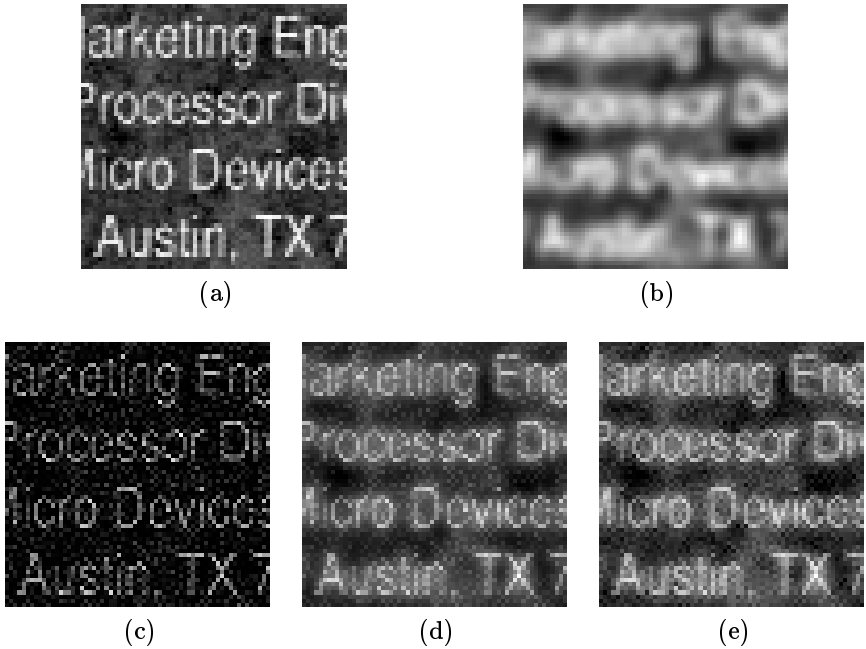


FIGURE 10.9

Iterative blind deconvolution of a severely blurred text image, digitized to 64×64 pixels and 256 gray levels: (a) original image; (b) blurred version of the test image with Gaussian-shaped blur function of radial standard deviation $\sigma_r = 5$ pixels, $MSE = 3,275$; (c) enhanced phase image; (d) initial estimate image, $MSE = 690$; (e) restored image after one iteration, $MSE = 618$. Reproduced with permission from T.F. Rabie, R.B. Paranjape, and R.M. Rangayyan, "Iterative method for blind deconvolution", *Journal of Electronic Imaging*, 3(3):245–250, 1994. © SPIE.

10.3 Homomorphic Deconvolution

Consider the case where we have an image that is given by the convolution of two component images, as expressed by the relation

$$g(x, y) = h(x, y) * f(x, y). \quad (10.74)$$

Similar to the case of the multiplicative homomorphic system described in Section 4.8, the goal in homomorphic deconvolution is to convert the convolution operation to addition. From the convolution property of the Fourier transform, we know that Equation 10.74 translates to

$$G(u, v) = H(u, v) \times F(u, v). \quad (10.75)$$

Thus, the application of the Fourier transform converts convolution to multiplication. Now, it is readily seen that the multiplicative homomorphic system may be applied to convert the multiplication above to addition. Taking the complex logarithm of $G(u, v)$, we have

$$\log[G(u, v)] = \log[H(u, v)] + \log[F(u, v)]; \quad H(u, v) \neq 0, F(u, v) \neq 0 \forall u, v. \quad (10.76)$$

{*Note:* $\log[F(u, v)] = \hat{F}(u, v) = \log[|F(u, v)|] + j\angle F(u, v)$.} A linear filter may now be used to separate the transformed components of f and h , with the assumption as before that they are separable in the transform space. A series of the inverses of the transformations applied initially will take us back to the original domain.

Figure 10.10 gives a block diagram of the steps involved in a homomorphic filter for convolved signals. Observe that the path formed by the first three blocks (the top row) transforms the convolution operation at the input to addition. The set of the last three blocks (the bottom row) performs the reverse transformation, converting addition to convolution. The filter in between thus deals with (transformed) images that are combined by simple addition. Practical application of the homomorphic filter is not simple. Figure 10.11 gives a detailed block diagram of the procedure [7, 239].

10.3.1 The complex cepstrum

The formal definition of the complex cepstrum states that it is the inverse z -transform of the complex logarithm of the z -transform of the input signal [7, 239]. (The name “cepstrum” was derived by transposing the syllables of the word “spectrum”; other transposed terms [7, 236, 239] are less commonly used.) In practice, the Fourier transform is used in place of the z -transform.

Given $g(x, y) = h(x, y) * f(x, y)$, it follows that

$$\hat{G}(u, v) = \hat{H}(u, v) + \hat{F}(u, v), \quad (10.77)$$

and furthermore, that the complex cepstra of the signals are related simply as

$$\hat{g}(x, y) = \hat{h}(x, y) + \hat{f}(x, y). \quad (10.78)$$

Here, the $\hat{}$ symbol over a function of frequency indicates the complex logarithm of the corresponding function, whereas the same symbol over a function of space indicates its complex cepstrum.

An important consideration in the evaluation of the complex logarithm of $G(u, v)$ relates to the phase of the signal. The phase spectrum computed as its principal value in the range $[0, 2\pi]$, given by $\tan^{-1} \left[\frac{\text{imaginary}\{G(u, v)\}}{\text{real}\{G(u, v)\}} \right]$, will almost always have discontinuities that will conflict with the requirements of the inverse Fourier transform to follow. Thus, $G(u, v)$ needs to be separated into its magnitude and phase components, the logarithmic operation applied to the magnitude, the phase corrected to be continuous by adding correction factors of $\pm 2\pi$ at discontinuities larger than π , and the two components combined again before the subsequent inverse transformation. Correcting the phase spectrum as above is referred to as *phase unwrapping* [239, 7, 866, 867, 868, 869]. In addition, it has been shown that a linear phase term, if present in the spectrum of the input, may cause rapidly decaying oscillations in the complex cepstrum [239]. It is advisable to remove the linear phase term, if present, during the phase-unwrapping step. The linear phase term may be added to the filtered result, if necessary.

Taxt [870] applied 2D homomorphic filtering to improve the quality of medical ultrasound images. The recorded image was modeled as the convolution of a field representing the sound-reflecting structures (the anatomical surfaces of interest) with a PSF related to the interrogating ultrasound pulse shape. A Wiener filter was then used to deconvolve the effects of the PSF.

An important application of 1D homomorphic filtering is in the extraction of the basic wavelet from a composite signal made up of quasiperiodic repetitions (or echoes) of the wavelet, such as a voiced-speech signal or seismic signal [7, 31, 237, 238, 871, 872]. An extension of this technique to the removal of visual echoes in images is described in Section 10.3.2.

10.3.2 Echo removal by Radon-domain cepstral filtering

Homomorphic deconvolution has been used to remove repeated versions of a basic pattern in an image (known as a visual echo) [444, 505] and other applications in image processing [873, 874, 875, 876, 877]. Martins and Rangayyan [444, 505] performed 1D homomorphic deconvolution on the Radon transform (projections) of the given image, instead of 2D homomorphic filtering: the final filtered image was obtained by applying the ART method of reconstruction from projections (see Section 9.4) to the filtered projections. The general scheme of the method is shown in Figure 10.12; the details of the method are described in the following paragraphs.

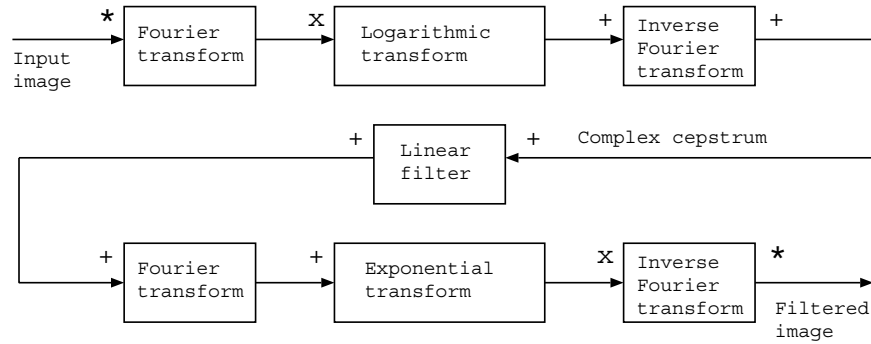


FIGURE 10.10

Operations involved in a homomorphic filter for convolved signals. The symbol at the input or output of each block indicates the operation that combines the signal components at the corresponding step. Reproduced with permission from R.M. Rangayyan, *Biomedical Signal Analysis: A Case-Study Approach*, IEEE Press and Wiley, New York, NY, 2002. © IEEE.

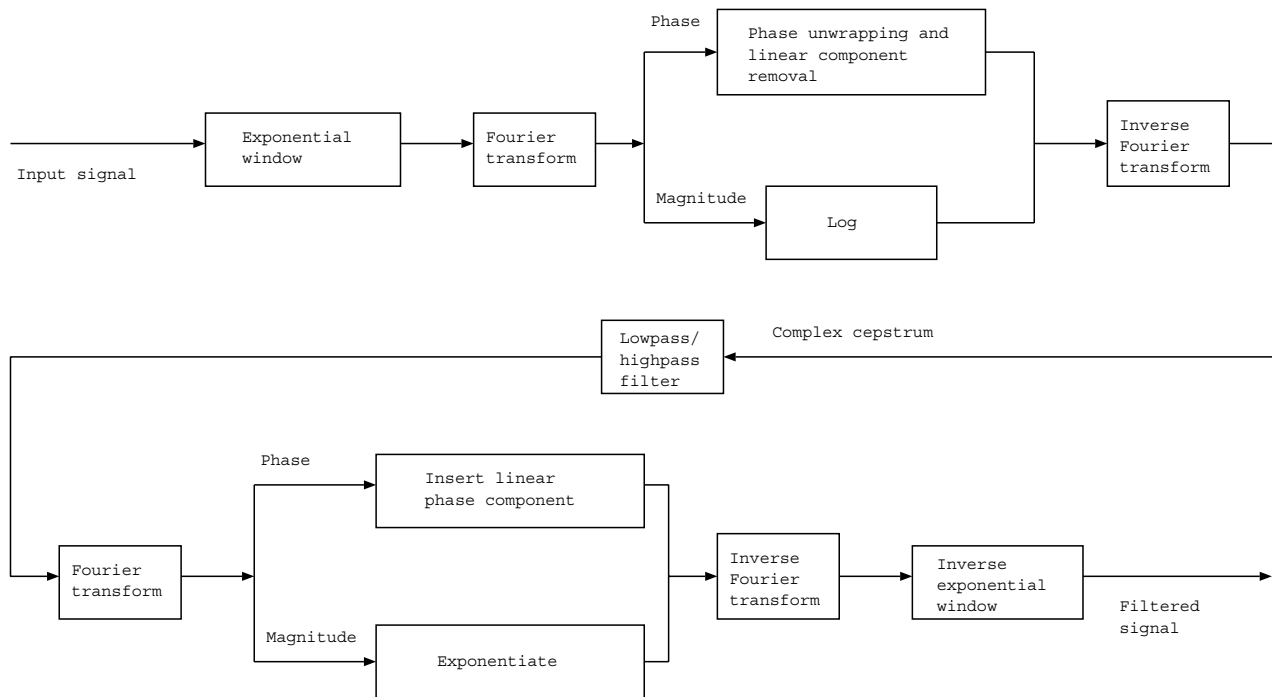


FIGURE 10.11

Detailed block diagram of the steps involved in deconvolution of signals using the complex cepstrum. Reproduced with permission from A.C.G. Martins and R.M. Rangayyan, “Complex cepstral filtering of images and echo removal in the Radon domain”, *Pattern Recognition*, 30(11):1931–1938, 1997. © Pattern Recognition Society. Published by Elsevier Science Ltd.




FIGURE 10.12

Homomorphic (complex cepstral) filtering of an image in the Radon domain. Reproduced with permission from A.C.G. Martins and R.M. Rangayyan, “Complex cepstral filtering of images and echo removal in the Radon domain”, *Pattern Recognition*, 30(11):1931–1938, 1997. © Pattern Recognition Society. Published by Elsevier Science Ltd.

Let $f_e(x, y)$ be an image given by

$$f_e(x, y) = f(x, y) * d(x, y), \tag{10.79}$$

where

$$d(x, y) = \delta(x, y) + a \delta(x - x_0, y - y_0), \tag{10.80}$$

with a being a scalar weighting factor. In the context of images with visual echoes, we could label $f(x, y)$ as the basic element image, $d(x, y)$ as a field of impulses at the positions of the echoes (including the original element), and $f_e(x, y)$ as the composite image. Applying the Radon transform (see Section 9.1), we have the projection $p_\theta(t)$ of the composite image $f_e(x, y)$ at angle θ given by

$$p_\theta(t) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f_e(x, y) \delta(x \cos \theta + y \sin \theta - t) dx dy, \tag{10.81}$$

which leads to

$$\begin{aligned} p_\theta(t) &= \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(\alpha, \beta) \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \delta(x - \alpha, y - \beta) \\ &\times \delta(x \cos \theta + y \sin \theta - t) dx dy d\alpha d\beta \\ &+ a \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(\alpha, \beta) \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \delta[(x - x_0) - \alpha, (y - y_0) - \beta] \\ &\times \delta(x \cos \theta + y \sin \theta - t) dx dy d\alpha d\beta. \end{aligned} \tag{10.82}$$

Here, t is the displacement between the projection samples (rays) in the Radon domain. Using the properties of the δ function (see Section 2.9), we get

$$\begin{aligned} p_\theta(t) &= \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(\alpha, \beta) \delta(\alpha \cos \theta + \beta \sin \theta - t) d\alpha d\beta \\ &+ a \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(\alpha, \beta) \delta[(x_0 + \alpha) \cos \theta + (y_0 + \beta) \sin \theta - t] d\alpha d\beta. \end{aligned} \tag{10.83}$$

Defining

$$n_\theta = x_0 \cos \theta + y_0 \sin \theta, \quad (10.84)$$

and

$$f_\theta(t) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(\alpha, \beta) \delta(\alpha \cos \theta + \beta \sin \theta - t) d\alpha d\beta, \quad (10.85)$$

we get

$$p_\theta(t) = f_\theta(t) + a f_\theta(t - n_\theta). \quad (10.86)$$

Here, $f_\theta(t)$ represents the projection of the basic element image $f(x, y)$ at angle θ , and n_θ is a displacement or shift (in the Radon domain) for the given angle θ and position (x_0, y_0) . Equation 10.86 indicates that the projection of the composite image at a given angle is the summation of the projections at the same angle of the basic element image and its replication (or echo) with an appropriate shift factor n_θ .

In order to demonstrate the effect of an echo in a signal on its complex cepstrum, we could apply the Fourier transform, the log operation, and then the inverse Fourier transform, as follows: Applying the Fourier transform to Equation 10.86, we get

$$P_\theta(w) = F_\theta(w) + a \exp(-j 2\pi w n_\theta) F_\theta(w), \quad (10.87)$$

where $P_\theta(w)$ and $F_\theta(w)$ are the Fourier transforms of $p_\theta(t)$ and $f_\theta(t)$, respectively, and w is the Fourier-domain variable corresponding to the space-domain (Radon-domain) variable t . Applying the natural log, we get

$$\hat{P}_\theta(w) = \hat{F}_\theta(w) + \ln[1 + a \exp(-j 2\pi w n_\theta)], \quad (10.88)$$

where the $\hat{}$ represents the log-transformed version of the variable under the symbol. If $a < 1$, the log function may be expanded into a power series as

$$\hat{P}_\theta(w) = \hat{F}_\theta(w) + a \exp(-j 2\pi w n_\theta) - \frac{[a \exp(-j 2\pi w n_\theta)]^2}{2} + \dots \quad (10.89)$$

Applying the inverse Fourier transform, we get the complex cepstrum of $p_\theta(t)$ as

$$\hat{p}_\theta(t) = \hat{f}_\theta(t) + a \delta(t - n_\theta) - \frac{a^2}{2} \delta(t - 2n_\theta) + \dots \quad (10.90)$$

Thus, the complex cepstrum $\hat{p}_\theta(t)$ of the projection $p_\theta(t)$ at angle θ of the image $f_e(x, y)$ is composed of the complex cepstrum $\hat{f}_\theta(t)$ of the projection at angle θ of the basic element image $f(x, y)$, and a train of impulses. If $\hat{f}_\theta(t)$ and the impulse train are sufficiently separated in the cepstral domain, the use of a lowpass filter, followed by the inverse cepstrum operation, gives the filtered projection $\hat{f}_\theta(t)$ of the basic element image. The impulse train may also be suppressed by using a notch or a comb function. After the filtered projections are obtained at several angles, it will be possible to reconstruct

the basic element image $f(x, y)$. If the effects of the echoes are considered to be a type of image distortion, filtering as above may be considered to be an operation for image restoration.

If $a \geq 1$, that is, if the echoes are of amplitude equal to or greater than that of the basic element image, the impulses in the cepstrum will appear with negative delays, and the filtering operation will lead to the extraction of an echo image [505]. The situation may be modified easily by applying decaying exponential weighting factors to the original image and/or the projections, such that the weighted echoes and/or their projections are attenuated to lower amplitudes. Then, the effect is equivalent to the situation with $a < 1$.

Example: A composite image with five occurrences of a simple circular object is shown in Figure 10.13 (a). If the object at the lower-right corner of the image is considered to be the original object, the remaining four objects could be viewed as echoes of the original object. Although echoes would, in general, be of lower amplitude (intensity) than the original object, they have been maintained at the same intensity as the original in this image. The test image is of size 101×101 pixels; the radius of each circle is 10 pixels, and the intensity of each circle is 100 on an inverted gray scale.

The Radon-domain homomorphic filter was applied to the test image. The image was multiplied by a weighting function given by y^3 , where y is the vertical axis of the image and the origin is at the top-left corner of the image. Furthermore, another weighting function, given by α^t with $\alpha = 0.985$, was applied to each projection before computing the complex cepstrum. The weighting functions were used in order to reduce the effect of the echoes and facilitate filtering of the cepstrum [239, 31]. The cepstrum was lowpass filtered with a window of 40 pixels. Eighteen projections in the range $[0^\circ, 170^\circ]$ in steps of 10° were computed, filtered, and used to reconstruct the basic element image. The result, shown in Figure 10.13 (b), was thresholded at the gray level of 100. It is seen that a single circular object has been extracted, with minimal residual artifact.

The method was extended to the extraction of the basic element in images with periodic texture, known as the texton, by Martins and Rangayyan [444]; see Section 7.7.1.

10.4 Space-variant Restoration

Several image restoration techniques, such as the Wiener and PSE filters, are based upon the assumption that the image can be modeled by a stationary (random) field. Restoration is achieved by filtering the degraded image with an LSI filter, the frequency response of which is a function of the PSD of the uncorrupted image. There are at least two difficulties in this approach: most




FIGURE 10.13

Illustration of homomorphic (complex cepstral) filtering of an image in the Radon domain to remove echoes in images. (a) Original image. (b) Filtered image after thresholding. Reproduced with permission from A.C.G. Martins and R.M. Rangayyan, “Complex cepstral filtering of images and echo removal in the Radon domain”, *Pattern Recognition*, 30(11):1931–1938, 1997. © Pattern Recognition Society. Published by Elsevier Science Ltd.

images are not stationary, and, at best, may be described as locally stationary; furthermore, in practice, the PSD of the uncorrupted original image is not given, and needs to be estimated.

A common procedure to estimate the PSD of a signal involves sectioning the signal into smaller, presumably stationary segments, and averaging their modified PSDs or periodograms [12, 825, 853, 862, 863, 865, 878, 879, 880, 881, 882]. The procedure assumes shift-invariance of the blur PSF, and averages out the nonstationary frequency content of the original image. In order for the sectioning approach to be valid, the blurring PSF must have an ROS (spatial extent) that is much smaller than the size of the subimages; as a result, the size of the subimages cannot be made arbitrarily small. Thus, the number of subimages that may be used to form the ensemble average is limited; consequently, the variance of the PSD estimate from the subimages could be high, which leads to a poor estimate of the PSD of the original image.

Another consequence of the assumption of image stationarity and the use of space-invariant filtering is the fact that the deblurred images suffer from artifacts at the boundaries of the image [853]. In a simple mathematical representation of this problem, it is assumed that the image is of infinite extent; however, practical images are of finite extent and the performance of a filter may vary depending upon the assumptions made regarding the edges of the image. The effect at the edges from deconvolution with incomplete information (due to the lack of information beyond the image boundary) could cause

different contributions from outside the image boundary during deblurring than those that were convolved into the image during the actual degradation process. This leads to a layer of boundary pixels taking incorrect values during deblurring, and consequently, artifacts at the image boundaries.

Attempts to overcome the problems caused by the inherent nonstationarity of images have led to several methods. Angel and Jain [883] proposed a technique to solve the general superposition integral iteratively by using a conjugate-gradient method. The method faces problems with convergence in the presence of noise. Techniques have been proposed to enhance the performance of nonadaptive filters by using radiometric and geometric transforms to generate images that are nearly stationary (block stationary) in the first and second moments [884]. The radiometric transform generates stationary mean and variance, whereas the geometric transform provides stationary autocorrelation.

Adaptive techniques for space-variant restoration have been proposed based upon sectioning the given image into smaller subsections and assuming different stationary models for each section [177, 865, 885, 886]. Two approaches to sectioned deblurring are described in Sections 10.4.1 and 10.4.2.

The Kalman filter is the most popular method for truly space-variant filtering, and is described in Section 10.4.3.

10.4.1 Sectioned image restoration

In the iterative sectioned MAP restoration technique proposed by Trussell and Hunt [177, 885], the input image is divided into small $P \times P$ sections and the MAP estimate of the uncorrupted section is developed and iterated upon for refinement. This procedure is carried out on each section using an overlap-save technique to reduce edge artifacts. Because sectioning an image presumably causes each individual section to be close to a stationary process, a simpler approach to sectioned deblurring could be based upon the use of a conventional LSI filter, such as the Wiener or PSE filter, to deblur each section individually; the deblurred sections may then be combined to form the final deblurred image. A technique to reduce edge effects between the sections is to center each subimage in a square region of size comparable to that of the input image, and then pad the region surrounding the centered subimage with its mean value, prior to filtering. Each mean-padded region may also be multiplied in the space domain with a smooth window function, such as the Hamming window.

Using the above argument and assuming that each section (of size $P \times P$ pixels) is large compared to the ROS of the blur PSF, but small compared to the actual image dimensions $M \times M$, each section of the image may be expressed as the convolution of the PSF with an equivalent section from the original undegraded image $f(m, n)$. Then, we have

$$g_l(m, n) \simeq h(m, n) * f_l(m, n) + \eta_l(m, n). \quad (10.91)$$

In the frequency domain, Equation 10.91 becomes

$$G_l(u, v) \simeq H(u, v) F_l(u, v) + \eta_l(u, v). \quad (10.92)$$

Now, applying the 2D Hamming window of size $M \times M$, given by

$$w_H(m, n) = \left[0.54 - 0.46 \cos \left(\frac{2 \pi m}{M-1} \right) \right] \left[0.54 - 0.46 \cos \left(\frac{2 \pi n}{M-1} \right) \right], \quad (10.93)$$

to each region padded by its mean to the size $M \times M$, we obtain

$$g_l(m, n) w_H(m, n) \simeq [h(m, n) * f_l(m, n)] w_H(m, n) + \eta_l(m, n) w_H(m, n), \quad (10.94)$$

or

$$g_l^w(m, n) \simeq h(m, n) * f_l^w(m, n) + \eta_l^w(m, n), \quad (10.95)$$

where the w notation represents the corresponding windowed sections. The PSD of each section $g_l(m, n)$ may be expressed as

$$\Phi_{g_l}(u, v) \simeq |H(u, v)|^2 \Phi_{f_l}(u, v) + \Phi_{\eta_l}(u, v). \quad (10.96)$$

The Wiener or PSE filter may then be applied to each section. The final restored image $\hat{f}(m, n)$ is obtained by extracting the individual sections from the center of the corresponding filtered mean-padded images, and placing them at their original locations.

Limitations exist in sectioned restoration due to the fact that the assumption of stationarity within the square sections may not be satisfied: sectioning using regularly spaced square blocks of equal size cannot discriminate between “flat” and “busy” areas of any given image. Furthermore, because of the limitations on the section size (that it must be large compared to the ROS of the blur PSF), sections cannot be made arbitrarily small. Thus, the mean value of a section could be significantly different from its pixel values, and consequently, artifacts could arise at section boundaries. To partially solve the problem of edge artifacts, the sections could be overlapped, for example, by one-half the section size in each dimension [862, 865]. This technique, however, will not reduce the effects at the image boundaries. An adaptive-neighborhood approach to address this limitation is described in Section 10.4.2, along with examples of application.

10.4.2 Adaptive-neighborhood deblurring

Rabie et al. [178] proposed an adaptive-neighborhood deblurring (AND) algorithm based on the use of adaptive-neighborhood regions determined individually for each pixel in the input image. (See Section 3.7.5 for a description of adaptive-neighborhood filtering.) In the AND approach, the image is treated as being made up of a collection of regions (features or objects) of relatively uniform gray levels. An adaptive neighborhood is determined for each pixel

in the image (called the seed pixel when it is being processed), being defined as the set of pixels 8-connected to the seed pixel and having a difference in gray level with respect to the seed that is within specified limits of tolerance. The tolerance used in AND is an additive factor, as given by Equation 3.164. Thus, the tolerance determines the maximum allowed deviation in the gray level from the seed pixel value within each adaptive neighborhood, and any deviation less than the tolerance is considered to be an intrinsic property of the adaptive-neighborhood region. The number of pixels in an adaptive-neighborhood region may be limited by a predetermined number Q ; however, there are no restrictions on the shape of the adaptive-neighborhood regions.

Assuming that each adaptive-neighborhood region grown is large compared to the ROS of the PSF, each such region may be expressed as the convolution of the PSF with an equivalent adaptive-neighborhood region grown in the original undegraded image $f(m, n)$. Thus, similar to Equation 10.91, we have

$$g_{m,n}(p, q) \simeq h(p, q) * f_{m,n}(p, q) + \eta_{m,n}(p, q), \quad (10.97)$$

where (m, n) is the seed pixel location for which the adaptive-neighborhood region $g_{m,n}(p, q)$ was grown, and (p, q) give the locations of the pixels within the region. It is assumed that regions corresponding to $g_{m,n}(p, q)$ may be identified in the original image and noise fields.

Next, each adaptive-neighborhood region is centered within a rectangular region of the same size as the input image ($M \times M$); the area surrounding the region is padded with its mean value in order to reduce edge artifacts and to enable the use of the 2D FFT. Thus, in the frequency domain, Equation 10.97 becomes

$$G_{m,n}(u, v) \simeq H(u, v) F_{m,n}(u, v) + \eta_{m,n}(u, v). \quad (10.98)$$

Applying the 2D Hamming window $w_H(p, q)$ (see Equation 10.93) to each mean-padded adaptive-neighborhood region, we obtain

$$g_{m,n}(p, q) w_H(p, q) \simeq [h(p, q) * f_{m,n}(p, q)] w_H(p, q) + \eta_{m,n}(p, q) w_H(p, q), \quad (10.99)$$

or

$$g_{m,n}^w(p, q) \simeq h(p, q) * f_{m,n}^w(p, q) + \eta_{m,n}^w(p, q), \quad (10.100)$$

where w represents the corresponding windowed regions. The PSD of the region $g_{m,n}(p, q)$ may be expressed as

$$\Phi_{g_{m,n}}(u, v) \simeq |H(u, v)|^2 \Phi_{f_{m,n}}(u, v) + \Phi_{\eta_{m,n}}(u, v). \quad (10.101)$$

In deriving the AND filter, the stationarity of the adaptive-neighborhood regions grown is taken into account. An estimate of the spectrum of the noise $\eta_{m,n}(u, v)$, within the current adaptive-neighborhood region grown from the seed pixel at (m, n) , is obtained as

$$\tilde{\eta}_{m,n}(u, v) = A_{m,n}(u, v) G_{m,n}(u, v), \quad (10.102)$$

where $A_{m,n}(u, v)$ is a frequency-domain, magnitude-only scale factor that depends on the spectral characteristics of the adaptive-neighborhood region grown. An estimate of $F_{m,n}(u, v)$ is obtained from Equation 10.98 by using the spectral estimate of the noise, $\tilde{\eta}_{m,n}(u, v)$, in place of $\eta_{m,n}(u, v)$. Then, we have

$$\tilde{F}_{m,n}(u, v) = \frac{G_{m,n}(u, v) - \tilde{\eta}_{m,n}(u, v)}{H(u, v)}, \quad (10.103)$$

which reduces to

$$\tilde{F}_{m,n}(u, v) = \frac{G_{m,n}(u, v)}{H(u, v)} [1 - A_{m,n}(u, v)]. \quad (10.104)$$

The spectral noise estimator $A_{m,n}(u, v)$ may be derived by requiring the PSD of the estimated noise $\Phi_{\tilde{\eta}}(u, v)$ to be equal to the original noise PSD $\Phi_{\eta}(u, v)$ for the current adaptive-neighborhood region. Thus, using Equation 10.102, we can describe the relationship between the noise PSD and the image PSD as

$$\Phi_{\eta_{m,n}}(u, v) = A_{m,n}^2(u, v) \Phi_{g_{m,n}}(u, v). \quad (10.105)$$

From Equation 10.101 and Equation 10.105, the spectral noise estimator $A_{m,n}(u, v)$ is given by

$$A_{m,n}(u, v) = \left(\frac{\Phi_{\eta}(u, v)}{|H(u, v)|^2 \Phi_{f_{m,n}}(u, v) + \Phi_{\eta}(u, v)} \right)^{1/2}, \quad (10.106)$$

where $\Phi_{\eta_{m,n}}(u, v) = \Phi_{\eta}(u, v) = \sigma_{\eta}^2$ for Gaussian white noise. The quantity in Equation 10.101 gives the denominator in Equation 10.106. Therefore, no additional information is required about the PSD of the original undegraded image.

The frequency-domain estimate of the uncorrupted adaptive-neighborhood region is obtained by using the value of $A_{m,n}(u, v)$, computed from Equation 10.106, in Equation 10.104. The spectral estimate of the original undegraded adaptive-neighborhood region is thus given by

$$\begin{aligned} \tilde{F}_{m,n}(u, v) &= \frac{G_{m,n}(u, v)}{H(u, v)} \left[1 - \left(\frac{\Phi_{\eta}(u, v)}{|H(u, v)|^2 \Phi_{f_{m,n}}(u, v) + \Phi_{\eta}(u, v)} \right)^{1/2} \right] \\ &= \frac{G_{m,n}(u, v)}{H(u, v)} \left[1 - \left(\frac{\Phi_{\eta}(u, v)}{\Phi_{g_{m,n}}(u, v)} \right)^{1/2} \right]. \end{aligned} \quad (10.107)$$

The space-domain estimate of the uncorrupted adaptive-neighborhood region $f_{m,n}(p, q)$ is obtained from the inverse Fourier transform of the expression above. By replacing the seed pixel at (m, n) with the deblurred pixel $\tilde{f}_{m,n}(m, n)$, and running the algorithm above for every pixel in the input image, we will obtain a deblurred image based on stationary adaptive regions.

A computational disadvantage of the algorithm above is the fact that it requires two $M \times M$ 2D FFT operations per pixel. An approach to circumvent this difficulty, to some extent, is provided by the intrinsic nature of the adaptive neighborhood: Because most of the pixels inside an adaptive-neighborhood region will have similar adaptive-neighborhood regions when they become seed pixels (because they lie within similar limits of tolerance), instead of growing an adaptive-neighborhood region for each pixel in the input image, we could grow adaptive-neighborhood regions only from those pixels that do not already belong to a previously grown region. Thus, after filtering an adaptive-neighborhood region, the entire adaptive-neighborhood region may be placed in the output image at the seed pixel location, instead of replacing only the single restored seed pixel. Note that the various adaptive-neighborhood regions grown as above could still overlap. This approach is a compromise to reduce the computational requirements.

Examples: The test image “Lenna”, of size 128×128 pixels and 256 gray levels, is shown in Figure 10.14 (a). The image, after degradation by a Gaussian-shaped blur PSF with a radial standard deviation $\sigma_r = 3$ pixels and additive white Gaussian noise to 35 dB SNR, is shown in part (b) of the figure. Parts (c) – (e) of the figure show three different fixed-neighborhood sections of the blurred image. Each section is of size 32×32 pixels, and is centered in a square region of the same size as the full image (128×128), with the surrounding area padded with the mean value of the section. Part (f) shows the windowed version of the region in part (e) after weighting with a Hamming window. It is evident from the images in parts (c) – (e) that the assumption of stationarity within a given section is not satisfied: each section contains a variety of image characteristics. The values of the mean-padded areas are also significantly different from the pixel values of the corresponding centered sections.

Three different adaptive-neighborhood regions of the blurred test image are shown in Figure 10.15 (c) – (e). Each adaptive-neighborhood region was allowed to grow to any size as long as the pixel values were within an adaptive tolerance given by $\frac{\sigma_g}{\sqrt{2}}$, where σ_g is an estimate of the standard deviation of the noise-free blurred image $g(m, n) = h(m, n) * f(m, n)$. Each adaptive-neighborhood region was centered in a square region of the same size as the full image (128×128), and the surrounding area was padded with the mean value of the region. Unlike the fixed square sections shown in Figure 10.14 (c) – (e), the adaptive-neighborhood regions in Figure 10.15 (c) – (e) do not contain large spatial fluctuations such as high-variance edges, but rather slow-varying and relatively smooth details. Thus, we may assume that each adaptive-neighborhood region approximates a stationary region in the input image. From Figure 10.15 (c) – (e), it should also be observed that the adaptive-neighborhood regions overlap.

Two restored versions of the blurred Lenna image, obtained using the sectioned Wiener filter with sections of size 16×16 and 32×32 pixels, with the

adjacent sections overlapped by one-half of the section size in both dimensions, are shown in Figure 10.16 (c) and (d), respectively. Overlapping the sections has effectively suppressed the artifacts associated with the inner sections of the image; however, it does not reduce the artifacts at the boundaries of the image because of the lack of information beyond the image boundaries.

Figure 10.16 (e) shows the test image deconvolved using the full image frame with the Wiener filter. In this result, less noise smoothing has been achieved as compared to the restored images using sections of size 32×32 or 16×16 pixels. This could be due to the nonstationarity of the full-frame image used to calculate the PSD required in the Wiener filter equation.

The restored image obtained using the AND filter of Equation 10.107 is shown in Figure 10.16 (f). The result has almost no edge artifacts, which is due to the overlapping adaptive-neighborhood regions used. The AND result has the lowest MSE of all of the results shown, as listed in Table 10.1.

The “Cameraman” test image, of size 128×128 pixels and 256 gray levels, is shown in Figure 10.17 (a). Part (b) of the figure shows the image after degradation by a PSF representing 9-pixel horizontal motion blurring and additive noise to $SNR = 35$ dB. Two restored images obtained using the sectioned Wiener filter with sections of size 16×16 and 32×32 pixels, with the adjacent sections overlapped by one-half of the section size in both dimensions, are shown in Figure 10.17 (c) and (d), respectively. Edge artifacts are apparent in these deblurred images; the artifacts are more pronounced around the vertical edges in the image than around the horizontal edges. This is due to the shape of the blur PSF, which is a 1D function along the horizontal axis.

The result of deconvolution using the full image frame and the Wiener filter is shown in Figure 10.17 (e). It is evident that edge artifacts do not exist within the boundaries of this result.


The test image restored by the application of the AND filter of Equation 10.107 is shown in Figure 10.17 (f). Almost no edge artifact is present in this image due to the use of adaptive-neighborhood regions. The image is sharper and cleaner than the other results shown in Figure 10.17; the AND result also has the lowest MSE, as listed in Table 10.1.

The results demonstrate that image stationarity plays an important role in the overall performance of restoration filters. The use of adaptive-neighborhood regions can improve the performance of restoration filters.

10.4.3 The Kalman filter

The Kalman filter is a popular approach to characterize dynamic systems in terms of state-space concepts [833, 887, 888, 889, 890]. The Kalman filter formulation could be used for filtering, restoration, prediction, and interpolation (smoothing).

Formulation of the Kalman filter: In the Kalman filter, the signals or items of information involved are represented as a state vector $\mathbf{f}(n)$ and an observation vector $\mathbf{g}(n)$, where n refers to the instant of time, or is an index

**FIGURE 10.14**

Sectioning of the Lenna image of size 128×128 pixels and gray-level range of $0 - 255$. (a) Original image. (b) Blurred image with a Gaussian-shaped blur function and noise to $SNR = 35$ dB; $MSE = 607$. (c), (d), and (e): Three 32×32 sections mean-padded to 128×128 pixels. (f) Hamming-windowed version of the region in (e). Reproduced with permission from T.F. Rabie, R.M. Rangayyan, and R.B. Paranjape, "Adaptive-neighborhood image deblurring", *Journal of Electronic Imaging*, 3(4): 368 - 378, 1994. © SPIE.





FIGURE 10.15

Adaptive-neighborhood segmentation of the Lenna image of size 128×128 pixels and gray-level range of $0 - 255$. (a) Original image. (b) Blurred image with a Gaussian-shaped blur function and noise to $SNR = 35 \text{ dB}$, $MSE = 607$. (c), (d), and (e): Three adaptive-neighborhood mean-padded regions. (f) Hamming-windowed version of the region in (e). Reproduced with permission from T.F. Rabie, R.M. Rangayyan, and R.B. Paranjape, "Adaptive-neighborhood image deblurring", *Journal of Electronic Imaging*, 3(4): 368 - 378, 1994. © SPIE.

**FIGURE 10.16**

(a) Lenna test image. (b) Blurred image with a Gaussian-shaped blur function and noise to $SNR = 35$ dB, $MSE = 607$. Sectioned deblurring with overlapped sections of size (c) 16×16 pixels, $MSE = 783$; and (d) 32×32 pixels, $MSE = 501$. (e) Full-frame Wiener filtering, $MSE = 634$. (f) Adaptive-neighborhood deblurring, $MSE = 292$. Reproduced with permission from T.F. Rabie, R.M. Rangayyan, and R.B. Paranjape, "Adaptive-neighborhood image deblurring", *Journal of Electronic Imaging*, 3(4): 368 – 378, 1994. © SPIE.




FIGURE 10.17

(a) Cameraman test image of size 128×128 pixels and gray-level range 0–255. (b) Image blurred by 9-pixel horizontal motion and degraded by additive Gaussian noise to $SNR = 35$ dB, $MSE = 1,247$. Deblurred images: (c) Sectioned deblurring with overlapped sections of size 16×16 pixels, $MSE = 539$. (d) Sectioned deblurring with overlapped sections of size 32×32 pixels, $MSE = 424$. (e) Full-frame Wiener filtering, $MSE = 217$. (f) Adaptive-neighborhood deblurring, $MSE = 181$. Reproduced with permission from T.F. Rabie, R.M. Rangayyan, and R.B. Paranjape, “Adaptive-neighborhood image deblurring”, *Journal of Electronic Imaging*, 3(4): 368 – 378, 1994. © SPIE.

TABLE 10.1

Mean-squared Errors of the Results of Sectioned and Adaptive-neighborhood Deblurring of the Lenna and Cameraman Images of Size 128×128 Pixels and 256 Gray Levels for Various Neighborhood Sizes and Two Different Blurring Functions, and Approximate Computer Processing Time Using a SUN/Sparc-2 Workstation.

Filter	Section size	Time (minutes)	Mean-squared error	
			Lenna	Cameraman
Degraded	—	—	607	1,247
Wiener	16×16	25	783	539
PSE	16×16		751	538
Wiener	32×32	10	501	424
PSE	32×32		513	425
Wiener	64×64	5	483	463
PSE	64×64		488	460
Wiener	Full frame	2	634	217
PSE	Full frame		605	220
Adaptive neighborhood	max = 16,384	15	292	181

Reproduced with permission from T.F. Rabie, R.M. Rangayyan, and R.B. Paranjape, "Adaptive-neighborhood image deblurring", *Journal of Electronic Imaging*, 3(4): 368 – 378, 1994. © SPIE.

related to the sequence of the state or observation vectors. It is assumed that the input is generated by a driving (or process) noise source $\eta_d(n)$, and that the output is affected by an observation noise source $\eta_o(n)$. A state transition matrix $\mathbf{a}(n+1, n)$ is used to indicate the modification of the state vector from one instant n to the next instant $(n+1)$. [Note: The argument in the notation $\mathbf{a}(n+1, n)$ is used to indicate the dependence of the variable (matrix) \mathbf{a} upon the instants of observation n and $(n+1)$, and not the indices of the matrix \mathbf{a} . This notation also represents the memory of the associated system.] An observation matrix $\mathbf{h}(n)$ is used to represent the mapping from the state vector to the observation vector.

Figure 10.18 illustrates the concepts described above in a schematic form. A state vector could represent a series of the values of a 1D signal, or a collection of the pixels of an image in a local ROS related to the pixel being processed; see Figure 10.19 for an illustration of two commonly used types of ROS in image filtering. The state vector should be composed of the minimal amount of data that would be adequate to describe the behavior of the system. As we have seen in Section 3.5, images may be represented using vectors, and image filtering or transformation operations may be expressed as multiplication with matrices. The state transition matrix $\mathbf{a}(n+1, n)$ could represent a linear prediction or autoregressive model (see Section 11.8) that characterizes the input state. The state vector would then be composed of a series of the signal or image samples that would be related to the order of the model, and be adequate to predict the subsequent values of the state and observation vectors (with minimal error). The observation matrix $\mathbf{h}(n)$ could represent a blurring PSF that degrades a given image. Observe that the Kalman filter formulation permits the representation of the signals and their statistics, as well as the operators affecting the signals, as functions that are nonstationary, dynamic, or varying with time (or space).

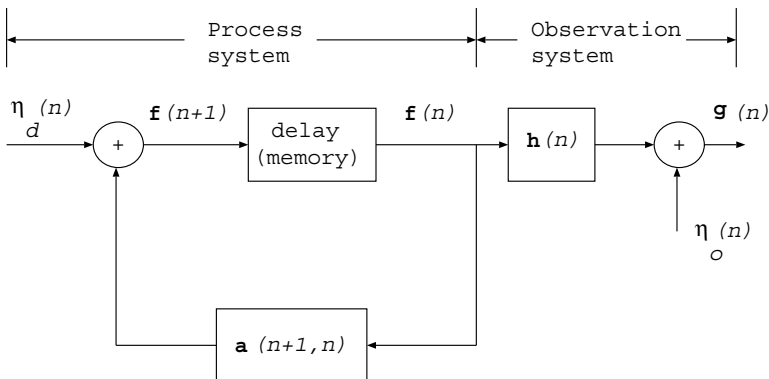


FIGURE 10.18

State-space representation of the basic Kalman filter formulation.

The Kalman filter formulation [833] represents a new or updated value of the state vector \mathbf{f} recursively in terms of its previous value and new input as

$$\mathbf{f}(n+1) = \mathbf{a}(n+1, n)\mathbf{f}(n) + \boldsymbol{\eta}_d(n). \quad (10.108)$$

This is known as the process equation; the corresponding model is also known as the plant, process, or message model. If we let the state vector \mathbf{f} be of size $M \times 1$, then the state transition matrix \mathbf{a} is of size $M \times M$, and the driving noise vector $\boldsymbol{\eta}_d$ is of size $M \times 1$. We may interpret the driving noise vector $\boldsymbol{\eta}_d$ as the source of excitation of the process system represented by the state vector \mathbf{f} and the state transition matrix \mathbf{a} . It is assumed that the noise process $\boldsymbol{\eta}_d$ is a zero-mean white-noise process that is statistically independent of the stochastic process underlying the state vector \mathbf{f} . The noise process $\boldsymbol{\eta}_d$ is characterized by its $M \times M$ correlation matrix (ACF) ϕ_{η_d} as

$$E[\boldsymbol{\eta}_d(n)\boldsymbol{\eta}_d^T(k)] = \begin{cases} \phi_{\eta_d}(n) & \text{if } n = k \\ \mathbf{0} & \text{otherwise.} \end{cases} \quad (10.109)$$

The state transition matrix $\mathbf{a}(n+1, n)$ is characterized by the following properties [833]:

$$\begin{aligned} \mathbf{a}(l, m)\mathbf{a}(m, n) &= \mathbf{a}(l, n) && \text{product rule;} \\ \mathbf{a}^{-1}(m, n) &= \mathbf{a}(n, m) && \text{inverse rule;} \\ \mathbf{a}(m, m) &= \mathbf{I} && \text{identity.} \end{aligned} \quad (10.110)$$

For a stationary system, the state transition matrix would be a constant matrix \mathbf{a} that is stationary or independent of time or space.

The output side of the dynamic system (see Figure 10.18) is characterized by a measurement or observation matrix $\mathbf{h}(n)$ that transforms the state vector $\mathbf{f}(n)$. The result is corrupted by a source of measurement or observation noise $\boldsymbol{\eta}_o$, which is assumed to be a zero-mean white-noise process that is statistically independent of the processes related to the state vector \mathbf{f} and the driving noise $\boldsymbol{\eta}_d$. It is assumed that the observation vector $\mathbf{g}(n)$ is of size $N \times 1$ (different from the size of the state vector \mathbf{f} which is $M \times 1$); then, the observation matrix $\mathbf{h}(n)$ is required to be of size $N \times M$, and the observation noise $\boldsymbol{\eta}_o(n)$ is required to be of size $N \times 1$. We then have the measurement or observation equation

$$\mathbf{g}(n) = \mathbf{h}(n)\mathbf{f}(n) + \boldsymbol{\eta}_o(n). \quad (10.111)$$

Similar to the characterization of the driving noise in Equation 10.109, the observation noise $\boldsymbol{\eta}_o(n)$ is characterized by its $N \times N$ correlation matrix (ACF) ϕ_{η_o} . Due to the assumption of statistical independence of $\boldsymbol{\eta}_d$ and $\boldsymbol{\eta}_o$, we have

$$E[\boldsymbol{\eta}_d(n)\boldsymbol{\eta}_o^T(k)] = \mathbf{0} \quad \forall n, k. \quad (10.112)$$

With the situation formulated as above, the Kalman filtering problem may be stated as follows: given a series of the observations $\mathcal{G}_n = \{\mathbf{g}(1), \mathbf{g}(2), \dots\}$,




FIGURE 10.19

Regions of support (ROS) for image filtering. (a) Illustration of the past, present, and future in filtering of an image, pixel by pixel, in raster-scan order. (b) Nonsymmetric half plane (NSHP) of order or size $P_1 \times P_2 \times P_3$. (c) Quarter plane (QP) of order or size $P_1 \times P_2$.

$\dots, \mathbf{g}(n)\}$, for each $n \geq 1$, find the MMSE estimate of the state vector $\mathbf{f}(l)$. The application is referred to as filtering if $l = n$; prediction if $l > n$; and smoothing or interpolation if $1 \leq l < n$. Given our application of interest in the area of image restoration, we would be concerned with filtering. (*Note:* The derivation of the Kalman filter presented in the following paragraphs closely follows those of Haykin [833] and Sage and Melsa [889].)

The innovation process: An established approach to obtain the solution to the Kalman filtering problem is via a recursive estimation procedure using a one-step prediction process, and the resultant difference referred to as the innovation process [833, 889]. Suppose that, based upon the set of observations $\mathcal{G}_{n-1} = \{\mathbf{g}(1), \mathbf{g}(2), \dots, \mathbf{g}(n-1)\}$, the MMSE estimate of $\mathbf{f}(n-1)$ has been obtained; let this estimate be denoted as $\tilde{\mathbf{f}}(n-1|\mathcal{G}_{n-1})$. Given a new observation $\mathbf{g}(n)$, we could update the previous estimate and obtain a new state vector $\tilde{\mathbf{f}}(n|\mathcal{G}_n)$. Because the state vector $\mathbf{f}(n)$ and the observation $\mathbf{g}(n)$ are related via the observation system, we may transfer the estimation procedure to the observation variable, and let $\tilde{\mathbf{g}}(n|\mathcal{G}_{n-1})$ denote the MMSE estimate of $\mathbf{g}(n)$ given \mathcal{G}_{n-1} . Then, the innovation process is defined as

$$\zeta(n) = \mathbf{g}(n) - \tilde{\mathbf{g}}(n|\mathcal{G}_{n-1}), \quad n = 1, 2, \dots \quad (10.113)$$

The innovation process $\zeta(n)$ represents the new information contained in $\mathbf{g}(n)$ that cannot be estimated from \mathcal{G}_{n-1} . Using the observation equation 10.111, we get

$$\begin{aligned} \tilde{\mathbf{g}}(n|\mathcal{G}_{n-1}) &= \mathbf{h}(n)\tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}) + \tilde{\boldsymbol{\eta}}_o(n|\mathcal{G}_{n-1}) \\ &= \mathbf{h}(n)\tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}), \end{aligned} \quad (10.114)$$

noting that $\tilde{\boldsymbol{\eta}}_o(n|\mathcal{G}_{n-1}) = \mathbf{0}$ because the observation noise is orthogonal to the past observations. Combining Equations 10.113 and 10.114, we have

$$\zeta(n) = \mathbf{g}(n) - \mathbf{h}(n)\tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}). \quad (10.115)$$

Using Equation 10.111, Equation 10.115 becomes

$$\zeta(n) = \mathbf{h}(n)\boldsymbol{\epsilon}_p(n, n-1) + \boldsymbol{\eta}_o(n), \quad (10.116)$$

where $\boldsymbol{\epsilon}_p(n, n-1)$ is the predicted state error vector at n using the information available up to $(n-1)$, given by

$$\boldsymbol{\epsilon}_p(n, n-1) = \mathbf{f}(n) - \tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}). \quad (10.117)$$

The innovation process has the following properties [833]:

- $\zeta(n)$ is orthogonal to the past observations $\mathbf{g}(1), \mathbf{g}(2), \dots, \mathbf{g}(n-1)$, and hence

$$E[\zeta(n)\mathbf{g}^T(m)] = \mathbf{0}, \quad 1 \leq m \leq n-1. \quad (10.118)$$

- The innovation process is a series of vectors composed of random variables that are mutually orthogonal, and hence

$$E [\zeta(n) \zeta^T(m)] = \mathbf{0}, \quad 1 \leq m \leq n-1. \quad (10.119)$$

- A one-to-one correspondence exists between the observations $\{\mathbf{g}(1), \mathbf{g}(2), \dots, \mathbf{g}(n)\}$ and the vectors of the innovation process $\{\zeta(1), \zeta(2), \dots, \zeta(n)\}$; that is, one of the two series may be derived from the other without any loss of information via linear operations.

The ACF matrix of the innovation process $\zeta(n)$ is given by

$$\begin{aligned} \phi_\zeta(n) &= E [\zeta(n) \zeta^T(n)] \\ &= \mathbf{h}(n) \phi_{\epsilon_p}(n, n-1) \mathbf{h}^T(n) + \phi_{\eta_o}(n), \end{aligned} \quad (10.120)$$

where

$$\phi_{\epsilon_p}(n, n-1) = E [\epsilon_p(n, n-1) \epsilon_p^T(n, n-1)] \quad (10.121)$$

is the ACF matrix of the predicted state error, and the property that $\epsilon_p(n, n-1)$ and $\eta_o(n)$ are mutually orthogonal has been used. The ACF matrix of the predicted state error provides a statistical representation of the error in the predicted state vector $\tilde{\mathbf{f}}(n|\mathcal{G}_{n-1})$.

Estimation of the state vector using the innovation process: The aim of the estimation process is to derive the MMSE estimate of the state vector $\mathbf{f}(l)$. Given that \mathbf{g} is related to \mathbf{f} via a linear transform, and that ζ is linearly related to \mathbf{g} , we may formulate the state vector as a linear transform of the innovation process as

$$\tilde{\mathbf{f}}(l|\mathcal{G}_n) = \sum_{k=1}^n \mathbf{L}_l(k) \zeta(k), \quad (10.122)$$

where $\mathbf{L}_l(k)$, $l = 1, 2, \dots, n$, is a series of transformation matrices. Now, the predicted state error vector is orthogonal to the innovation process:

$$\begin{aligned} E [\epsilon_p(l, n) \zeta^T(m)] &= E [\{\mathbf{f}(l) - \tilde{\mathbf{f}}(l|\mathcal{G}_n)\} \zeta^T(m)] \\ &= \mathbf{0}, \quad m = 1, 2, \dots, n. \end{aligned} \quad (10.123)$$

Using Equations 10.122, 10.123, and 10.119, we get

$$\begin{aligned} E [\mathbf{f}(l) \zeta^T(m)] &= \mathbf{L}_l(m) E [\zeta(m) \zeta^T(m)] \\ &= \mathbf{L}_l(m) \phi_\zeta(m). \end{aligned} \quad (10.124)$$

Consequently, we get

$$\mathbf{L}_l(m) = E [\mathbf{f}(l) \zeta^T(m)] \phi_\zeta^{-1}(m). \quad (10.125)$$

Using the expression above for $\mathbf{L}_l(m)$ in Equation 10.122, we have

$$\tilde{\mathbf{f}}(l|\mathcal{G}_n) = \sum_{k=1}^n E [\mathbf{f}(l) \boldsymbol{\zeta}^T(k)] \boldsymbol{\phi}_{\zeta}^{-1}(k) \boldsymbol{\zeta}(k), \quad (10.126)$$

from which it follows that

$$\begin{aligned} \tilde{\mathbf{f}}(l|\mathcal{G}_n) &= \sum_{k=1}^{n-1} E [\mathbf{f}(l) \boldsymbol{\zeta}^T(k)] \boldsymbol{\phi}_{\zeta}^{-1}(k) \boldsymbol{\zeta}(k) \\ &+ E [\mathbf{f}(l) \boldsymbol{\zeta}^T(n)] \boldsymbol{\phi}_{\zeta}^{-1}(n) \boldsymbol{\zeta}(n). \end{aligned} \quad (10.127)$$

For $l = n + 1$, we obtain

$$\begin{aligned} \tilde{\mathbf{f}}(n+1|\mathcal{G}_n) &= \sum_{k=1}^{n-1} E [\mathbf{f}(n+1) \boldsymbol{\zeta}^T(k)] \boldsymbol{\phi}_{\zeta}^{-1}(k) \boldsymbol{\zeta}(k) \\ &+ E [\mathbf{f}(n+1) \boldsymbol{\zeta}^T(n)] \boldsymbol{\phi}_{\zeta}^{-1}(n) \boldsymbol{\zeta}(n). \end{aligned} \quad (10.128)$$

Using the process equation 10.108, we have, for $0 \leq k \leq n$,

$$\begin{aligned} E [\mathbf{f}(n+1) \boldsymbol{\zeta}^T(n)] &= E [\{\mathbf{a}(n+1, n) \mathbf{f}(n) + \boldsymbol{\eta}_d(n)\} \boldsymbol{\zeta}^T(n)] \\ &= \mathbf{a}(n+1, n) E [\mathbf{f}(n) \boldsymbol{\zeta}^T(n)]. \end{aligned} \quad (10.129)$$

In arriving at the expression above, the property that $\boldsymbol{\eta}_d(n)$ and $\boldsymbol{\zeta}(k)$ are mutually orthogonal for $0 \leq k \leq n$ has been used. Now, using Equation 10.129 and Equation 10.126 with $l = n$, we get

$$\begin{aligned} &\sum_{k=1}^{n-1} E [\mathbf{f}(n+1) \boldsymbol{\zeta}^T(k)] \boldsymbol{\phi}_{\zeta}^{-1}(k) \boldsymbol{\zeta}(k) \\ &= \mathbf{a}(n+1, n) \sum_{k=1}^{n-1} E [\mathbf{f}(n) \boldsymbol{\zeta}^T(k)] \boldsymbol{\phi}_{\zeta}^{-1}(k) \boldsymbol{\zeta}(k) \\ &= \mathbf{a}(n+1, n) \tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}). \end{aligned} \quad (10.130)$$

Interpretation of the expression in Equation 10.128 is made easier by the following formulations.

The Kalman gain: Let

$$\mathbf{K}(n) = E [\mathbf{f}(n+1) \boldsymbol{\zeta}^T(n)] \boldsymbol{\phi}_{\zeta}^{-1}(n); \quad (10.131)$$

this is a matrix of size $M \times N$, whose significance will be apparent after a few steps. The expectation in the equation above represents the cross-correlation matrix between the state vector $\mathbf{f}(n+1)$ and the innovation process $\boldsymbol{\zeta}(n)$. Using Equations 10.131 and 10.130, Equation 10.128 may be simplified to

$$\tilde{\mathbf{f}}(n+1|\mathcal{G}_n) = \mathbf{a}(n+1, n) \tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}) + \mathbf{K}(n) \boldsymbol{\zeta}(n). \quad (10.132)$$

This is an important result, indicating that we may obtain the MMSE estimate of the state vector $\tilde{\mathbf{f}}(n+1|\mathcal{G}_n)$ by applying the state transition matrix $\mathbf{a}(n+1, n)$ to the previous estimate of the state vector $\tilde{\mathbf{f}}(n|\mathcal{G}_{n-1})$ and adding a correction term. The correction term $\mathbf{K}(n)\zeta(n)$ includes the innovation process $\zeta(n)$ multiplied with the matrix $\mathbf{K}(n)$; for this reason, and in recognition of the original developer of the underlying procedures, the matrix $\mathbf{K}(n)$ is referred to as the Kalman gain.

In order to facilitate practical implementation of the steps required to compute the Kalman gain matrix, we need to examine a few related entities, as follows. Using Equation 10.117, we can write

$$\begin{aligned} & E [\epsilon_p(n, n-1) \epsilon_p^T(n, n-1)] \\ &= E [\mathbf{f}(n) \epsilon_p^T(n, n-1)] - E [\tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}) \epsilon_p^T(n, n-1)] \\ &= E [\mathbf{f}(n) \epsilon_p^T(n, n-1)], \end{aligned} \quad (10.133)$$

where the last step follows from the property that the estimated state vector $\tilde{\mathbf{f}}(n|\mathcal{G}_{n-1})$ and the predicted state error vector $\epsilon_p(n, n-1)$ are mutually orthogonal. Using Equations 10.129, 10.116, and 10.133, we have

$$\begin{aligned} & E [\mathbf{f}(n+1) \zeta^T(n)] \\ &= \mathbf{a}(n+1, n) E [\mathbf{f}(n) \zeta^T(n)] \\ &= \mathbf{a}(n+1, n) E [\mathbf{f}(n) \{\mathbf{h}(n) \epsilon_p(n, n-1) + \boldsymbol{\eta}_o(n)\}^T] \\ &= \mathbf{a}(n+1, n) E [\mathbf{f}(n) \epsilon_p^T(n, n-1)] \mathbf{h}^T(n) \\ &= \mathbf{a}(n+1, n) E [\epsilon_p(n, n-1) \epsilon_p^T(n, n-1)] \mathbf{h}^T(n) \\ &= \mathbf{a}(n+1, n) \phi_{\epsilon_p}(n, n-1) \mathbf{h}^T(n). \end{aligned} \quad (10.134)$$

In arriving at the result above, Equation 10.121 has been used; use has also been made of the property that \mathbf{f} and $\boldsymbol{\eta}_o$ are independent processes. Using Equation 10.134 in Equation 10.131, we get the Kalman gain matrix as

$$\mathbf{K}(n) = \mathbf{a}(n+1, n) \phi_{\epsilon_p}(n, n-1) \mathbf{h}^T(n) \phi_{\zeta}^{-1}(n), \quad (10.135)$$

which upon the use of the expression in Equation 10.120 for $\phi_{\zeta}(n)$ gives

$$\mathbf{K}(n) = \mathbf{a}(n+1, n) \phi_{\epsilon_p}(n, n-1) \mathbf{h}^T(n) [\mathbf{h}(n) \phi_{\epsilon_p}(n, n-1) \mathbf{h}^T(n) + \phi_{\boldsymbol{\eta}_o}(n)]^{-1}. \quad (10.136)$$

This expression for the Kalman gain matrix may be used with Equation 10.132 to update the estimate of the state vector.

Practical computation of the Kalman gain matrix may be facilitated further by the following derivations [833]. Extending Equation 10.117 one step further, we have

$$\epsilon_p(n+1, n) = \mathbf{f}(n+1) - \tilde{\mathbf{f}}(n+1|\mathcal{G}_n). \quad (10.137)$$

Putting together Equations 10.108, 10.115, 10.132, and 10.137, we have

$$\begin{aligned} \epsilon_p(n+1, n) &= \mathbf{a}(n+1, n) \left[\mathbf{f}(n) - \tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}) \right] \\ &\quad - \mathbf{K}(n) \left[\mathbf{g}(n) - \mathbf{h}(n)\tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}) \right] + \boldsymbol{\eta}_d(n). \end{aligned} \quad (10.138)$$

Using the measurement equation 10.111 and Equation 10.117, the equation above may be modified to

$$\begin{aligned} \epsilon_p(n+1, n) &= \mathbf{a}(n+1, n) \epsilon_p(n, n-1) \\ &\quad - \mathbf{K}(n) \left[\mathbf{h}(n) \mathbf{f}(n) + \boldsymbol{\eta}_o(n) - \mathbf{h}(n) \tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}) \right] + \boldsymbol{\eta}_d(n) \\ &= \mathbf{a}(n+1, n) \epsilon_p(n, n-1) \\ &\quad - \mathbf{K}(n) \mathbf{h}(n) \left[\mathbf{f}(n) - \tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}) \right] + \boldsymbol{\eta}_d(n) - \mathbf{K}(n) \boldsymbol{\eta}_o(n) \\ &= [\mathbf{a}(n+1, n) - \mathbf{K}(n) \mathbf{h}(n)] \epsilon_p(n, n-1) + \boldsymbol{\eta}_d(n) - \mathbf{K}(n) \boldsymbol{\eta}_o(n). \end{aligned} \quad (10.139)$$

Putting together Equations 10.121 and 10.139, and noting the property that the processes ϵ_p , $\boldsymbol{\eta}_d$, and $\boldsymbol{\eta}_o$ are mutually uncorrelated, we have the ACF matrix of the predicted state error $\epsilon_p(n+1, n)$ as

$$\begin{aligned} \phi_{\epsilon_p}(n+1, n) &= E \left[\epsilon_p(n+1, n) \epsilon_p^T(n+1, n) \right] \\ &= [\mathbf{a}(n+1, n) - \mathbf{K}(n) \mathbf{h}(n)] \phi_{\epsilon_p}(n, n-1) \\ &\quad \times [\mathbf{a}(n+1, n) - \mathbf{K}(n) \mathbf{h}(n)]^T + \phi_{\boldsymbol{\eta}_d}(n) + \mathbf{K}(n) \phi_{\boldsymbol{\eta}_o}(n) \mathbf{K}^T(n) \\ &= \mathbf{a}(n+1, n) \phi_{\epsilon_p}(n, n-1) \mathbf{a}^T(n+1, n) \\ &\quad - \mathbf{K}(n) \mathbf{h}(n) \phi_{\epsilon_p}(n, n-1) \mathbf{a}^T(n+1, n) \\ &\quad - \mathbf{a}(n+1, n) \phi_{\epsilon_p}(n, n-1) \mathbf{h}^T(n) \mathbf{K}^T(n) \\ &\quad + \mathbf{K}(n) \mathbf{h}(n) \phi_{\epsilon_p}(n, n-1) \mathbf{h}^T(n) \mathbf{K}^T(n) \\ &\quad + \phi_{\boldsymbol{\eta}_d}(n) + \mathbf{K}(n) \phi_{\boldsymbol{\eta}_o}(n) \mathbf{K}^T(n) \\ &= \mathbf{a}(n+1, n) \phi_{\epsilon_p}(n, n-1) \mathbf{a}^T(n+1, n) \\ &\quad - \mathbf{K}(n) \mathbf{h}(n) \phi_{\epsilon_p}(n, n-1) \mathbf{a}^T(n+1, n) \\ &\quad - \mathbf{a}(n+1, n) \phi_{\epsilon_p}(n, n-1) \mathbf{h}^T(n) \mathbf{K}^T(n) \\ &\quad + \mathbf{K}(n) [\mathbf{h}(n) \phi_{\epsilon_p}(n, n-1) \mathbf{h}^T(n) + \phi_{\boldsymbol{\eta}_o}(n)] \mathbf{K}^T(n) + \phi_{\boldsymbol{\eta}_d}(n) \\ &= \mathbf{a}(n+1, n) \phi_{\epsilon_p}(n, n-1) \mathbf{a}^T(n+1, n) \end{aligned}$$

$$\begin{aligned}
& - \mathbf{K}(n) \mathbf{h}(n) \boldsymbol{\phi}_{\epsilon_p}(n, n-1) \mathbf{a}^T(n+1, n) \\
& - \mathbf{a}(n+1, n) \boldsymbol{\phi}_{\epsilon_p}(n, n-1) \mathbf{h}^T(n) \mathbf{K}^T(n) \\
& + \mathbf{K}(n) \boldsymbol{\phi}_{\zeta}(n) \mathbf{K}^T(n) + \boldsymbol{\phi}_{\eta_d}(n),
\end{aligned} \tag{10.140}$$

which results in

$$\boldsymbol{\phi}_{\epsilon_p}(n+1, n) = \mathbf{a}(n+1, n) \boldsymbol{\phi}_{\epsilon_p}(n) \mathbf{a}^T(n+1, n) + \boldsymbol{\phi}_{\eta_d}(n). \tag{10.141}$$

In arriving at the result above, Equations 10.120 and 10.135 have been used. A new matrix $\boldsymbol{\phi}_{\epsilon_p}(n)$, of size $M \times M$, has been introduced, defined as

$$\boldsymbol{\phi}_{\epsilon_p}(n) = \boldsymbol{\phi}_{\epsilon_p}(n, n-1) - \mathbf{a}(n, n+1) \mathbf{K}(n) \mathbf{h}(n) \boldsymbol{\phi}_{\epsilon_p}(n, n-1); \tag{10.142}$$

use has been made of the property $\mathbf{a}^{-1}(n+1, n) = \mathbf{a}(n, n+1)$, which follows from the inverse rule in Equation 10.110. Equation 10.141 is known as the Riccati equation, and assists in the recursive computation of the ACF matrix of the predicted state error.

The procedures developed to this point are referred to as Kalman's one-step or one-stage prediction algorithm [833, 889]. The algorithm is represented by, in order, Equations 10.135, 10.120, 10.115, 10.132, 10.142, and 10.141.

Application to filtering: In filtering, the aim is compute the estimate $\tilde{\mathbf{f}}(n|\mathcal{G}_n)$. The one-step prediction algorithm developed in the preceding paragraphs may be extended to the filtering application as follows.

Because the processes \mathbf{f} and $\boldsymbol{\eta}_d$ are mutually independent, it follows from Equation 10.108 that the MMSE estimate of $\mathbf{f}(n+1)$ given \mathcal{G}_n is

$$\begin{aligned}
\tilde{\mathbf{f}}(n+1|\mathcal{G}_n) &= \mathbf{a}(n+1, n) \tilde{\mathbf{f}}(n|\mathcal{G}_n) + \boldsymbol{\eta}_d(n|\mathcal{G}_n) \\
&= \mathbf{a}(n+1, n) \tilde{\mathbf{f}}(n|\mathcal{G}_n).
\end{aligned} \tag{10.143}$$

Premultiplying both sides by $\mathbf{a}(n, n+1)$, and using the inverse rule in Equation 10.110, we get

$$\begin{aligned}
\mathbf{a}(n, n+1) \tilde{\mathbf{f}}(n+1|\mathcal{G}_n) &= \mathbf{a}(n, n+1) \mathbf{a}(n+1, n) \tilde{\mathbf{f}}(n|\mathcal{G}_n) \\
&= \tilde{\mathbf{f}}(n|\mathcal{G}_n)
\end{aligned} \tag{10.144}$$

or

$$\tilde{\mathbf{f}}(n|\mathcal{G}_n) = \mathbf{a}(n, n+1) \tilde{\mathbf{f}}(n+1|\mathcal{G}_n). \tag{10.145}$$

Thus, given the result of the one-step prediction algorithm $\mathbf{f}(n+1|\mathcal{G}_n)$, we can derive the filtered estimate $\tilde{\mathbf{f}}(n|\mathcal{G}_n)$.

Let us now consider the filtered estimation error, defined as

$$\boldsymbol{\epsilon}_e(n) = \mathbf{g}(n) - \mathbf{h}(n) \tilde{\mathbf{f}}(n|\mathcal{G}_n). \tag{10.146}$$

Using Equations 10.115, 10.132, and 10.145, we can modify Equation 10.146 as follows:

$$\boldsymbol{\epsilon}_e(n) = \mathbf{g}(n) - \mathbf{h}(n) \left[\mathbf{a}(n, n+1) \tilde{\mathbf{f}}(n+1|\mathcal{G}_n) \right]$$

$$\begin{aligned}
&= \mathbf{g}(n) - \mathbf{h}(n) \left[\mathbf{a}(n, n+1) \{ \mathbf{a}(n+1, n) \tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}) + \mathbf{K}(n) \boldsymbol{\zeta}(n) \} \right] \\
&= \mathbf{g}(n) - \mathbf{h}(n) \tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}) - \mathbf{h}(n) \mathbf{a}(n, n+1) \mathbf{K}(n) \boldsymbol{\zeta}(n) \\
&= \boldsymbol{\zeta}(n) - \mathbf{h}(n) \mathbf{a}(n, n+1) \mathbf{K}(n) \boldsymbol{\zeta}(n) \\
&= [\mathbf{I} - \mathbf{h}(n) \mathbf{a}(n, n+1) \mathbf{K}(n)] \boldsymbol{\zeta}(n).
\end{aligned} \tag{10.147}$$

This expression indicates that the filtered estimation error $\boldsymbol{\epsilon}_e(n)$ is related to the innovation process $\boldsymbol{\zeta}(n)$ through a conversion factor that is given by the matrix within the square brackets. Using Equations 10.120 and 10.135, we may simplify the expression above as follows:

$$\begin{aligned}
\boldsymbol{\epsilon}_e(n) &= \left[\mathbf{I} - \mathbf{h}(n) \mathbf{a}(n, n+1) \mathbf{a}(n+1, n) \boldsymbol{\phi}_{\epsilon_p}(n, n-1) \mathbf{h}^T(n) \boldsymbol{\phi}_{\zeta}^{-1}(n) \right] \boldsymbol{\zeta}(n) \\
&= \left[\mathbf{I} - \mathbf{h}(n) \boldsymbol{\phi}_{\epsilon_p}(n, n-1) \mathbf{h}^T(n) \boldsymbol{\phi}_{\zeta}^{-1}(n) \right] \boldsymbol{\zeta}(n) \\
&= \left[\boldsymbol{\phi}_{\zeta}(n) - \mathbf{h}(n) \boldsymbol{\phi}_{\epsilon_p}(n, n-1) \mathbf{h}^T(n) \right] \boldsymbol{\phi}_{\zeta}^{-1}(n) \boldsymbol{\zeta}(n) \\
&= \boldsymbol{\phi}_{n_o}(n) \boldsymbol{\phi}_{\zeta}^{-1}(n) \boldsymbol{\zeta}(n).
\end{aligned} \tag{10.148}$$

The difference between the true state vector $\mathbf{f}(n)$ and the filtered estimate $\tilde{\mathbf{f}}(n|\mathcal{G}_n)$, labeled as the filtered state error vector $\boldsymbol{\epsilon}_f(n)$, is given by

$$\boldsymbol{\epsilon}_f(n) = \mathbf{f}(n) - \tilde{\mathbf{f}}(n|\mathcal{G}_n). \tag{10.149}$$

Using Equations 10.132 and 10.145, we may modify the equation above as

$$\begin{aligned}
\boldsymbol{\epsilon}_f(n) &= \mathbf{f}(n) - \mathbf{a}(n, n+1) \tilde{\mathbf{f}}(n+1|\mathcal{G}_n) \\
&= \mathbf{f}(n) - \mathbf{a}(n, n+1) \left[\mathbf{a}(n+1, n) \tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}) + \mathbf{K}(n) \boldsymbol{\zeta}(n) \right] \\
&= \mathbf{f}(n) - \mathbf{a}(n, n+1) \mathbf{a}(n+1, n) \tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}) - \mathbf{a}(n, n+1) \mathbf{K}(n) \boldsymbol{\zeta}(n) \\
&= \mathbf{f}(n) - \tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}) - \mathbf{a}(n, n+1) \mathbf{K}(n) \boldsymbol{\zeta}(n) \\
&= \boldsymbol{\epsilon}_p(n, n-1) - \mathbf{a}(n, n+1) \mathbf{K}(n) \boldsymbol{\zeta}(n).
\end{aligned} \tag{10.150}$$

Equation 10.137 has been used in the last step above, where $\boldsymbol{\epsilon}_p(n, n-1)$ is the predicted state error vector at n using the information provided up to $(n-1)$.

The ACF matrix of the filtered state error vector $\boldsymbol{\epsilon}_f(n)$ is obtained as follows:

$$\begin{aligned}
E [\boldsymbol{\epsilon}_f(n) \boldsymbol{\epsilon}_f^T(n)] &= E [\boldsymbol{\epsilon}_p(n, n-1) \boldsymbol{\epsilon}_p^T(n, n-1)] \\
&\quad + \mathbf{a}(n, n+1) \mathbf{K}(n) E [\boldsymbol{\zeta}(n) \boldsymbol{\zeta}^T(n)] \mathbf{K}^T(n) \mathbf{a}^T(n, n+1) \\
&\quad - E [\boldsymbol{\epsilon}_p(n, n-1) \boldsymbol{\zeta}^T(n)] \mathbf{K}^T(n) \mathbf{a}^T(n, n+1) \\
&\quad - \mathbf{a}(n, n+1) \mathbf{K}(n) E [\boldsymbol{\zeta}(n) \boldsymbol{\epsilon}_p^T(n, n-1)].
\end{aligned} \tag{10.151}$$

The expectation in the third term in the expression above may be simplified as

$$\begin{aligned}
E [\boldsymbol{\epsilon}_p(n, n-1) \boldsymbol{\zeta}^T(n)] &= E \left[\{ \mathbf{f}(n) - \tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}) \} \boldsymbol{\zeta}^T(n) \right] \\
&= E [\mathbf{f}(n) \boldsymbol{\zeta}^T(n)],
\end{aligned} \tag{10.152}$$

because $\tilde{\mathbf{f}}(n|\mathcal{G}_{n-1})$ is orthogonal to $\boldsymbol{\zeta}(n)$. Using Equation 10.129 with $k = n$ and premultiplying both sides with $\mathbf{a}^{-1}(n+1, n) = \mathbf{a}(n, n+1)$, we get

$$\begin{aligned} E [\mathbf{f}(n) \boldsymbol{\zeta}^T(n)] &= \mathbf{a}(n, n+1) E [\mathbf{f}(n+1) \boldsymbol{\zeta}^T(n)] \\ &= \mathbf{a}(n, n+1) \mathbf{K}(n) \boldsymbol{\phi}_\zeta(n). \end{aligned} \quad (10.153)$$

Equation 10.131 has been used for the second step above. Therefore, we have

$$E [\boldsymbol{\epsilon}_p(n, n-1) \boldsymbol{\zeta}^T(n)] = \mathbf{a}(n, n+1) \mathbf{K}(n) \boldsymbol{\phi}_\zeta(n). \quad (10.154)$$

Using a similar procedure, the expectation in the fourth term of Equation 10.151 may be modified as

$$E [\boldsymbol{\zeta}(n) \boldsymbol{\epsilon}_p^T(n, n-1)] = \boldsymbol{\phi}_\zeta(n) \mathbf{K}^T(n) \mathbf{a}^T(n, n+1). \quad (10.155)$$

Substituting Equations 10.153 and 10.155 in Equation 10.151, we get

$$E [\boldsymbol{\epsilon}_f(n) \boldsymbol{\epsilon}_f^T(n)] = \boldsymbol{\phi}_{\epsilon_p}(n, n-1) - \mathbf{a}(n, n+1) \mathbf{K}(n) \boldsymbol{\phi}_\zeta(n) \mathbf{K}^T(n) \mathbf{a}^T(n, n+1). \quad (10.156)$$

From Equation 10.135, we get

$$\mathbf{K}(n) \boldsymbol{\phi}_\zeta(n) = \mathbf{a}(n+1, n) \boldsymbol{\phi}_{\epsilon_p}(n, n-1) \mathbf{h}^T(n), \quad (10.157)$$

using which we can modify Equation 10.156 as

$$\begin{aligned} E [\boldsymbol{\epsilon}_f(n) \boldsymbol{\epsilon}_f^T(n)] &= \boldsymbol{\phi}_{\epsilon_p}(n, n-1) - \mathbf{a}(n, n+1) \mathbf{a}(n+1, n) \boldsymbol{\phi}_{\epsilon_p}(n, n-1) \\ &\quad \times \mathbf{h}^T(n) \mathbf{K}^T(n) \mathbf{a}^T(n, n+1) \\ &= \boldsymbol{\phi}_{\epsilon_p}(n, n-1) - \boldsymbol{\phi}_{\epsilon_p}(n, n-1) \mathbf{h}^T(n) \mathbf{K}^T(n) \mathbf{a}^T(n, n+1). \end{aligned} \quad (10.158)$$

The inverse rule of the state transition matrix (see Equation 10.110) has been used in the last step above. Because ACF matrices are symmetric, we may transpose the expression in Equation 10.158 and obtain

$$\begin{aligned} E [\boldsymbol{\epsilon}_f(n) \boldsymbol{\epsilon}_f^T(n)] &= \boldsymbol{\phi}_{\epsilon_p}(n, n-1) - \mathbf{a}(n, n+1) \mathbf{K}(n) \mathbf{h}(n) \boldsymbol{\phi}_{\epsilon_p}(n, n-1) \\ &= [\mathbf{I} - \mathbf{a}(n, n+1) \mathbf{K}(n) \mathbf{h}(n)] \boldsymbol{\phi}_{\epsilon_p}(n, n-1) \\ &= \boldsymbol{\phi}_{\epsilon_p}(n), \end{aligned} \quad (10.159)$$

where Equation 10.142 has been used for the last step. This result indicates that the matrix $\boldsymbol{\phi}_{\epsilon_p}(n)$ introduced in the Riccati equation 10.141 is the ACF matrix of the filtered state error.

Initial conditions: In practice, the initial state of the process equation 10.108 will not be known. However, it may be possible to describe it in a statistical manner, in terms of the mean and ACF of the state vector (or estimates thereof). The initial conditions given by

$$\tilde{\mathbf{f}}(1|\mathcal{G}_0) = \mathbf{0} \quad (10.160)$$

and

$$\phi_{\epsilon_p}(1, 0) = E[\mathbf{f}(1)\mathbf{f}^T(1)]. \quad (10.161)$$

result in an unbiased filtered estimate [833].

Summary of the Kalman filter: The following description of the Kalman filter, based upon one-step prediction, summarizes the main principles and procedures involved, and assists in implementing the filter [833].

Data available: The observation vectors $\mathcal{G}_n = \{\mathbf{g}(1), \mathbf{g}(2), \dots, \mathbf{g}(n)\}$.

System parameters assumed to be known:

- The state transition matrix $\mathbf{a}(n+1, n)$.
- The observation system matrix $\mathbf{h}(n)$.
- The ACF matrix of the driving noise $\phi_{\eta_d}(n)$.
- The ACF matrix of the observation noise $\phi_{\eta_o}(n)$.

Initial conditions:

- $\tilde{\mathbf{f}}(1|\mathcal{G}_0) = E[\mathbf{f}(1)] = \mathbf{0}$;
- $\phi_{\epsilon_p}(1, 0) = \mathbf{D}_0$, a diagonal matrix with values of the order of 10^{-2} .

Recursive computational steps: For $n = 1, 2, 3, \dots$, do the following:

1. Using Equation 10.136, compute the Kalman gain matrix as

$$\begin{aligned} \mathbf{K}(n) &= \mathbf{a}(n+1, n) \phi_{\epsilon_p}(n, n-1) \mathbf{h}^T(n) \\ &\quad \times [\mathbf{h}(n) \phi_{\epsilon_p}(n, n-1) \mathbf{h}^T(n) + \phi_{\eta_o}(n)]^{-1}. \end{aligned} \quad (10.162)$$

2. Obtain the innovation process vector using Equation 10.115 as

$$\boldsymbol{\zeta}(n) = \mathbf{g}(n) - \mathbf{h}(n) \tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}). \quad (10.163)$$

3. Using Equation 10.132, update the estimate of the state vector as

$$\tilde{\mathbf{f}}(n+1|\mathcal{G}_n) = \mathbf{a}(n+1, n) \tilde{\mathbf{f}}(n|\mathcal{G}_{n-1}) + \mathbf{K}(n) \boldsymbol{\zeta}(n). \quad (10.164)$$

4. Compute the ACF matrix of the filtered state error, given by Equation 10.142 as

$$\phi_{\epsilon_p}(n) = \phi_{\epsilon_p}(n, n-1) - \mathbf{a}(n, n+1) \mathbf{K}(n) \mathbf{h}(n) \phi_{\epsilon_p}(n, n-1). \quad (10.165)$$

5. Using Equation 10.141, update the ACF matrix of the predicted state error as

$$\phi_{\epsilon_p}(n+1, n) = \mathbf{a}(n+1, n) \phi_{\epsilon_p}(n) \mathbf{a}^T(n+1, n) + \phi_{\eta_d}(n). \quad (10.166)$$


The Kalman filter formulation is a general formulation of broad scope, relevance, and application. Haykin [833] provides a discussion on the Kalman filter as the unifying basis for recursive least-squares (RLS) filters. Sage and Melsa [889] present the derivation of a stationary Kalman filter and relate it to the Wiener filter.

Extension to of the Kalman filter to image restoration: Extending the Kalman filter to 2D image filtering poses the following challenges:

- defining the state vector;
- determining the size (spatial extent) of the state vector;
- deriving the dynamic (space-variant) state transition matrix (process, phenomenon, or model);
- obtaining the dynamic (space-variant) observation matrix (system);
- estimating the driving and observation noise correlation matrices; and
- dealing with the matrices of large size due to the large number of elements in the state vector.

Woods and Radewan [891, 892] and Woods and Ingle [893] proposed a scalar processor that they called as the reduced-update Kalman filter (RUKF), discussing in detail the options for the ROS of the filter [in particular the NSHP ROS, see Figure 10.19 (b)] as well as the problems associated with the boundary conditions (see also Tekalp et al. [894, 895, 896]). Boulfefel et al. [750] modified the RUKF algorithm of Woods and Ingle [893] for application to the restoration of SPECT images with the following main characteristics:

- The state transition matrix was derived using a 2D AR model (see Section 11.8).
- The observation matrix was composed by selecting one of 16 PSFs depending upon the distance of the pixel being processed from the center of the image (see Figure 10.20 and Section 10.5 for details).
- Filtering operations were performed in a QP ROS, as illustrated in Figure 10.19 (c). The image was processed pixel-by-pixel in raster-scan order, by moving the QP ROS window from one pixel to the next.
- The size of the ROS was determined as the larger of the AR model order (related to the state transition matrix) and the width of the PSF.
- With the assumption that the observation noise follows the Poisson PDF, the variance of the observation noise was estimated as the total photon count in a local window.

**FIGURE 10.20**

In order to apply the nonstationary Kalman filter to SPECT images, the image plane is divided into 16 zones based upon the distance from the center of the axis of rotation of the camera [750]. A different PSF (MTF) is used to process the pixels in each zone. The width of the PSF increases with distance from the center.

The steps and equations of the scalar RUKF algorithm are as follows:

Update index n to $n + 1$; at the end of a row, reset $n = 1$, and update m to $m + 1$.

1. Project the previous estimate of the state vector one step forward by using the dynamic system model as

$$\tilde{f}_b^{(m,n)}(m, n) = \sum_{\alpha} \sum_{\beta} a^{(m,n)}(\alpha, \beta) \tilde{f}_a^{(m,n-1)}(m - \alpha, n - \beta). \quad (10.167)$$

2. Project the error ACF matrix one step forward as

$$\phi_b^{(m,n)}(m, n; \alpha, \beta) = \sum_r \sum_s a^{(m,n)}(r, s) \phi_a^{(m,n-1)}(m - r, n - s; \alpha, \beta) \quad (10.168)$$

and

$$\begin{aligned} \phi_b^{(m,n)}(m, n; m, n) &= \sum_{\alpha} \sum_{\beta} a^{(m,n)}(\alpha, \beta) \phi_b^{(m,n)}(m, n; m - \alpha, n - \beta) \\ &\quad + \sigma_{\eta_a}^2(m, n). \end{aligned} \quad (10.169)$$

3. Compute the updated Kalman gain matrix as

$$K^{(m,n)}(p, q) = \left[\sum_{\alpha} \sum_{\beta} h^{(m,n)}(\alpha, \beta) \phi_b^{(m,n)}(m - \alpha, n - \beta; m - p, n - q) \right] \div \left[\sum_{\alpha} \sum_{\beta} \sum_p \sum_q h^{(m,n)}(\alpha, \beta) h^{(m,n)}(p, q) \phi_b^{(m,n)}(m - \alpha, n - \beta; m - p, n - q) + \sigma_{\eta_o}^{2(m,n)} \right]. \quad (10.170)$$

4. Update the state vector as

$$\begin{aligned} \tilde{f}_a^{(m,n)}(p, q) &= \tilde{f}_b^{(m,n)}(p, q) + K^{(m,n)}(m - p, n - q) \\ &\quad \times \left[g(m, n) - \sum_{\alpha} \sum_{\beta} h^{(m,n)}(\alpha, \beta) \tilde{f}_b^{(m,n)}(m - \alpha, n - \beta) \right]. \end{aligned} \quad (10.171)$$

5. Update the error ACF matrix as

$$\begin{aligned} \phi_a^{(m,n)}(p, q; \alpha, \beta) &= \phi_b^{(m,n)}(p, q; \alpha, \beta) - K^{(m,n)}(m - p, n - q) \\ &\quad \times \sum_r \sum_s h^{(m,n)}(r, s) \phi_b^{(m,n)}(m - r, n - s; \alpha, \beta). \end{aligned} \quad (10.172)$$

In the notation used above, the superscript (m, n) indicates the filtering step (position of the QP filtering window); the indices within the argument of a variable indicate the spatial coordinates of the pixels of the variable being used or processed; the subscript b indicates the corresponding variable before updating; the subscript a indicates the corresponding variable after updating; and all summations are over the chosen ROS for the filtering operations. The subscript ϵ_p has been dropped from the error ACF ϕ .

With reference to the basic Kalman algorithm documented on page 915, the RUKF procedure described above has the following differences:

- The sequence of operations in the RUKF algorithm above follows the Kalman filter algorithm described by Sage and Melsa [889] (p268), and is different from that given by Haykin [833] as well as the algorithm on page 915.
- Equation 10.167 computes only the first part of the right-hand side of Equation 10.164.

- Equations 10.168 and 10.169 together perform the operations represented by Equation 10.166.
- Equation 10.170 is equivalent to Equation 10.162 except for the presence of the state transition matrix $\mathbf{a}(n+1, n)$ in the latter.
- Equation 10.171 is similar to Equation 10.164, with the observation that the innovation process $\zeta(n)$ is given by the expression within the brackets on the right-hand side of the latter.
- Equation 10.172 is equivalent to Equation 10.165 except for the presence of the state transition matrix $\mathbf{a}(n, n+1)$ in the latter. Observe that the term $\mathbf{a}(n+1, n)$ in Equation 10.162 is cancelled by the term $\mathbf{a}(n, n+1)$ in Equation 10.165 due to the inverse rule given in Equation 10.110.

Illustrations of application of the 2D Kalman filter (the RUKF algorithm) to the restoration of SPECT images are provided in Section 10.5.

10.5 Application: Restoration of Nuclear Medicine Images

Nuclear medicine images, including planar, SPECT, and PET images, are useful in functional imaging of several organs, such as the heart, brain, thyroid, and liver; see Section 1.7 for an introduction to the principles of nuclear medicine imaging. However, nuclear medicine images are severely affected by several factors that degrade their quality and resolution. Some of the important causes of image degradation in nuclear medicine imaging are briefly described below, along with suitable methods for image restoration [46, 897, 898].

- **Poor quality control:** SPECT images could be blurred by misalignment of the axis of rotation of the system in the image reconstruction process. Data in the projection images could become inaccurate due to defects and nonuniformities in the detector. Patient motion during image data acquisition leads to blurring.
- **Poor statistics:** The number of photons acquired in a nuclear medicine image will be limited due to the constrained dose of the radiopharmaceutical administered, the time that the patient can remain immobile for the imaging procedure, the time over which the distribution and activity of the radiopharmaceutical within the patient will remain stable, the low efficiency of detection of photons imposed by the collimator, the limited photon-capture efficiency of the detection system, and the

attenuation of the photons within the body. These factors lead to low SNR in the images. SPECT images have poorer statistics due to the limited time of acquisition of each projection image. See Figure 3.68 for an illustration of increased levels of noise in planar images due to limited imaging periods.

- **Photon-counting (Poisson) noise:** The emission of gamma-ray photons is an inherently random process that follows the Poisson distribution (see Section 3.1.2). When the photon count is large, the Poisson PDF tends toward the Gaussian PDF (see Figure 3.8). The detection of gamma-ray photons is also a random process, governed by the probabilities of interaction of photons with matter. The noise in SPECT images is further affected by the filters used in the process of image reconstruction from projections: the noise is amplified and modified to result in a mottled appearance of relatively uniform regions in SPECT images.
- **Gamma-ray attenuation:** Photons are continuously lost as a gamma ray passes through an object due to scattering and photoelectric absorption [899] (p140). The effect of these phenomena is an overall attenuation of the gamma ray. The number of photons that arrive at the detector will depend upon the attenuating effect of the tissues along the path of the ray, and hence will not be a true representation of the strength at the source (the organ that collected the radiopharmaceutical in relatively higher proportion).
- **Compton scattering:** Compton scattering occurs when a gamma-ray photon collides with an outer-shell electron in the medium through which it passes: the (scattered or secondary) gamma-ray photon continues in a different direction at a lower energy, and the electron that gained the energy is released and scattered [899] (pp 140 – 142). Gamma-ray photons affected by Compton scattering cause counts in images at locations that correspond to no true emitting source, and appear as background noise. However, because the affected photon arrives at the detector with a lower energy than that at which it was emitted at the source (which is known), it may be rejected based upon this knowledge.
- **Poor spatial resolution:** The spatial resolution of a gamma camera is affected by two different factors: the intrinsic resolution of the detector and the geometric resolution of the collimator. The intrinsic resolution of the detector is related to the diameter of the PMTs, the thickness of the crystal, and the energy of the gamma rays used. The intrinsic resolution may be improved by using more PMT tubes, a thinner crystal, and gamma rays of higher energy. However, a thinner crystal would lead to lower efficiency in the detection of gamma-ray photons.

The geometric resolution of the collimator is a function of the depth (thickness) of the collimator, the diameter of the holes, and the source-

to-collimator distance. The geometric resolution of a collimator may be improved by making it thicker and reducing the size of the holes; however, these measures would reduce the number of photons that reach the crystal, and hence reduce the overall efficiency of photon detection.

Digital image processing techniques may be applied to correct for some of the degrading phenomena mentioned above; filters may also be designed to remove some of the effects [900, 901]. The following sections provide details of a few methods to improve the quality of nuclear medicine images; a schematic representation of the application of such techniques is shown in Figure 10.21.

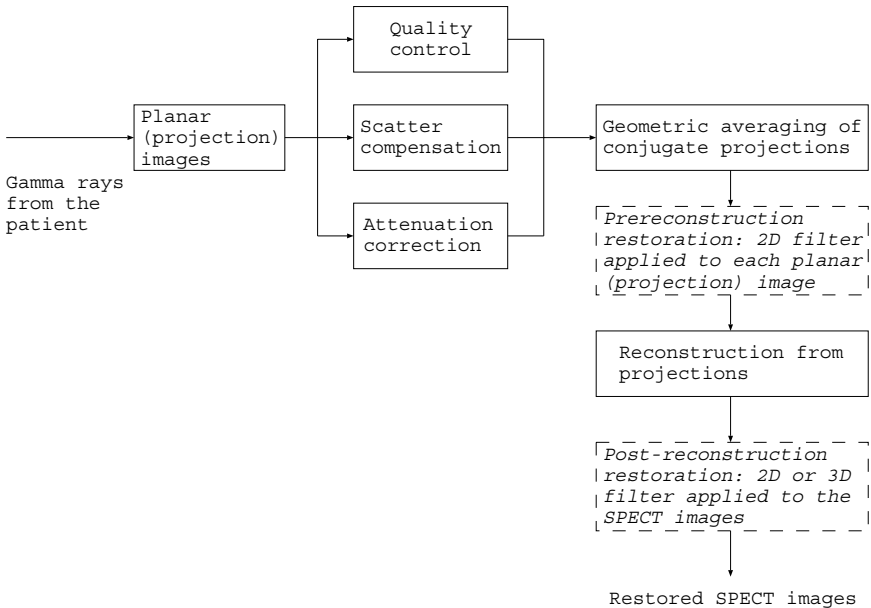


FIGURE 10.21

Schematic representation of the application of several techniques to improve the quality of nuclear medicine images. Attenuation correction may be achieved via averaging of conjugate projections, modifications to the reconstruction algorithm, or other methods applied to the planar projections or the reconstructed SPECT data. The blocks are shown separately to emphasize the various procedures for quality improvement. Only one of the two blocks shown in dashed lines — pre-reconstruction or post-reconstruction processing — would be used.

10.5.1 Quality control

The quality of images obtained using a gamma camera is affected by imperfections in the detection system of the camera. System misalignment, detector nonlinearity, and detector nonuniformity are the major contributors to image degradation in this category. System misalignment errors arise when an incorrect axis of rotation is used in the reconstruction process, or due to nonlinearities in the field of view of the camera, and result in smearing of the SPECT image [46, 902]. System misalignment may be corrected through regular calibration of the camera system [902, 903]. Detector nonlinearity arises due to the compression of events located near the centers of the PMTs and an expansion between the PMTs [904, 905]. Detector nonlinearities are usually not perceptible; however, they can have significant effects on image nonuniformities and misalignment. The finite number of PMTs in the detector is also a source of image nonuniformity, which results in variations in counts in the acquired image of a uniform object [904, 906, 907].

The common approach to correct nonuniformity is by acquiring an image of a flood-field (a uniform source) and then using it as a correction matrix for other images [46, 905, 908]. However, this method only corrects for variations in amplitude. A more sophisticated method stores correction matrices for regional differences in pulse-height spectra and for positions over the entire area of the detector. The correction matrices are then used on an event-by-event basis to compensate for regional differences by adjusting either the system gain or the pulse-height window, and to compensate for nonlinearities by accurately repositioning each event [904].

10.5.2 Scatter compensation

Compton scattering results in a broad, symmetric distribution centered about the primary wavelength or energy level of the gamma rays at the source. One approach to reduce the effect of scatter is by energy discrimination in the camera system: by rejecting photons at all energy levels outside a narrow window centered at the photo-peak of the radio-isotope, most of the Compton-scattered photons will be rejected; however, this rejection is not complete [909].

Jaszczak et al. [910] acquired projection images in two different energy windows, one placed at the photo-peak of the radio-isotope and the other over the Compton portion of the energy spectrum. A fraction of the second image was subtracted from the photo-peak image. The procedure was shown to eliminate most of the scatter; however, difficulty was encountered in determining the fraction of the scatter image to be subtracted. Egbert and May [911] modeled Compton scattering using the integral transport equation. Correction was performed in an iterative procedure using an attenuation-corrected reconstructed image as the initial estimate. The next image estimate was computed as a subtraction of the product of the Chang point-wise attenua-

tion correction operator [912], with a scattering operator determined from a given energy threshold from the previous estimate of the image. This technique suffers from the necessity of having to determine the scattering operator for each distribution of the scattering medium, photon energy, and threshold. Axelsson et al. [913] proposed a method in which the scatter distribution is modeled as the convolution of the projection image with an exponential kernel, and correction is performed by subtracting the estimate of the scatter from the acquired projection image. A similar technique was described by Floyd et al. [914, 915], in which the correction is performed by deconvolution rather than by subtraction.

See Rosenthal et al. [898], Buvat et al. [916], and Ljungberg et al. [917] for other methods for scatter compensation.

10.5.3 Attenuation correction

When gamma-ray photons pass through body tissues, several photons get attenuated and scattered. The amount of attenuation depends upon the attenuation coefficient of the tissues and the depth of the source of the photons. The attenuation in nuclear medicine imaging may be represented as

$$I_d = I_s \exp[-\mu x], \quad (10.173)$$

where I_s is the intensity of the gamma ray at the source, I_d is the intensity at the detector, μ is the attenuation coefficient of the attenuating medium or tissue (assumed to be uniform in the expression above), and x is the distance traveled by the gamma ray through the attenuating medium. Several methods for attenuation correction have been described in the literature; the methods may be divided into three categories: preprocessing correction methods, intrinsic correction methods, and post-processing correction methods [46, 912, 918, 919, 920, 921, 922, 923, 924, 925, 926].

The most common preprocessing correction procedure is to estimate the source-to-collimator distance at each point in the image, and use Equation 10.173 to compute the gamma-ray intensity at the source [921]. Other approaches include using the arithmetic mean or the geometric mean of conjugate projections to correct projections for attenuation before reconstruction [46]; see Section 10.5.5. These two methods are employed in most SPECT systems, and give acceptable results for head and liver SPECT images, where the attenuation coefficients may be assumed to be constant.

In the intrinsic correction methods, attenuation correction is incorporated into the reconstruction algorithm. Gullberg and Budinger [924] proposed a technique where the solution to the attenuated Radon integral for the case of a uniformly attenuating medium is used. This technique is limited to images with high statistics because it can amplify noise and introduce artifacts in low-count images. Censor et al. [927] proposed a correction method where both the attenuation coefficient map and the radio-isotope distribution are

estimated using discrete estimation theory; this approach provides noisier images but allows correction of nonuniform attenuation.

A commonly used post-processing correction method is the Chang iterative technique [912], in which each pixel of the SPECT image is multiplied by the reciprocal of the average attenuation along all rays from the pixel to the boundaries of the attenuating medium. The results are then iteratively refined by reprojecting the image using the assumed attenuation coefficients [925]. The difference between the true (acquired) projections and those obtained by reprojection is used to correct the image. This method does not amplify noise, and performs well in the case of inhomogeneous attenuating media.

See Rosenthal et al. [898] and King et al. [928] for other methods for attenuation correction.

10.5.4 Resolution recovery

Image restoration techniques could be applied to nuclear medicine images to perform resolution recovery (deblurring) and noise removal. Most restoration methods assume the presence of additive signal-independent white noise and a shift-invariant blurring function. However, nuclear medicine images are corrupted by Poisson noise that is correlated with or dependent upon the image data, and the blurring function is shift-variant. Furthermore, the noise present in the planar projection images is amplified in SPECT images by the reconstruction algorithm. The SNR of nuclear medicine images is low, which makes recovery or restoration difficult.

Several methods have been proposed for restoring nuclear medicine images [86, 87, 132, 749, 750, 751, 834, 835, 842, 844, 845, 846, 847, 848, 849, 851, 929, 930, 931, 932, 933, 934, 935, 936, 937, 938, 939]. SPECT images may be either restored after reconstruction (post-reconstruction restoration) [132, 834, 844, 842, 851, 929], or the projection (planar) images may be restored first and SPECT images reconstructed later (prereconstruction restoration) [87, 751, 835, 845, 851, 930]; see Figure 10.21.

Boardman [929] applied a constrained deconvolution filter to restore scintigraphic images of the brain: the filter required only the PSF as a priori information; it was assumed that there were no discontinuities in the object. Madsen and Park [930] applied a Fourier-domain filter on the projection set for the enhancement of SPECT images (prereconstruction restoration); they assumed the PSF to be a 2D isotropic Gaussian function. King et al. [834, 844, 845] applied the Wiener filter and a count-dependent Metz filter to restore both projection and SPECT images, and reported that prereconstruction restoration showed better results than post-reconstruction restoration [845]. However, a study of image contrast and percent fractional standard deviation of counts in regions of uniform activity in the test images used in their studies does not show a clear difference between the two methods. Webb et al. [842] attempted post-reconstruction restoration of liver SPECT images; they proposed a general formulation that unifies a number of filters, among which are

the inverse, maximum entropy, parametric Wiener, homomorphic, Philips, and hybrid filters. A study of image contrast in their work indicated that properly tuned maximum-entropy or homomorphic filters provide good results. Honda et al. [835] also attempted post-reconstruction restoration of myocardial SPECT images using a combination of Wiener and Butterworth filters; they assumed the SNR to be a constant and the system transfer function to be Gaussian.

Boulfefel et al. [751, 851] performed a comparison of prereconstruction and post-reconstruction restoration filters applied to myocardial images. The results obtained showed that prereconstruction restored images present a significant decrease in RMS error and an increase in contrast over post-reconstruction restored images. Examples from these studies are presented in Section 10.5.6.

An important consideration in the restoration of SPECT images is the stationarity of the PSF [46, 940, 941, 942, 943]. Hon et al. [132] and Boulfefel et al. [86, 749, 750] conducted several experiments to derive the FWHM values of the PSF of several imaging configurations, and tabulated the variation in the parameter with source-to-collimator distance and attenuating media.

Larsson [46] investigated the effects of using the arithmetic and geometric mean of LSFs taken at opposing (conjugate) angular positions in air and water to improve stationarity. It was found that the arithmetic mean of opposing LSFs resulted in an LSF of approximately the same FWHM as the LSF at the center of rotation of the camera, although there were significant amplitude differences. However, the geometric mean of opposing LSFs provided comparable FWHMs and amplitudes. Furthermore, SPECT images reconstructed by using the geometric means of opposing LSFs did not show any significant distortion due to the spatial variation in the detector response with distance from the collimator, as compared to the results with the arithmetic means [46]. However, Larsson found that when more than one source is employed, the use of the geometric mean resulted in interference artifacts between the sources. Axelsson et al. [913] investigated the shape of the LSF resulting from the geometric mean of opposing planar projections, and concluded that the LSF was approximately stationary in the central two-thirds of their phantom. Msaki et al. [944] also found that the stationarity of the 2D MTF improved in their study of the variation in the PSF of the geometric means of opposing views for source positions both orthogonal and parallel to the collimator axis. Boulfefel et al. [87] evaluated the use of the geometric mean of opposing projections in prereconstruction restoration, and found that this preprocessing step could lead to improved results; the details of their methods and results are presented in Section 10.5.5.

Coleman et al. [945] used the arithmetic and geometric means of opposing projections to calculate the 2D MTF and scatter fraction as a function of the camera angle and source location. Their results showed that the geometric mean provided an approximately stationary 2D MTF and scatter fraction, except near the edges of the phantom. However, it was observed by Coleman

et al. that the arithmetic mean provided more nonstationary results than the geometric mean, which were both less nonstationary than the MTF related to the planar projections. It was concluded that the arithmetic mean may be preferred to the geometric mean in the restoration of SPECT images because the latter is nonlinear. Although Coleman et al. recommended the use of averaging in restoration, their study did not perform any restoration experiments. King et al. [946], continuing the study of Coleman et al. [945], used the means of conjugate views (both arithmetic and geometric) for prereconstruction attenuation correction in their study on the use of the scatter degradation factor and Metz filtering for the improvement of activity quantitation; however, the Metz filter they used was predefined and did not explicitly make use of an MTF related to the averaging procedure. Glick et al. [942, 943, 947] and Coleman et al. [945] also investigated the effect of averaging (arithmetic and geometric) of opposing views on the stationarity of the PSF. Whereas these studies concluded that the 2D PSF is not stationary and is shift-variant, the 3D PSF has been shown to be more stationary and only slightly affected by the source-to-collimator distance [942].

The PSF of SPECT images has been shown to have a 3D spread [942, 943, 947]; regardless, most of the restoration methods proposed in the literature assume a 2D PSF. The use of a 2D PSF results in only a partial restoration of SPECT images, because the inter-slice blur is ignored. Boulfefel et al. [86, 749, 948] and Rangayyan et al. [935] studied the 3D nature of the PSF and proposed 3D filters to address this problem; examples from these works are presented in Section 10.5.6.

It has been shown that discretization of the filtered backprojection process can cause the MTF related to the blurring of SPECT images to be anisotropic and nonstationary, especially near the edges of the camera's field of view [942, 943, 947]. Furthermore, the Poisson noise present in nuclear medicine images is nonstationary. Shift-invariant restoration techniques will fail in the restoration of large images because they do not account for such variations in the MTF and the noise. Therefore, restoration methods for SPECT images should support the inclusion of a shift-variant MTF and nonstationary noise characterization. Boulfefel et al. [86, 750, 948] investigated the use of a shift-variant 2D Kalman filter for the restoration of SPECT images. The Kalman filter allows for the use of different MTFs and noise parameters at each pixel, and was observed to perform better than shift-invariant filters in the restoration of large objects and organs. Examples of this application are presented in Section 10.5.6.

10.5.5 Geometric averaging of conjugate projections

Geometric averaging of conjugate projections (reviewed in Section 10.5.4 in brief, and described in more detail in the following paragraphs) may be viewed as a predistortion mapping technique that consists of a transformation applied to the degraded image in such a way that the shift-variant blurring function

that caused the degradation becomes shift-invariant. The widely used coordinate transformation method [827, 828, 949] is a predistortion mapping scheme that eliminates the shift-variance of a blurring function by changing to a system of coordinates in which the blur becomes shift-invariant. Shift-invariant restoration may then be used, with the image changed back to the original coordinates by an inverse coordinate transformation. Coordinate transformation methods have been applied for shift-variant blurs such as motion blur, where the shift-variance is linear. However, this technique is not applicable in situations where the blur varies with the depth of field, or where a projection (integration) operation is involved as in planar nuclear medicine images. Geometric averaging of conjugate projections could be interpreted as a predistortion mapping scheme that reduces the shift-variance of the blur (but not eliminating it completely). Furthermore, because the procedure combines two nearly identical planar projections images acquired from opposite directions and averages them, only the blur function is affected, and the restored image does not need any processing for coordinate change.

In the work of Boulfefel et al. [86, 87], the MTF of the gamma camera was modeled as

$$H(u, v) = \exp \left\{ -2 \left[\frac{\pi \sigma \sqrt{(u^2 + v^2)}}{N f_s} \right]^2 \right\}, \quad (10.174)$$

where N is the width of the MTF in pixels, f_s is the sampling frequency, and σ is the standard deviation of the Gaussian in the model of the PSF. The FWHM of the PSF was experimentally measured using a line source (see Section 2.9 and Figure 2.21) for various source-to-collimator distances. The variation of the FWHM with source-to-collimator distance d was found to be linear, and modeled as

$$FWHM = a + b d, \quad (10.175)$$

where a and b are the parameters of the linear model that were determined from experimental measurements of FWHM. The standard deviation σ of the Gaussian model of the PSF and MTF is related to FWHM as

$$\sigma = \frac{FWHM}{\kappa} = \frac{a + b d}{\kappa}, \quad (10.176)$$

where $\kappa = 2.355$ is a constant of proportionality.

Let us assume that the camera is rotating around a point source that is located away from the center of rotation of the camera, as shown in Figure 10.22. If the radius of rotation is R , and the camera is at a distance d from the point source when at the closest position, then rotating the camera by 180° places the point source at a distance $(2R - d)$ from the camera. When the point source is at a distance d from the camera, the MTF according to Equation 10.174 is

$$H_\theta(u, v) = \exp \left\{ -2 \left[\frac{\pi (a + b d) \sqrt{(u^2 + v^2)}}{\kappa N f_s} \right]^2 \right\}, \quad (10.177)$$

where θ refers to the angle of the camera. The MTF at the conjugate position is

$$H_{\theta^*}(u, v) = \exp \left\{ -2 \left[\frac{\pi [a + b(2R - d)] \sqrt{(u^2 + v^2)}}{\kappa N f_s} \right]^2 \right\}, \quad (10.178)$$

where $\theta^* = \theta + 180^\circ$ refers to the angle of the camera. The geometric mean of the two MTFs given above is given by

$$H_g(u, v) = [H_\theta(u, v) H_{\theta^*}(u, v)]^{1/2}. \quad (10.179)$$

Therefore, we have

$$\begin{aligned} H_g^2(u, v) &= \exp \left\{ -\frac{2\pi^2(u^2 + v^2)}{\kappa^2 N^2 f_s^2} \{ (a + bd)^2 + [a + b(2R - d)]^2 \} \right\} \\ &= \exp \left\{ -\frac{2\pi^2(u^2 + v^2)}{\kappa^2 N^2 f_s^2} \{ 2a^2 + 4Rab + b^2(2d^2 - 4Rd + 4R^2) \} \right\}, \end{aligned} \quad (10.180)$$

which leads to

$$H_g(u, v) = \exp \left\{ -\frac{2\pi^2(u^2 + v^2)}{\kappa^2 N^2 f_s^2} \{ a^2 + 2Rab + b^2(d^2 - 2Rd + 2R^2) \} \right\}. \quad (10.181)$$

If the point source is located at the center of rotation of the camera, we have $d = R$, and the MTF is reduced to

$$H_0(u, v) = \exp \left\{ -\frac{2\pi^2(u^2 + v^2)}{\kappa^2 N^2 f_s^2} \{ a^2 + 2Rab + b^2 R^2 \} \right\}. \quad (10.182)$$

Letting

$$B(u, v) = -\frac{2\pi^2(u^2 + v^2)}{\kappa^2 N^2 f_s^2}, \quad (10.183)$$

we have


$$H_\theta(u, v) = \exp \{ B(u, v) (a^2 + 2abd + b^2 d^2) \}, \quad (10.184)$$

$$H_g(u, v) = \exp \{ B(u, v) [a^2 + 2Rab + b^2 (d^2 - 2Rd + 2R^2)] \}, \quad (10.185)$$

and

$$H_0(u, v) = \exp \{ B(u, v) (a^2 + 2Rab + b^2 R^2) \}. \quad (10.186)$$

Equations 10.184, 10.185, and 10.186 are, respectively, the MTF related to a point source located at a distance d from a camera rotating around a circle of radius R , the geometric mean of the MTFs related to two opposing projections of a point source located at distances d and $(2R - d)$ from the camera, and the MTF of a point source located at the center of rotation of the camera with

**FIGURE 10.22**

Gamma-camera imaging geometry illustrating conjugate projections being obtained for a point source at distances d and $2R - d$ from the camera in the two views, where R is the radius of rotation of the camera. The same principles apply to imaging with a dual-camera or multi-camera imaging system. Reproduced with permission from D. Boulfefel, R.M. Rangayyan, L.J. Hahn, and R. Kloiber, "Use of the geometric mean of opposing planar projections in prereconstruction restoration of SPECT images", *Physics in Medicine and Biology*, 37(10): 1915–1929, 1992. © IOP Publishing Ltd.

$d = R$. Equation 10.184 shows clearly that the MTF is highly dependent on the source-to-collimator distance, whereas Equation 10.185 suggests that the geometric averaging procedure makes the MTF less dependent on the source-to-collimator distance (only the last factor in the equation is a function of the source-to-collimator distance d), and that it is close to the MTF of a point source located at the center of rotation (Equation 10.186).

Comparing Equations 10.184, 10.185, and 10.186 leads to the comparison of the following three functions involving the source-to-collimator distance d and the other parameters encountered above:

$$f_{\theta}(d) = a^2 + 2abd + b^2 d^2, \quad (10.187)$$

$$f_g(d) = a^2 + 2Rab + b^2 (d^2 - 2Rd + 2R^2), \quad (10.188)$$

and

$$f_0(d) = a^2 + 2Rab + b^2 R^2. \quad (10.189)$$

The subscripts of the function $f(d)$ relate to the subscripts of $H(u, v)$ in Equations 10.184, 10.185, and 10.186. Figure 10.23 shows plots of the three functions $f_{\theta}(d)$, $f_g(d)$, and $f_0(d)$ for radius of rotation $R = 20.0 \text{ cm}$, and the FWHM model parameters $a = 0.383762$ and $b = 0.0468488$ (from experimental measurements using line sources with the Siemens Rota camera at the Foothills Hospital, Calgary). It is seen that $f_0(d) = 1.74$ (a constant), $f_{\theta}(d)$ varies from 0.15 to 5.09, and $f_g(d)$ varies between 1.74 and 2.62. The

plots and the ranges of the values of the three functions show that geometric averaging reduces the space-variance of the MTF.




FIGURE 10.23

Plots of the distance functions $f_{\theta}(d)$ in Equation 10.187 (solid line), $f_g(d)$ in Equation 10.188 (dashed line), and $f_0(d)$ in Equation 10.189 (dotted line). Reproduced with permission from D. Boulfefel, R.M. Rangayyan, L.J. Hahn, and R. Kloiber, "Use of the geometric mean of opposing planar projections in prereconstruction restoration of SPECT images", *Physics in Medicine and Biology*, 37(10): 1915–1929, 1992. © IOP Publishing Ltd.

Figure 10.24 shows profiles of the MTFs related to a point source at $d = 20\text{ cm}$ and $d = 40\text{ cm}$ with $R = 30\text{ cm}$, the averaged MTF with $d = 20\text{ cm}$ and $d = 40\text{ cm}$, and the MTF at the center of rotation of the camera ($d = R$), computed using Equations 10.184, 10.185, and 10.186, respectively. The figure shows that the averaged MTF and the MTF at the center of rotation are close to each other, as compared to the MTFs for the point source at $d = 20\text{ cm}$ and $d = 40\text{ cm}$.

Boulfefel et al. [87, 86] performed experimental measurements with a line source to verify the validity of the theoretical results described above. The line source was constructed with a thin plastic tube of internal radius of 1 mm and filled with 1 mCi of ^{99m}Tc . No scattering medium was used. Figure 10.25




FIGURE 10.24

Computed profiles of MTFs related to point sources in gamma-camera imaging: (a) averaged MTF with $d = 20\text{ cm}$ and $d = 40\text{ cm}$, (b) MTF at the center of rotation of the camera ($d = R = 30\text{ cm}$), (c) $d = 20\text{ cm}$, and (d) $d = 40\text{ cm}$. Reproduced with permission from D. Boulfefel, R.M. Rangayyan, L.J. Hahn, and R. Kloiber, "Use of the geometric mean of opposing planar projections in prereconstruction restoration of SPECT images", *Physics in Medicine and Biology*, 37(10): 1915–1929, 1992. © IOP Publishing Ltd.

shows profiles of the PSF derived from the LSF for source-to-collimator distances $d = 20, 30$, and 40 cm, obtained using the ADAC GENESYS camera with the low-energy general-purpose collimator at the Foothills Hospital, Calgary. The averaged PSF for $d = 20$ cm and 40 cm is also plotted in the figure. It is seen that the averaged PSF matches closely the PSF at the central position of $d = 30$ cm.




FIGURE 10.25

Experimentally measured profiles of PSFs in gamma-camera imaging: (a) $d = 20$ cm, (b) at the center of rotation of the camera ($d = R = 30$ cm), (c) $d = 40$ cm, and (d) averaged PSF with $d = 20$ cm and $d = 40$ cm. Reproduced with permission from D. Boulfefel, R.M. Rangayyan, L.J. Hahn, and R. Kloiber, "Use of the geometric mean of opposing planar projections in prereconstruction restoration of SPECT images", *Physics in Medicine and Biology*, 37(10): 1915–1929, 1992. © IOP Publishing Ltd.

The preceding derivations and arguments were based upon a point source being imaged. In order to extend the arguments to a combination of distributed sources, we could consider a planar source image $p(x, y)$ that is parallel to the plane of the camera. A 3D source may then be modeled as a collection of several planar sources, with the individual planes being parallel to the plane of the camera. Then, the acquired image of each plane could be

modeled as being convolved with the blur PSF for the corresponding distance to the camera. The net projection of the 3D source would be the sum of the blurred images of the constituent planes.

For a planar source $p(x, y)$ placed away from the center of the axis of rotation of the camera, let us consider two planar images acquired, at an angle θ and its conjugate θ^* . In the frequency domain, we may represent the planar images as

$$P_{\theta}(u, v) = H_{\theta}(u, v) P(u, v), \quad (10.190)$$

and

$$P_{\theta^*}(u, v) = H_{\theta^*}(u, v) P(u, v). \quad (10.191)$$

The geometric mean of the pair of conjugate planar images is given as

$$\begin{aligned} P_g(u, v) &= [H_{\theta}(u, v) P(u, v) H_{\theta^*}(u, v) P(u, v)]^{1/2} \\ &= [H_{\theta}(u, v) H_{\theta^*}(u, v)]^{1/2} P(u, v) \\ &= H_g(u, v) P(u, v). \end{aligned} \quad (10.192)$$

Therefore, the geometric mean of a pair of conjugate planar images of a planar source is equal to the original source distribution blurred by an MTF that is given by the geometric mean of the individual MTFs. This implies that geometric means of pairs of conjugate planar images may be deconvolved by using the geometric mean of the corresponding MTFs. In practice, it may be appropriate to assume that the MTF is independent of the angular position of the camera(s); then, the same averaged MTF may be used for all angles.

Pixel-by-pixel geometric averaging of opposing projections before prereconstruction restoration can improve the performance of the restoration filter because it reduces the space-variance of the blur; furthermore, the averaging procedure reduces the effects of scatter and attenuation. The averaging technique is applicable when the object to be restored is of medium size, over which the averaged PSF (or MTF) may be assumed to be space-invariant (see Figure 10.23). Prereconstruction restoration requires large computing resources because each projection image needs to be restored. Averaging reduces the filtering time for restoration by 50%, because each opposing pair of projections is replaced by a single image. Geometric averaging of opposing projections performs well in applications where the object is not located in a corner of the field of view of the camera. It is required that projections be acquired through the full range of $0^\circ - 360^\circ$.

Geometric averaging reduces the shift-variance of the blur function, but does not completely eliminate the variance. Therefore, artifacts due to the shift-variance of the blur function may remain in regions situated close to the edges of the field of view of the camera. Prereconstruction restoration filtering assumes that the blur function for all averaged projections is the same as for points located at the axis of rotation of the camera. Geometric averaging and prereconstruction restoration procedures are well-suited to SPECT imaging

of the brain, where the image is centered and does not occupy the full field of view of the camera.

Examples of the application of geometric averaging as a preprocessing step prior to the restoration of SPECT images are presented in Section 10.5.6.

10.5.6 Examples of restoration of SPECT images

Boulfelfel et al. [86, 87, 749, 750, 751, 851, 935, 948] conducted several studies on the restoration of SPECT images using the Wiener, PSE, Metz, and Kalman filters, including the options of prereconstruction restoration, post-reconstruction restoration, and geometric averaging of the projections, as well as the application of the filters in 2D or 3D. The MTFs of the imaging systems used were experimentally measured using a line source to obtain the LSF for various imaging parameters and configurations; see Sections 2.9 and 2.12. Some of their experiments and results are described in the following paragraphs.

Images of a tubular phantom: A tubular phantom was constructed with an acrylic tube of length 40 cm and internal diameter 35 mm, within which was introduced a solid acrylic rod of diameter 15 mm. The tube was filled with a solution containing ^{99m}Tc . Sixty planar projections, each of size 64×64 pixels, were acquired, spanning the angular range of $[0^\circ, 180^\circ]$, with the source-to-collimator distances of 5 cm and 20 cm, using a Siemens Rota camera with a low-energy all-purpose collimator. No scattering medium was used around the phantom (other than the ambient atmosphere). SPECT images of size 64×64 pixels each were reconstructed using the FBP algorithm.

Models of the PSD of the phantom were mathematically derived, for both planar and SPECT imaging, from the known geometry of the phantom. Blob-like artifacts could arise in filtered images due to the discontinuities that are present in discrete mathematical models as above [950]. In order to prevent such artifacts, the model PSDs were smoothed with a Gaussian window [951] of the form

$$W(r) = \exp \left[-\frac{1}{2} \left(\alpha \frac{r}{N/2} \right)^2 \right], \quad (10.193)$$

where $r = \sqrt{u^2 + v^2}$ is the index of radial frequency in the 2D Fourier space, N is the width of the PSD array in pixels, and α is a scale factor. Under the condition of MMSE of the restored images of known test images, the optimal value of α was determined to be 0.4. The PSD of the noise was derived from the known total count in the image being restored.

Due to the small size of the phantom (in cross-section), and due to its placement at the center of the field of imaging with respect to the axis of rotation of the gamma camera, it is reasonable to assume that the PSF of the imaging system is stationary. Hence, shift-invariant filters may be applied for restoration.

The Wiener and PSE filters were used for prereconstruction restoration of the planar images; in post-reconstruction restoration, the same filters were also applied to the SPECT image. Two sample projection images of the phantom are shown in Figure 10.26 for source-to-collimator distances of 5 *cm* and 20 *cm*; it is evident that the latter image is blurred to a greater extent than the former. The restored version of the planar image with the source-to-collimator distance of 20 *cm*, using the Wiener filter, is shown in part (c) of Figure 10.26; profiles of the original image and the restored image are shown in Figure 10.27. The restored image clearly demonstrates the expected reduction in counts at the center of the image (due to the solid rod).

Figure 10.28 shows the original and restored SPECT images for various cases of filtering. Prereconstruction restoration gave better results than post-reconstruction restoration. The Gaussian window applied to the object model effectively reduced the hot spots (blobs) seen in the restored images without the window. Observe that the blobs appear with significant amplitude in the post-reconstruction restored images, but are reduced in the prereconstruction restored images. RMS error values were computed for the original and restored images with respect to the ideal (known) cross-section of the phantom (shown in Figure 10.28 b). It was found that, whereas the acquired image had an RMS error of 69.12, the error for the post-reconstruction Wiener-restored image was 48.58, and that for the prereconstruction Wiener-restored image was 26.52; application of the Gaussian window to the model PSD further decreased the error to 23.37. The results of the PSE filter and Metz filter (not shown) had comparable RMS errors.

Images of a cardiac phantom: A cardiac phantom (Data Spectrum Corporation [912]) was used to simulate myocardial perfusion images with “defect” inserts to simulate myocardial ischemia and infarction. A sectional view showing the dimensions of the phantom is illustrated in Figure 10.29. The phantom consists of two U-shaped barrels of diameter 4.1 *cm* and 6.1 *cm*. The space between the barrels was filled with 37 *MBq* (1 *mCi*) of ^{201}Tl -chloride. Several types of defect arrangements were used; Figure 10.29 shows, in cross-section, a case with a 45° solid defect and a 45° defect with 50% activity. The phantom was held inclined at 45° with respect to the axis of rotation of the Rota camera with the low-energy all-purpose collimator, and 60 projections of size 64 × 64 pixels each, were acquired over 180°. The acquisition time for each projection was 5 *s*, and the radius of rotation was 20 *cm*. The SPECT images of different transaxial slices were reconstructed, and the 3D information was used to derive oblique sectional images at 45° using the Siemens Micro-Delta software. The total count for each image was between 15,000 and 17,000, which corresponds to low-statistics images in clinical studies.

In performing restoration of the projection images of the cardiac phantom, an object PSD model needs to be derived. If the model is derived directly from the physical shape of the phantom, it will match the actual phantom images, but will not be applicable for the restoration of real myocardial images with unknown defects. A model that matches closely the general case is a hollow

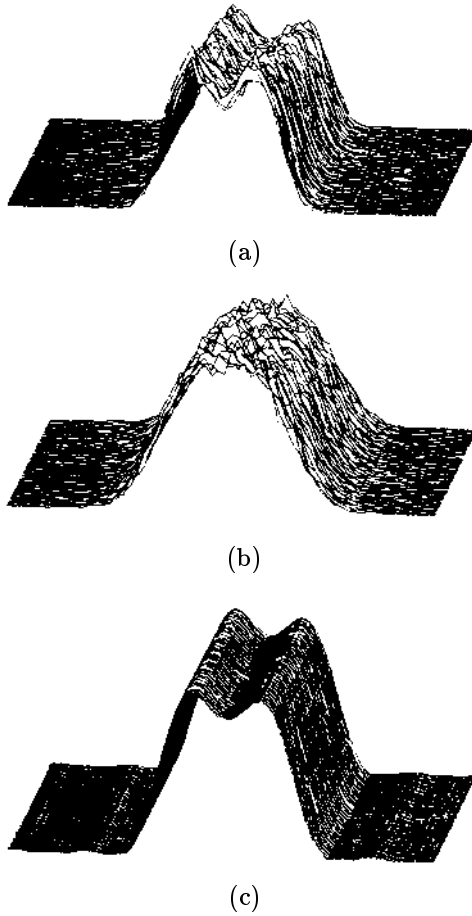


FIGURE 10.26

Projection (planar) images of the tubular phantom with the source-to-collimator distance of (a) 5 cm and (b) 20 cm. (c) Result of Wiener restoration of the image in (b). Figures courtesy of D. Boulfefel [86]. See also Figures 10.27 and 10.28.




FIGURE 10.27

Profiles of planar images of the tubular phantom across the center. Solid line: ideal (true) profile. Dotted line: profile of the acquired planar image with the source-to-collimator distance of 20 cm; see Figure 10.26 (b). Dashed line: profile of the Wiener-filtered planar image; see Figure 10.26 (c). Figure courtesy of D. Boulfefel [86].

cylinder of appropriate length inclined in the same way as the phantom. In the works of Boulfefel et al., for each projection, a rotation was performed on such a model, and the Radon integral was computed to model the projection.

The acquired SPECT image of the phantom, for the defect arrangement illustrated in Figure 10.29, as well as the restored images obtained by pre-reconstruction restoration using the Wiener and PSE filters, are shown in Figure 10.30. The contrast values of the defect regions were computed as defined in Equation 2.7. Groups of pixels representing the defect regions as well as suitable background regions were selected manually over 16×16 subimages using the physical model of the phantom as depicted in Figure 10.29. The restoration filters resulted in more than doubling of the contrast values (with respect to the contrast of the same defect in the acquired image) [86].

Cardiac SPECT images: One of the major applications of nuclear medicine imaging is in cardiology. With the use of ^{201}Tl , myocardial diseases and defects such as ischemia and necrosis are detected in SPECT slices as cold spots of reduced tracer concentration. Myocardial perfusion SPECT images, in the short-axis view (see Figure 1.27 a), appear similar to the rings of activity in the images of the tubular phantom in Figure 10.28. The cardiac phantom, described in the preceding paragraphs, may also be used to simulate myocardial perfusion SPECT images related to various pathological conditions; see Figures 10.29 and 10.30.






FIGURE 10.28

(a) SPECT image of a tubular phantom; size 64×64 pixels; RMS error = 69.12 with respect to the known, ideal version of the image, shown in (b). Post-reconstruction restoration of the image in (a) using (c) the Wiener filter (RMS error = 48.58), and (d) the PSE filter (RMS error = 54.12). SPECT image after preconstruction restoration of the planar images using (e) the Wiener filter (RMS error = 26.52), and (f) the PSE filter (RMS error = 29.31). (g) and (h) correspond to (e) and (f) with the inclusion of Gaussian smoothing of the image PSD model (RMS errors = 23.37 and 26.57, respectively). Images courtesy of D. Boulfefel [86].

**FIGURE 10.29**

Schematic representation of the cardiac phantom (Data Spectrum Corporation [912]) used to simulate myocardial perfusion images with “defect” inserts to simulate myocardial ischemia and infarction. The dimensions shown are in *mm*.

**FIGURE 10.30**

(a) Acquired SPECT image of the cardiac phantom with the defect arrangement illustrated in Figure 10.29. Result of pre-reconstruction restoration using (b) the Wiener and (c) PSE filter. Reproduced with permission from D. Boulfelfel, R.M. Rangayyan, L.J. Hahn, and R. Kloiber, “Pre-reconstruction restoration of single photon emission computed tomography images”, *IEEE Transactions on Medical Imaging*, 11(3): 336 – 341, 1992. © IEEE.


Unfortunately, myocardial SPECT images possess poor statistics, because only a small fraction of the injected activity will accumulate in the myocardium. Furthermore, as the peak photon energy of ^{201}Tl is about 80 keV, scattering has a serious effect on image quality. Boulfefel et al. [86, 751] applied prereconstruction restoration and post-reconstruction restoration techniques using the Wiener, PSE, and Metz filters to myocardial SPECT images; examples from their works are presented and discussed in the following paragraphs.

In the procedure to acquire myocardial SPECT images of human patients, 74 MBq (2 mCi) of ^{201}Tl was injected into the body. After accumulation of the tracer in the myocardium, 44 planar projections, each of size 64×64 pixels, spanning the full range of $[0^\circ, 180^\circ]$, were acquired. The time for the acquisition of each projection was 30 s. Each projection image had a total count in the range 10,000 to 20,000. The projections were acquired in an elliptic trajectory with the average distance from the heart being about 20 cm. Two energy peaks were used in the acquisition of the projections in order to perform scatter correction using the dual-energy-window subtraction technique. No attenuation correction was performed as the organ is small. Given the imaging protocol as above, and the fact that the organ being imaged is small, it is valid to assume that the blurring function is nearly shift-invariant; hence, it becomes possible to apply shift-invariant filters, such as the Wiener, PSE, and Metz filters, for the restoration of myocardial planar and SPECT images.

Figures 10.31 and 10.32 show one projection image each of two patients. Transverse SPECT images were reconstructed, and oblique slices perpendicular to the long axis of the heart were then computed from the 3D data available. Figures 10.31 and 10.32 show one representative oblique section image in each case, along with several restored versions of the images. The parts of the myocardium with reduced activity (cold spots) are seen more clearly in the restored images than in the original images. The results of prereconstruction restoration applied to the planar images are better than those of post-reconstruction restoration filtering in terms of noise content as well as improvement in sharpness and clarity.

SPECT images of the brain: Radionuclide brain scanning has been used extensively in the study of neurological and psychiatric diseases. The main area of application is the detection of pathology in the cerebral hemispheres and the cerebellum. A number of radiopharmaceuticals are used for brain scanning; however, $^{99\text{m}}\text{Tc}$ -based materials are most widely used.

An advantage in brain imaging is that the patient's head can be positioned at the center of rotation of the camera, which allows imaging over 360° , rather than over only 180° as in the case of myocardial imaging. Although the brain may be considered to be a large organ, the homogeneity of the (scattering) medium and the ability to image it from a short distance using a circular orbit allow the use of geometric averaging of the planar images as a preprocessing

**FIGURE 10.31**

Top: A sample planar projection image of a patient. (a) Short-axis SPECT image showing the myocardium of the left ventricle in cross-section. Results of post-reconstruction restoration applied to the SPECT image using (b) the Wiener, (c) the PSE, and (d) Metz filters. Results of prereconstruction restoration applied to the planar images using (e) the Wiener, (f) the PSE, and (g) Metz filters. Images courtesy of D. Boulfefel, L.J. Hahn, and R. Kloiber, Foothills Hospital, Calgary [86].




FIGURE 10.32

Top: A sample planar projection image of a patient. (a) Short-axis SPECT image showing the myocardium of the left ventricle in cross-section. Results of post-reconstruction restoration applied to the SPECT image using (b) the Wiener, (c) the PSE, and (d) Metz filters. Results of prereconstruction restoration applied to the planar images using (e) the Wiener, (f) the PSE, and (g) Metz filters. Images courtesy of D. Boulfefel, L.J. Hahn, and R. Kloiber, Foothills Hospital, Calgary [86].

step to reduce both attenuation and shift-variance of the blur before restoration. Brain images are also not as low in statistics as myocardial images.

In the procedure for nuclear medicine imaging of the brain, after the ^{99m}Tc -chloride administered to the patient had accumulated in the brain, 44 planar projections, each of size 64×64 pixels, were acquired. The time for the acquisition of each projection was 30 s. The projections were acquired over the full range of 360° in a circular trajectory with the radius of rotation of 20 cm. Two energy peaks were used in the acquisition of the projections to perform scatter correction using the dual-energy-window subtraction technique.

Figures 10.33 and 10.34 show a set of opposing projection images as well as their geometric mean for two patients. Transverse SPECT images were reconstructed after performing geometric averaging of conjugate projections and (prereconstruction) restoration using the Wiener, PSE, and Metz filters [86]. Figures 10.33 and 10.34 show one representative SPECT image in each case, along with several restored versions. The results show that averaging of conjugate projections improves the quality of the restored images, which are sharper than the images restored without averaging.

Images of a resolution phantom: Boulfefel et al. [86, 87, 749, 750, 935] conducted several restoration experiments with SPECT images of a “resolution” phantom. The phantom contains nine pairs of hot spots of diameters 39, 22, 17, 14, 12, 9, 8, 6, and 5 mm in the “hot lesion” insert (Nuclear Associates), with a total diameter of 200 mm; see Figure 3.68 for related illustrations. The phantom was filled with 1 mCi of ^{201}Tl -chloride, centered at the axis of rotation of the gamma camera at a distance of 217 mm, and 120 projections, each of size 128×128 pixels, were acquired over 360° . SPECT images of different transaxial slices were reconstructed using the Siemens Micro-Delta software.

Given the large size of the phantom, it would be inappropriate to assume that the degradation phenomena are shift-invariant. Boulfefel et al. [750, 948] applied the Kalman filter for restoration of SPECT images of the resolution phantom. Figure 10.35 shows a representative SPECT image of the phantom, along with (post-reconstruction) restoration of the image using the Kalman filter; the results of application of the shift-invariant Wiener and PSE filters are also shown for comparison. It is evident that the shift-variant Kalman filter has provided better results than the other filters: the Kalman-restored image clearly shows seven of the nine pairs of hot spots, whereas the results of the Wiener and PSE filters show only four or five pairs. For the sake of comparison, the results of prereconstruction restoration of the resolution phantom image obtained by applying the shift-invariant Wiener, PSE, and Metz filters after geometric averaging of conjugate projections are shown in Figure 10.36. Observe that the orientation of these results is different from that of the images in Figure 10.35 due to the alignment procedure required for averaging. Although the results show some of the hot spots with more clarity than the original image in Figure 10.35 (a), they are of lower quality than the result of Kalman filtering, shown in Figure 10.35 (d).





FIGURE 10.33

Top row: A sample pair of conjugate projections of a patient, along with their geometric mean. (a) SPECT image showing the brain in cross-section. Results of prereconstruction restoration applied to the planar images using (b) the Wiener, (c) the PSE, and (d) Metz filters. Results of geometric averaging and prereconstruction restoration applied to the planar images using (e) the Wiener, (f) the PSE, and (g) Metz filters. The orientation of the images in (e) - (g) is different from that of the images in (a) - (d) due to the alignment of conjugate projection images for geometric averaging. Images courtesy of D. Boulfelfel, L.J. Hahn, and R. Kloiber, Foothills Hospital, Calgary [86].

**FIGURE 10.34**

Top row: A sample pair of conjugate projections of a patient, along with their geometric mean. (a) SPECT image showing the brain in cross-section. Results of prereconstruction restoration applied to the planar images using (b) the Wiener, (c) the PSE, and (d) Metz filters. Results of geometric averaging and prereconstruction restoration applied to the planar images using (e) the Wiener, (f) the PSE, and (g) Metz filters. The orientation of the images in (e) – (g) is different from that of the images in (a) – (d) due to the alignment of conjugate projection images for geometric averaging. Images courtesy of D. Boulfefel, L.J. Hahn, and R. Kloiber, Foothills Hospital, Calgary [86].




FIGURE 10.35

(a) Acquired SPECT image (128×128 pixels) of the resolution phantom. Post-reconstruction restored versions using (b) the Wiener filter; (c) the PSE filter; and (d) the Kalman filter. The images (a) – (c) were enhanced by gamma correction with $\gamma = 0.8$; the image (d) was enhanced with $\gamma = 0.3$ (see Section 4.4.3). See also Figure 3.68. Reproduced with permission from D. Boulfefel, R.M. Rangayyan, L.J. Hahn, R. Kloiber, and G.R. Kuduvalli, “Restoration of single photon emission computed tomography images by the Kalman filter”, *IEEE Transactions on Medical Imaging*, 13(1): 102 – 109, 1994. © IEEE.




FIGURE 10.36

Prereconstruction restoration of the SPECT image of the resolution phantom shown in Figure 10.35 (a) after geometric averaging of conjugate projection images, using (a) the Wiener, (b) the PSE, and (c) the Metz filters. The orientation of the images in this figure is different from that of the images in Figure 10.35 due to the alignment of conjugate projection images for geometric averaging. Images courtesy of D. Boulfefel, L.J. Hahn, and R. Kloiber, Foothills Hospital, Calgary [86].

SPECT images of the liver and spleen: Liver and spleen images are difficult to restore because of their large size and irregular shape. The liver and spleen are imaged together when radiopharmaceuticals that are trapped by the reticulo-endothelial cell system are used. The most commonly used radiopharmaceutical for this purpose is a ^{99m}Tc -based label. In the procedure for imaging the liver and spleen, 2 mCi of a ^{99m}Tc -based radiopharmaceutical was given to the patient. After the isotope accumulated in the liver and spleen, 44 projections, each of size 64×64 pixels, were acquired. The time for the acquisition of each projection was 40 s . The projections were acquired over the full range of 360° in a circular trajectory, with the average radius of rotation of 25 cm . Two energy peaks were used in the acquisition of the projections in order to perform scatter correction using the dual-energy-window subtraction technique. Transverse SPECT images were reconstructed after averaging and correcting for attenuation using the Siemens Micro-Delta processor. The Chang algorithm was used for attenuation correction.

Figures 10.37 and 10.38 show a sample SPECT slice of the liver and spleen of two patients, along with its restored version using the Kalman filter. The restored images demonstrate the full outlines of the liver and spleen with improved clarity, and show a few cold spots within the organs with increased contrast as compared to the original images. The clinical validity of this observation was not confirmed.

3D restoration of SPECT images: Boulfefel et al. [86, 750, 948, 935] applied 3D filters for the restoration of SPECT images, including 3D extensions of the Wiener, PSE, and Metz filters, as well as a combination of a 2D Kalman filter in the SPECT plane and a 1D Metz filter in the inter-slice direction. Figures 10.39 and 10.40 show a sample planar image of the liver and




FIGURE 10.37

(a) Acquired SPECT image of the liver and spleen of a patient. (b) Restored image obtained by the application of the Kalman filter. Images courtesy of D. Boulfefel, L.J. Hahn, and R. Kloiber, Foothills Hospital, Calgary [86].




FIGURE 10.38

(a) Acquired SPECT image of the liver and spleen of a patient. (b) Restored image obtained by the application of the Kalman filter. Images courtesy of D. Boulfefel, L.J. Hahn, and R. Kloiber, Foothills Hospital, Calgary [86].

spleen each of two patients, a sample SPECT image in each case, and restored images of the SPECT slices after 3D restoration of the entire SPECT volumes using the Wiener, PSE, and Metz filters. Figures 10.41 and 10.42 show a sample SPECT image and the corresponding restored image after 3D restoration of the entire SPECT volume using the 2D Kalman filter in the SPECT plane and a 1D Metz filter in the inter-slice direction. The restored images show more cold spots within the liver, with increased contrast; however, the clinical validity of this observation was not confirmed. A sample SPECT image and the corresponding restored version after 3D restoration of the entire SPECT volume using the Kalman-Metz filter combination as above are shown in Figure 10.43. Compared to the result of 2D filtering shown in Figure 10.35, the 3D filtering procedure appears to have yielded a better image.

10.6 Remarks

The widespread occurrence of image degradation in even the most sophisticated and expensive imaging systems has continually frustrated and challenged researchers in imaging and image processing. The field of image restoration has attracted a high level of activity from researchers with several different perspectives [8, 11, 822, 823, 824, 825, 952, 953, 954, 955]. In this chapter, we have studied a small selection of techniques that are among the popular approaches to this intriguing problem. Most of the restoration techniques require detailed and specific information about the original undegraded image and the degradation phenomena. Several additional constraints may also be applied, based upon a priori and independent knowledge about the desired image. However, it is often difficult to obtain accurate information as above. The quality of the result obtained is affected by the accuracy of the information provided and the appropriateness of the constraints applied. The nature of the problem is characterized very well by the title of a special meeting held on this subject: “Signal recovery and synthesis with incomplete information and partial constraints” [954, 955]. Regardless of the difficulties and challenges involved, researchers in the field of image restoration have demonstrated that a good understanding of the problem can often lead to usable solutions.




FIGURE 10.39

Top: A sample planar projection image of a patient. (a) SPECT image showing the liver and spleen. Results of post-reconstruction 3D restoration applied to the entire SPECT volume using (b) the Wiener, (c) the PSE, and (d) Metz filters. Images courtesy of D. Boulfefel, L.J. Hahn, and R. Kloiber, Foothills Hospital, Calgary [86].




FIGURE 10.40

Top: A sample planar projection image of a patient. (a) SPECT image showing the liver and spleen. Results of post-reconstruction 3D restoration applied to the entire SPECT volume using (b) the Wiener, (c) the PSE, and (d) Metz filters. Images courtesy of D. Boulfefel, L.J. Hahn, and R. Kloiber, Foothills Hospital, Calgary [86].




FIGURE 10.41

(a) Acquired SPECT image of the liver and spleen of a patient. (b) Restored image obtained by the application of the 3D Kalman-Metz combined filter. Images courtesy of D. Boulfefel, L.J. Hahn, and R. Kloiber, Foothills Hospital, Calgary [86].




FIGURE 10.42

(a) Acquired SPECT image of the liver and spleen of a patient. (b) Restored image obtained by the application of the 3D Kalman-Metz combined filter. Images courtesy of D. Boulfefel, L.J. Hahn, and R. Kloiber, Foothills Hospital, Calgary [86].




FIGURE 10.43

(a) Acquired SPECT image of the resolution phantom. (b) Restored image obtained by the application of the 3D Kalman-Metz combined filter. Images courtesy of D. Boulfefel, L.J. Hahn, and R. Kloiber, Foothills Hospital, Calgary [86].

10.7 Study Questions and Problems

Selected data files related to some of the problems and exercises are available at the site

www.enel.ucalgary.ca/People/Ranga/enel697

1. Using mathematical expressions and operations as required, explain how a degraded image of an edge may be used to derive the MTF $H(u, v)$ of an imaging system. State clearly any assumptions made, and explain their relevance or significance.
2. Given $\mathbf{g} = \mathbf{h}\mathbf{f} + \boldsymbol{\eta}$ and $\tilde{\mathbf{f}} = \mathbf{L}\mathbf{g}$, where \mathbf{g} is a degraded image, \mathbf{f} is the original image, $\boldsymbol{\eta}$ is the noise process, \mathbf{h} is the PSF of the blurring system, $\tilde{\mathbf{f}}$ is the restored image, and \mathbf{L} is the restoration filter, expand $\epsilon^2 = E \left\{ \text{Tr} \left[(\mathbf{f} - \tilde{\mathbf{f}}) (\mathbf{f} - \tilde{\mathbf{f}})^T \right] \right\}$ and simplify the result to contain only \mathbf{h} , \mathbf{L} , and the ACF matrices of \mathbf{f} and $\boldsymbol{\eta}$. Give the reasons for each step, and explain the significance and implications of the assumptions made in deriving the Wiener filter.
3. With reference to the Wiener filter for image restoration (deblurring), explain the role of the signal-to-noise spectral ratio. How does this ratio control the performance of the filter?
4. Prove that the PSE filter is the geometric mean of the inverse and Wiener filters.
5. List the various items of information required in order to implement the Wiener filter for deblurring a noisy image. Explain how you would derive each item in practice.

10.8 Laboratory Exercises and Projects

1. Create or acquire a test image including components with sharp edges. Blur the image by convolution with a Gaussian PSF. Add Gaussian-distributed random noise to the blurred image.

Derive the MTF of the blurring function and the PSD of the noise. Pay attention to the scale factors involved in the Fourier transform.

Restore the degraded image using the inverse, Wiener, and PSE filters. You may have to restrict the inverse filter to a certain frequency limit in order to prevent the amplification of noise.

How would you derive or model the ideal object PSD required in the design of the Wiener and PSE filters?

2. Using a camera that is not in focus, capture a blurred image of a test image containing a sharp line. Derive the PSF and the MTF of the imaging system.
3. Using a camera that is not in focus, capture a blurred image of a scene, such as your laboratory, including a person and some equipment. Ensure that the scene includes an object with a sharp edge (for example, the edge of a door frame or a blackboard), as well as a uniform area (for example, a part of a clean wall or board with no texture).

Derive the PSF and MTF of the imaging system by manual segmentation of the edge spread function and further analysis as required. Estimate the noise PSD by using segments of areas expected to be uniform. Design the Wiener and PSE filters and restore the image.

How would you derive or model the ideal PSD of the original scene?

4. Restore the image in the preceding exercise by designing the blind deblurring version of the PSE filter.

Image Coding and Data Compression

High spatial resolution and fine gray-scale quantization are often required in biomedical imaging. Digital mammograms are typically represented in arrays of $4,096 \times 4,096$ pixels with 12 *b/pixel*, leading to raw-data files of the order of 32 MB per image. Volumetric data obtained by CT and MRI could be of size $512 \times 512 \times 64$ voxels with 16 *b/voxel*, occupying 32 MB per examination. Patients with undetermined or multiple complications may undergo several examinations via different modalities such as X-ray imaging, ultrasound scanning, CT scanning, and nuclear medicine imaging, resulting in large collections of image files.

Most health-care jurisdictions require medical records, including images, of adults to be stored for durations of the order of seven years from the date of acquisition. Children's records and images are required to be maintained until at least the time they reach adulthood.

With the view to improve the efficiency of storage and access, several imaging centers and hospitals have moved away from film-based storage toward electronic storage. Furthermore, most medical imaging systems have moved to direct digital image acquisition with adequate resolution, putting aside the debate on the quality of an original film-based image versus that of its scanned (digitized) representation. Since 1980, an entire series of conferences has been dedicated to PACS: see the PACS volumes of the SPIE Medical Imaging conference series [956]. Networks and systems for PACS are integrated into the infrastructure of most modern hospitals. The major advantages and disadvantages of digital and film-based archival systems are listed below.

- Films deteriorate with age and handling. Digital images are unaffected by these factors.
- Despite elaborate indexing schemes, films tend to get lost or misplaced. Digital image files are less likely to face these problems.
- Digital image files may be accessed simultaneously by several users. Although multiple copies of film-based images may be made, it would be an expensive option that adds storage and handling complexities.
- With the proliferation of computers, digital images may be viewed and manipulated at several convenient locations, including a surgical suite, a patient's bedside, and one's home or office. Viewing film-based im-

ages with detailed attention requires specialized viewing consoles under controlled lighting conditions.

- Digital PACS require significant initial capital outlay, as well as routine maintenance and upgrading of the computer, storage, and communication systems. However, these costs may be offset by the savings in the continuing costs of film, as well as the associated chemical processing systems and disposal. The environmental concerns related to film processing are also removed by digital PACS.
- *Digital images may be compressed via image coding and data compression techniques so as to occupy less storage space.*

The final point above forms the topic of the present chapter.

Although the discussion above has been in the context of image storage or archival, similar concerns regarding the size of image files and the desirability of compression arise in the communication of image data. In this chapter, we shall study the basic concepts of information theory that apply to image coding, compression, and communication. We shall investigate several techniques for encoding image data, including decorrelation procedures to modify the statistical characteristics of the data so as to permit efficient representation, coding, and compression.

The representation of the significant aspects of an image in terms of a small number of numerical features for the purpose of pattern classification may also be viewed as image coding or data compression; however, we shall treat this topic separately (see Chapter 12).

11.1 Considerations Based on Information Theory

Image data compression is possible due to the following basic characteristics:

- Code redundancy — all code words (pixel values) do not occur with equal probability.
- Spatial redundancy — the values of neighboring pixels tend to lie within a small dynamic range, and exhibit a high level of correlation.
- Psychovisual redundancy — human analysts can recognize the essential nature and components of an image from severely reduced versions such as caricatures, edges, and regions, and need not (or do not) pay attention to precise numerical values.

Information-theoretic considerations are based upon the notion of information as related to the statistical uncertainty of the occurrence of an event

(such as a signal, an image, or a pixel value), rather than the structural, symbolic, pictorial, semantic, or diagnostic content of the entity. The measure of entropy is based upon the probabilities of occurrence of the various symbols involved in the representation of a message or image: see Section 2.8. Despite the mathematical and theoretical powers of measures such as entropy, the standpoint of viewing an image as being composed of discrete and independent symbols (numerical values) removes the analyst from the real-world and physical properties of the image. The use of the underlying assumptions also lead to severe limitations in entropy-based source coding, with lossless compression factors often limited to the order of 2:1. Additional techniques based upon decorrelation of the image data via the identification and modeling of the underlying image-generation phenomena, or the use of pattern recognition techniques, could assist in improving the performance of image compression procedures.

11.1.1 Noiseless coding theorem for binary transmission

Given a code with an alphabet of two symbols and a source A with an alphabet of two symbols, the average length of the code words per source symbol may be made arbitrarily close to the lower bound (entropy) $H(A)$ by encoding sequences of source symbols instead of encoding individual symbols [9, 126]. The average length $L(n)$ of encoded n -symbol sequences is bounded by

$$H(A) \leq \frac{L(n)}{n} \leq H(A) + \frac{1}{n}. \quad (11.1)$$

Difficulties exist in estimating the true entropy of a source due to the fact that pixels are statistically dependent, that is, correlated, from pixel to pixel, row to row, and frame to frame of real-life images. The computation of the true entropy requires that symbols be considered in blocks over which the statistical dependence is negligible. In practice, this would translate to estimating joint PDFs of excessively long vectors. Values of entropy estimated with single pixels or small blocks of pixels would result in over-estimates of the source entropy. If blocks of pixels are chosen such that the sequence-entropy estimates converge rapidly to the limit, then block-coding methods may provide results close to the minimum length given by Equation 11.1. Run-length coding may be viewed as an adaptive block-coding technique; see Section 11.3.2.

11.1.2 Lossy versus lossless compression

A coding or compression method is considered to be *lossless* if the original image data can be recovered, with no error, from the coded and compressed data. Such a technique may also be referred to as a *reversible*, bit-preserving, or error-free compression technique.

A compression technique becomes *lossy* or *irreversible* if the original data cannot be recovered, with complete pixel-by-pixel numerical accuracy, from the compressed data. In the case of images, the human visual system can tolerate significant numerical differences or error, in the sense that the degraded image recovered from the compressed data is perceived to be *essentially the same as the original image*. This arises from the fact that a human observer will, typically, not examine the numerical values of individual pixels, but instead assess the semantic or pictorial information conveyed by the data. Furthermore, a human analyst may tolerate more error, noise, or distortion in the uniform areas of an image than around its edges that attract visual attention. Data compression techniques may be designed to exploit these aspects to gain significant advantages in terms of highly compressed representation, with high levels of loss of numerical accuracy while remaining *perceptually lossless*. On the same token, in medical imaging, if the numerical errors in the retrieved and reconstructed images do not cause any change in the diagnostic results obtained by using the degraded images, one could achieve high levels of numerically lossy compression while remaining *diagnostically lossless*.

In the quest to push the limits of numerically lossy compression techniques while remaining practically lossless under some criterion, the question arises as to the worth of such practice. Medical practice in the present highly litigious society could face large financial penalties and loss due to errors. Radiological diagnosis is often based upon the detection of minor deviations from the normal (or average) patterns expected in medical images. If a lossy data compression technique were to cause such a faint deviation to be less perceptible in the compressed (and reconstructed) image than in the original image, and the diagnosis based upon the reconstructed image were to be in error, the financial compensation to be paid would cost several times the amount saved in data storage; the loss in professional standing and public confidence could be even more damaging. In addition, defining the fidelity of representation in terms of the closeness to the original image or distortion measures is a difficult and evasive activity. Given the high levels of the professional care and concern, as well as the fiscal and emotional investment, that are part of medical image acquisition procedures, it would be undesirable to use a subsequent procedure that could cause any degradation of the image. In this spirit, only lossless coding and compression techniques will be described in the present chapter. Regardless, it should be noted that any lossy compression technique may be made lossless by providing the numerical error between the original image and the degraded image reconstructed from the compressed data. Although this step will lead to additional storage or transmission requirements, the approach can facilitate the rapid retrieval or transmission of an initial, low-quality image, followed by completely lossless recovery: such a procedure is known as *progressive transmission*, especially when performed over multiple stages of image quality or fidelity.

11.1.3 Distortion measures and fidelity criteria

Although we have stated our interest in lossless coding of biomedical images, other processes, such as the transmission of large quantities of data over noisy channels, may lead to some errors in the received images. Hence, it would be relevant to consider the characterization of the distortion so introduced, and analyze the fidelity of the received image with respect to the original [9].

The *binary symmetric channel* is characterized by a single parameter: the bit-error probability p (see Figure 11.1). The channel capacity is given by

$$C = 1 + p \log p + q \log q, \quad (11.2)$$

where $q = 1 - p$.




FIGURE 11.1

Transmission error probabilities in a binary symmetric channel [9].

The least-squares single-letter fidelity criterion is defined as [9]

$$\rho_n(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \sum_{l=1}^n (x_l - y_l)^2 2^{l-1}, \quad (11.3)$$

where \mathbf{x} and \mathbf{y} are the transmitted and received n -bit vectors (blocks or words), respectively.

The Hamming distance between the vectors \mathbf{x} and \mathbf{y} is defined as

$$D_H(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \sum_{l=1}^n (x_l - y_l)^2. \quad (11.4)$$

Measures of fidelity may also be defined based upon entire images by defining an error image as

$$e(m, n) = g(m, n) - f(m, n), \quad (11.5)$$

where $g(m, n)$ is the received (degraded) version of the original (transmitted) image $f(m, n)$, and then defining the RMS value of the error as

$$e_{\text{RMS}} = \sqrt{\frac{1}{N^2} \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} [g(m, n) - f(m, n)]^2}, \quad (11.6)$$

or SNR as

$$\text{SNR} = \frac{\sum_{m=0}^{N-1} \sum_{n=0}^{N-1} g^2(m, n)}{\sum_{m=0}^{N-1} \sum_{n=0}^{N-1} e^2(m, n)}. \quad (11.7)$$

See Section 2.13 for more details on measures of SNR.

11.2 Fundamental Concepts of Coding

In general, coding could be defined as the use of symbols to represent information. The following list provides the definitions of a few basic terms and concepts related to coding [9]:

- An *alphabet* is a predefined set of symbols.
- A *word* is a finite sequence of symbols from an alphabet.
- A *code* is a mapping of words from a source alphabet into the words of a code alphabet.
- A code is said to be *distinct* if each code word is distinguishable from the other code words.
- A distinct code is *uniquely decodable* if every code word is identifiable when immersed in a sequence of code words (with no separators between the words).
- A desirable property of a uniquely decodable code is that it be decodable on a word-to-word basis. This is ensured if no code word may be a prefix to another; the code is then *instantaneously decodable*.
- A code is said to be *optimal* if it is instantaneously decodable and has the minimum average length for a given source PDF.

Examples of symbols are $\{0, 1\}$ in the binary alphabet; $\{0, 1, 2, 3, 4, 5, 6, 7\}$ in the octal system; $\{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$ in the decimal system; $\{0, 1, 2, 3, 4, 5, 6, 7, 8, 9, A, B, C, D, E, F\}$ in the hexadecimal system; $\{I, V, X, L, C, D, M\}$ in the Roman system (with the decimal equivalents of 1, 5, 10, 50, 100, 500, and 1,000, respectively); and $\{A - Z, a - z\}$ in the English alphabet (not

considering punctuation marks and special symbols). An example of a word in the context of image coding is 00001011 in 8 b binary coding, standing for the gray level 11 in the decimal system. Table 11.1 lists the codes for integers in the range $[0, 20]$ in the Roman, decimal, binary, Gray [957], octal, and hexadecimal codes [958]. The Gray code has the advantageous feature that only one digit is changed from one number to the next. Observe that, in general, all of the codes described here (including the English language) fail the conditions defined above for an optimal code.

11.3 Direct Source Coding

Pixels generated by real-life sources of images bear limitations in dynamic range and variability within a small spatial neighborhood. Therefore, codes used to represent pixel data at the source may be expected to demonstrate certain patterns of limited variation and high correlation. Furthermore, real-life sources of images do not generate random, uncorrelated values that are equally likely; instead, it is common to encounter PDFs of gray levels that are nonuniform. Some of these characteristics may be exploited to achieve efficient representation of images by designing coding systems tuned to specific properties of the source. Because the coding method is applied directly to pixel values generated by the source (without processing them by an algorithm to generate a different series of values), such techniques are categorized as direct source coding procedures.

11.3.1 Huffman coding

Huffman [9, 959] proposed a coding system to exploit the occurrence of some pixel values with higher probabilities than other pixels. The basic idea in Huffman coding is to use short code words for values with high probabilities of occurrence, and longer code words to represent values with lower probabilities of occurrence. This implies that the code words used will be of variable length; the method also presumes prior knowledge of the PDF of the source symbols (gray levels). It is required that the code words be uniquely decodable on a word-by-word basis, which implies that no code word may be a prefix to another. Huffman devised a coding scheme to meet these requirements and lead to average code-word lengths lower than that provided by fixed-length codes. Huffman coding provides an average code-word length L that is limited by the zeroth-order entropy of the source H_0 (see Equation 2.18) and $H_0 + 1$:

$$H_0 \leq L \leq H_0 + 1. \quad (11.8)$$

The procedure to generate the Huffman code is as follows [9, 959]:

TABLE 11.1

Integers in the Range [0, 20] in Several Alphabets or Codes [957, 958].

English	Portuguese	Roman	Decimal	Binary	Gray	Octal	Hex
Zero	Zero		0	00000	00000	000	0
One	Um/Uma	I	1	00001	00001	001	1
Two	Dois/Duas	II	2	00010	00011	002	2
Three	Três	III	3	00011	00010	003	3
Four	Quatro	IV	4	00100	00110	004	4
Five	Cinco	V	5	00101	00111	005	5
Six	Seis	VI	6	00110	00101	006	6
Seven	Sete	VII	7	00111	00100	007	7
Eight	Oito	VIII	8	01000	01100	010	8
Nine	Nove	IX	9	01001	01101	011	9
Ten	Dez	X	10	01010	01111	012	A
Eleven	Onze	XI	11	01011	01110	013	B
Twelve	Doze	XII	12	01100	01010	014	C
Thirteen	Treze	XIII	13	01101	01011	015	D
Fourteen	Catorze	XIV	14	01110	01001	016	E
Fifteen	Quinze	XV	15	01111	01000	017	F
Sixteen	Dezesseis	XVI	16	10000	11000	020	10
Seventeen	Dezessete	XVII	17	10001	11001	021	11
Eighteen	Dezoito	XVIII	18	10010	11011	022	12
Nineteen	Dezenove	XIX	19	10011	11010	023	13
Twenty	Vinte	XX	20	10100	11110	024	14


Leading zeros have been removed in the decimal and hexadecimal (Hex) codes, but retained in the binary, Gray, and octal codes.

1. Prepare a table listing the symbols (gray levels) in the source (image) sorted in decreasing order of the probabilities of their occurrence.
2. Combine the last two probabilities. The list of probabilities now has one less entry than before.
3. Copy the reduced list over to a new column, rearranging (as necessary) such that the probabilities are in decreasing order.
4. Repeat the procedure above until the list of probabilities is reduced to only two entries.
5. Assign the code digits 0 and 1 to the two entries in the final column of probabilities. (*Note:* There are two possibilities of this assignment that will lead to two different codes; however, their performance will be identical.)
6. Working backwards through the columns of probabilities, assign additional bits of 0 and 1 to the two entries that resulted in the last compounded entry in the column.
7. Repeat the procedure until the first column of probabilities is reached and all symbols have been assigned a code word.

It should be noted that a Huffman code is optimal for only the given source PDF; a change in the source PDF would require the design of a different code in order to be optimal. A disadvantage of the Huffman code is the increasing length of its code words, especially for sources with several symbols. The method does not perform any decorrelation of the data, and is limited in average code-word length by the zeroth-order entropy of the source.

Example: Figure 11.2 shows a 16×16 part of the image in Figure 2.1 (a), quantized to 3 b/pixel . The gray levels in the image are in the range $[0, 7]$, and would require 3 b/pixel with straight binary coding. The histogram of the image is shown in Figure 11.3; it is evident that some of the pixel values occur with low probabilities.

The procedure for accumulating the probabilities of occurrence of the source symbols is illustrated in Figure 11.4. The Huffman coding process is shown in Figure 11.5. Note that a different code with equivalent performance may be generated by reversing the order of assignment of the code symbols 0 and 1 at each step. The average code-word length is 2.69 b/pixel , which is slightly above the zeroth-order entropy of 2.65 b of the image. The advantage is relatively small due to the fact that the source in the example uses only eight symbols with 3 b/pixel , and has a relatively well-spread histogram (PDF). However, simple representation of the data using ASCII coding would require a minimum of 8 b/pixel ; the savings with reference to this requirement are significant. Larger advantages may be gained by Huffman coding of sources with more symbols and narrow PDFs.



```

1 1 1 1 1 1 1 1 1 1 2 3 2 2 1 2
0 1 1 1 1 1 1 1 1 1 1 2 2 3 4 5
1 0 0 0 1 1 1 1 1 1 1 1 2 2 4 6
2 2 3 5 4 3 1 0 1 1 1 1 1 2 3 5
4 6 5 4 3 1 1 2 2 1 1 1 1 1 2 4
5 5 2 1 2 3 2 2 2 3 3 4 3 2 1 3
4 3 1 2 1 1 1 2 2 2 1 2 2 2 3 5
2 0 2 0 1 3 1 3 5 3 3 2 2 3 3 6
1 1 2 2 1 2 1 2 3 3 3 4 4 6 5 6
1 1 2 4 1 0 0 1 3 4 5 5 5 4 4 6
1 1 1 4 2 1 2 3 5 5 5 4 4 3 4 6
1 1 1 4 4 4 5 6 6 5 4 3 2 3 5 6
1 1 2 5 5 4 5 5 4 3 3 2 3 4 5 6
2 1 4 5 5 5 5 4 3 1 1 1 4 6 5 6
2 2 5 5 5 4 3 2 2 1 1 4 6 6 6 7
4 4 4 4 3 2 2 1 0 1 5 6 6 6 6 7

```


```

111111111123221201111111111223451000111111112246223543101111235
4654311221111124552123222334321343121112221222352020131353322336
1122121233344656112410013455544611142123555443461114445665432356
1125545543323456214555543111465622555432211466674444322101566667

```

FIGURE 11.2

Top to bottom: A 16×16 part of the image in Figure 2.1 (a) quantized to 3 *b/pixel*, shown as an image, a 2D array, and as a string of integers with the gray-level values of every pixel. The line breaks in the string format have been included only for the sake of printing within the width of the page.

**FIGURE 11.3**

Gray-level histogram of the image in Figure 11.2. Zeroth-order entropy $H_0 = 2.65 b$.




FIGURE 11.4

Accumulation of probabilities (Prob.) of occurrence of gray levels in the derivation of the Huffman code for the image in Figure 11.2. The probabilities of occurrence of the symbols have been rounded to two decimal places, and add up to unity.

Symbol	Prob.	Code	Step 1	Step 2	Step 3	Step 4	Step 5	Step 6
1	0.30	<i>00</i>	0.30 <i>00</i>	0.30 <i>00</i>	0.30 <i>00</i>	0.30 <i>00</i>	0.44 <i>1</i>	0.56 <i>0</i>
2	0.19	<i>11</i>	0.19 <i>11</i>	0.19 <i>11</i>	0.25 <i>10</i>	0.26 <i>01</i>	0.30 <i>00</i>	0.44 <i>1</i>
4	0.13	<i>010</i>	0.13 <i>010</i>	0.13 <i>010</i>	0.19 <i>11</i>	0.25 <i>10</i>	0.26 <i>01</i>	
3	0.13	<i>011</i>	0.13 <i>011</i>	0.13 <i>011</i>	0.13 <i>010</i>	0.19 <i>11</i>		
5	0.12	<i>101</i>	0.12 <i>101</i>	0.13 <i>100</i>	0.13 <i>011</i>			
6	0.08	<i>1000</i>	0.08 <i>1000</i>	0.12 <i>101</i>				
0	0.04	<i>10010</i>	0.05 <i>1001</i>					
7	0.01	<i>10011</i>						

FIGURE 11.5

Steps in the derivation of the Huffman code for the image in Figure 11.2. (Prob. = probabilities of occurrence of the gray levels.) The binary words in bold italics are the Huffman code words at the various stages of their derivation. See also Figure 11.4.

Inter-pixel correlation may be taken into account in Huffman coding by considering combinations of pixels (gray levels) as symbols. If we were to consider pairs of gray levels in the example above, with gray levels quantized to 3 b/pixel , we would have a total of $8 \times 8 = 64$ possibilities; see Table 11.2. The first-order entropy of the image, considering pairs of gray levels, is $H_1 = 2.25 \text{ b}$; an average code-word length close to this value may be expected if Huffman coding is applied to pairs of gray levels.

TABLE 11.2

Counts of Occurrence of Pairs of Pixels in the Image in Figure 11.2.

Current pixel	Next pixel in the same row								
	0	1	2	3	4	5	6	7	
0	3	6	1	0	0	0	0	0	0
1	4	46	17	4	5	1	0	0	
2	2	12	16	12	3	2	0	0	
3	0	5	8	6	6	6	1	0	
4	0	1	1	10	8	6	7	0	
5	0	0	1	1	9	12	6	0	
6	0	0	0	0	0	4	6	2	
7	0	0	0	0	0	0	0	0	

For example, the pair (1, 2) occurs 17 times in the image. The last pixel in each row was not paired with any pixel. The first-order entropy of the image, considering the probabilities of occurrence of pairs of gray-level values as in Equation 2.23, is $H_1 = 2.25 \text{ b}$. The zeroth-order entropy is $H_0 = 2.65 \text{ b}$.

Although the performance of Huffman coding is limited when applied directly to source symbols, the method may be applied to decorrelated data with significant advantage, due to the highly nonuniform or concentrated PDFs of such data. The performance of Huffman coding as a post-encoder following decorrelation methods is discussed in several sections to follow.

11.3.2 Run-length coding

Images with high levels of correlation may be expected to contain strings of repeated occurrences of the same gray level: such strings are known as *runs*. Data compression may be achieved by coding such runs of gray levels. For example, the first three rows of the image in Figure 11.2 may be represented as follows:

Row 1: (1, 10), (2, 1), (3, 1), (2, 2), (1, 1), (2, 1);
 Row 2: (0, 1), (1, 10), (2, 2), (3, 1), (4, 1), (5, 1);
 Row 3: (1, 1), (0, 3), (1, 8), (2, 2), (4, 1), (6, 1).

In the code above, each pair of values represents a run, with the first value standing for the gray level and the second value giving the number of times the value has occurred in the run. The coding procedure is interrupted at the end of each row to permit synchronization in case of errors in the reception and decoding of run values. Direct coding of the $16 \times 3 = 48$ pixels in the three rows considered, at 3 b/pixel , leads to a total code length of 144 b . In the run-length code given above, if we were to use three bits per gray level and four bits per run-length value, we get a total code length of $18 \times 7 = 126 \text{ b}$. The savings are small in this case, due to the “busy” nature of the image, which represents an eye.

Run-length coding is best suited for the compression of bilevel images, where long runs may be expected of the two symbols 0 and 1. Images with fine details, intricate texture, and high-resolution quantization with large numbers of bits per pixel may not present long runs of the same gray level; run-length coding in such cases may lead to data expansion rather than compression. (This is the case with the image in Figure 11.2 past the third row.)

Run-length coding may be advantageously applied to bit planes of gray-level and color images. The use of Gray coding (see Table 11.1) improves the chances of long runs in the bit planes due to the feature that the Gray code changes in only one bit from one numerical value to the next.

Errors in run length could cause severe degradation of the reconstructed image due to the loss of pixel position. Synchronization at the end of each row can avoid the carrying over of such errors beyond the affected row.

Runs may also be defined over 2D areas. However, images with fine details do not present such uniform areas with large numbers of occurrence to lend much coding advantage.

11.3.3 Arithmetic coding

Arithmetic coding [960] is a family of codes that treat input symbols as magnitudes. Shannon [126] presented the basic idea of representing a string of symbols as the sum of their scaled probabilities. Most of the development towards practical arithmetic coding has been due to Langdon and Rissa-

nen [960, 961, 962]. The basic advantage of arithmetic coding over Huffman coding is that it does not suffer by the limitation that each symbol should have a unique code word that is at least one bit in length.

The mechanism of arithmetic coding is illustrated in Figure 11.6 [338, 963]. The symbols of the source string are represented by their individual probabilities p_l and cumulative probabilities (sum of the probabilities of all symbols up to, but not including, the current symbol) P_l . At any given stage in coding, the source string is represented by a code point C_k and an interval A_k . The code point C_k represents the cumulative probability of the current symbol on a scale of interval size A_k . A new symbol (being appended to the source string) is encoded by scaling the interval by the probability of the current symbol as

$$A_{k+1} = A_k p_l, \quad (11.9)$$


and defining the new code point as

$$C_{k+1} = C_k + A_k P_l. \quad (11.10)$$

Decoding is performed by the reverse of the procedure described above. The interval A_0 is initialized to unity, and the current code point is determined by the range into which the final code point C_{final} falls. The scaling of the interval and code point is performed as during encoding. The encoding procedure ensures that no future code point C_k exceeds the current value $C_k + A_k$. Thus, a carry over to a given bit position (in the binary representation of C_{final}) occurs at most once during encoding. This fact is made use of for incremental coding and for using finite-precision arithmetic. Finite precision is used by employing a technique known as *bit stuffing* [960], where, if a series of ones longer than the specified precision occurs in the binary-fraction representation of C_k , a zero is inserted; this ensures that further carries do not propagate into the series of ones. Witten et al. [963] provide an implementation of arithmetic coding using integer arithmetic.

Direct arithmetic coding of an image consists of an initial estimation of the probabilities of the gray values in the image, followed by row-wise arithmetic coding of pixels. Direct coding does not take into account the correlation between adjacent pixels. Arithmetic coding can be modified to make use of the correlation between pixels to some extent by using conditional probabilities of occurrence of gray levels.

In a version of arithmetic coding known as Q-coding [964], the individual bit planes of an image are coded using probabilities conditioned on the surrounding bits in the same plane as the context. A more efficient procedure is to perform decorrelation of the pixels of the image separately, and to use the basic arithmetic coder as a post-encoder on the decorrelated set of symbols (see, for example, Rabbani and Jones [965]). The performance of arithmetic coding as a post-encoder after the application of decorrelation methods is discussed in several sections to follow.

**FIGURE 11.6**

Arithmetic coding procedure. The range A_0 is initialized to unity. Each symbol is represented by its individual probability p_l and cumulative probability P_l . The string being encoded is represented by the code point C_k on the current range A_k . The range is scaled down by the probability p_l of the current symbol, and the process is repeated. One symbol is reserved for the end of the string [338, 963]. Figure courtesy of G.R. Kuduvali [338].

Example: The symbols used to represent the image in Figure 11.2 and their individual as well as cumulative probabilities are listed in Table 11.3. The intervals representing the symbols are also provided in the table. Let us consider the first three symbols $\{4, 6, 5\}$ in the fifth row of the image. The procedure to derive the arithmetic code for this string of symbols is shown in Figure 11.7.

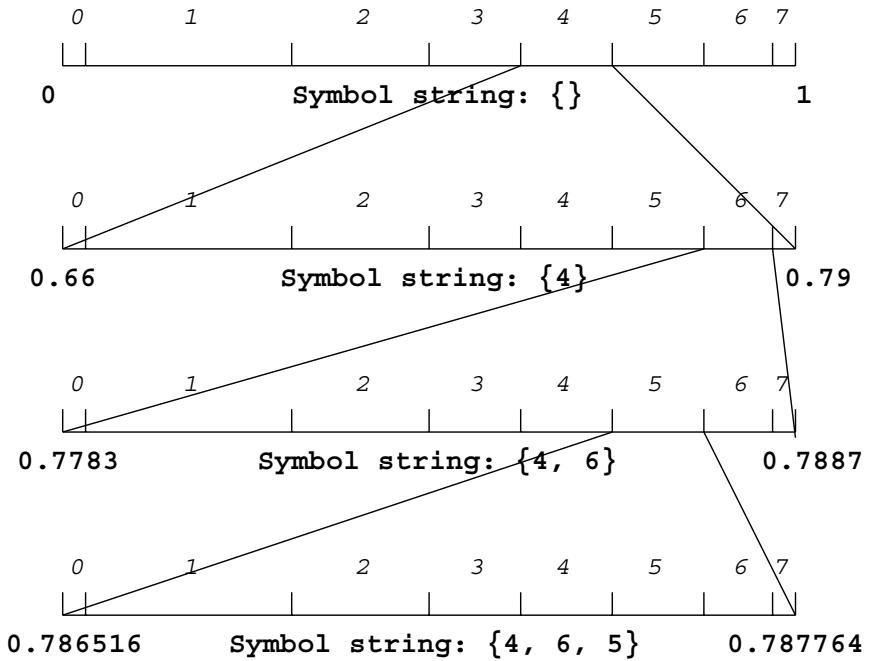
TABLE 11.3

The Symbols Used in the Image in Figure 11.2, Along with Their Individual Probabilities of Occurrence p_l , Cumulative Probabilities P_l , and Intervals Used in Arithmetic Coding.

Symbol l	Count	p_l	P_l	Interval
0	10	0.04	0.00	[0.00, 0.04)
1	77	0.30	0.04	[0.04, 0.34)
2	48	0.19	0.34	[0.34, 0.53)
3	33	0.13	0.53	[0.53, 0.66)
4	34	0.13	0.66	[0.66, 0.79)
5	32	0.12	0.79	[0.79, 0.91)
6	20	0.08	0.91	[0.91, 0.99)
7	2	0.01	0.99	[0.99, 1.00)

The initial code point is $C_0 = 0$, and the initial interval is $A_0 = 1$. When the first symbol “4” is encountered, the code point and interval are updated as $C_1 = C_0 + A_0 P_4 = 0 + 0.66 = 0.66$; $A_1 = A_0 p_4 = 1 \times 0.13 = 0.13$. For the next symbol “6”, we get $C_2 = C_1 + A_1 P_6 = 0.66 + 0.13 \times 0.91 = 0.7783$, and $A_2 = A_1 p_6 = 0.13 \times 0.08 = 0.0104$. With the third symbol “5” appended to the string, we have $C_3 = C_2 + A_2 P_5 = 0.7783 + 0.0104 \times 0.79 = 0.786516$, and $A_3 = A_2 p_5 = 0.0104 \times 0.12 = 0.001248$.

The code points have been given in decimal code to full precision, as required, in this example; the individual code probabilities have been rounded to two decimal places. In actual application, the code points need to be represented in binary code with finite precision. The average code-word length

**FIGURE 11.7**

The arithmetic coding procedure applied to the string {4, 6, 5} formed by the first three symbols in the fifth row of the image in Figure 11.2. See Table 11.3 for the related probabilities and intervals; see also Figure 11.6. All intervals are shown mapped to the same physical length, although their true values decrease from the interval at the top of the figure to that at the bottom. The numerals 0, 1, 2, ..., 7 in italics indicate the symbols (gray levels) in the image. The values in bold at the ends of each interval give the values of C_k and $(C_k + A_k)$ at the corresponding stage of coding.

per symbol is reduced by encoding long strings of symbols, such as an entire row of pixels in the given image.

11.3.4 Lempel–Ziv coding

Ziv and Lempel [966] proposed a universal coding scheme for encoding symbols from a discrete source when their probabilities of occurrence are not known a priori. The coding scheme consists of a rule for parsing strings of symbols from the source into substrings or words, and mapping the substrings into uniquely decipherable code words of fixed length. Thus, unlike the Huffman code where codes of fixed length are mapped into variable-length codes, the Lempel–Ziv code maps codes of variable length (corresponding to symbol strings of variable length) into codes of fixed length.

The Lempel–Ziv coding scheme is illustrated in Figure 11.8 [338]. The coding procedure starts with a buffer of length


$$n = L_s \alpha^\gamma L_s, \quad (11.11)$$

where L_s is the maximum length of the input symbol strings being parsed, α is the cardinality of the symbol source (in the case of image coding, the number of possible gray levels), and γ is chosen such that $0 < \gamma < 1$. The buffer is initially filled with $n - L_s$ zeros and the first L_s symbols from the source. The buffer is then parsed for the string whose length l_k is less than L_s , but is the maximum of all such strings from 0 to $n - L_s - 1$, and which has an identical string in the buffer starting at position $n - L_s$. The code to be mapped consists of the beginning position b_k of this string in the buffer from position 0 to $n - L_s - 1$, the length of the string l_k , and the last symbol s_k following the end of the string. The total length of the code for a straight binary representation is

$$l = \lceil \log_2(n - L_s) + \log_2(L_s) + \log_2(\alpha) \rceil, \quad (11.12)$$

where $\lceil x \rceil$ is the smallest integer $\geq x$. After coding the string, the buffer is advanced by l_k number of symbols. Ziv and Lempel [966] showed that, as the total length of the input symbols tends to ∞ , the average bit rate for coding the string approaches that of an optimal code with complete knowledge of the statistics of the source.

The Lempel–Ziv coding procedure may be viewed as a search through a fixed-size, variable-content dictionary for words that match the current string. A modification of this procedure, known as Lempel–Ziv–Welch (LZW) coding [967], consists of using a variable-sized dictionary with every new string encountered in the source string added to the dictionary. The dictionary is initialized to single-symbol strings, made up of the entire symbol set. This eliminates the need for including the symbol s_k in the code words. The LZW string table has the prefix property: for every string of symbols in the table, its prefix is also present in the table.

**FIGURE 11.8**

The Lempel–Ziv coding procedure. At each iteration, the buffer is scanned for strings of length $l_k \leq L_s$ for a match in the substring of length $(n - L_s)$ within the buffer. The matched string location b_k is encoded and transmitted. Figure courtesy of G.R. Kuduvali [338].

Kuduvali [338] implemented a slight variation of the LZW code, in which the first symbol of the current string is appended as the last symbol of the previously parsed string, and the new string is added to the string table. With this method, the decoded strings are generated in the same order as the encoded strings. The string table itself is addressed during the encoding procedure as a link-list. Each string contains the address of every other string of which it is a prefix. Such a link-list is not necessary during decoding, because the addresses of the strings are directly available to the decoder.

LZW coding may be applied directly for source coding, or applied to decorrelated data. The following sections provide examples of application of the LZW code.

Example: Let us consider again the image in Figure 11.2. The image has eight symbols in the range $[0, 7]$, each of which will be an item in the LZW table, shown in Table 11.4; the eight basic symbols may be represented by their own code. Index1 and Index2 represent two possibilities of coding. Consider the string $\{2, 2, 3\}$, which occurs five times in the image. In order to exploit this feature, we need to add the strings $\{2, 2\}$ and $\{2, 2, 3\}$ to the table; we could use the codes “8” for the former, and “9” or “83” for the latter (“83” represents the symbol “3” being appended to the string represented by the code “8”). The string $\{2, 2, 3, 5\}$ present at the beginning of the fourth row in the image may be represented as “95”. In this manner, long strings of symbols are encoded with short code words. The code index in the dictionary (table) is used to represent the symbol string. A predefined limit may be applied to the length of the dictionary.

TABLE 11.4
 Development of the
 Lempel–Ziv–Welch (LZW)
 Code Table for the Image
 in Figure 11.2.

String	Index1	Index2
0	0	0
1	1	1
2	2	2
3	3	3
4	4	4
5	5	5
6	6	6
7	7	7
⋮	⋮	⋮
22	8	22
223	9	83
2235	10	95
⋮	⋮	⋮

11.3.5 Contour coding

Given that a digital image includes only a finite number of gray levels, we could expect strings of the same values to occur in some form of 2D contours or patterns in the image [589]. The same expectation may be arrived at if we were to consider the gray level to represent height: an image may then be viewed as a relief map, with iso-intensity contours representing steps or plateaus (as in elevation maps of mountains). Information related to all such contours may then be used to encode the image. Although the idea is appealing in principle, fine quantization could result in low probabilities of occurrence of large contours with simple patterns.

Each iso-intensity contour would require the encoding of the coordinates of the starting point of the contour, the associated gray level, and the sequence of steps (movement) needed to trace the contour. A consistent rule is required for repeatable tracing of contours. The left-most-looking rule [589] for tracing a contour is as follows:

Select a pixel that is not already a member of a contour.

Look at the pixel to the left relative to the direction of entry to the current pixel on the contour.

If the pixel has the same gray level as the current pixel, move to the pixel and encode the type of movement.

If not, look at the pixel straight ahead.

If the pixel has the same gray level as the current pixel, move to the pixel and encode the type of movement.

If not, look at the pixel to the right.

If the pixel has the same gray level as the current pixel, move to the pixel and encode the type of movement.

If not, move to the pixel behind the current pixel.

Repeat; the procedure will trace a closed loop and return to the starting point.

Repeat the procedure until every pixel in the image belongs to a contour.

The movements allowed are only to the left, straight ahead, to the right, and back; the four possibilities may be encoded using the Freeman chain code illustrated in Figure 6.4.

Example: The image in Figure 11.2 is shown in Figure 11.9 along with the tracings of three iso-intensity contours. With reference to the contour with the value “1” at the top-left corner of the image, observe that several spatially connected pixels with the same value and lying within the contour have not been included in the contour: these pixels require additional contours. The contour-coding procedure needs to be applied repeatedly until every pixel in the image belongs to a contour. The encoding of short strings or isolated occurrences of pixel values could require several coding steps, and lead to increased code length.

The data required to represent the contour with the gray level “1” starting with the pixel in the first row and first column would be as follows:

Initial point: Coordinates [1, 1]. Gray level 1.

Freeman code: 0000000003030303022221212233201122122212.

Contour code requirement: four bits for each coordinate; three bits for the pixel value; two bits per Freeman code. Total $4 + 4 + 3 + 40 \times 2 = 91$ b.

Direct binary code requirement for the 15 pixels on the contour at 8 b/pixel = 120 b; at 3 b/pixel = 45 b.

The data required to represent the contour with the gray level “2” starting with the pixel in the fifth row and eighth column would be as follows:

Initial point: Coordinates [5, 8]. Gray level 2.

Freeman code: 0330221201.

Contour code requirement: Total 31 b. Direct binary code requirement for the eight pixels on the contour at 8 b/pixel = 64 b; at 3 b/pixel = 24 b.

The data required to represent the contour with the gray level “1” starting with the pixel in the ninth row and first column would be as follows:

Initial point: Coordinates [9, 1]. Gray level 1.

Freeman code: 03303233121111.

Contour code requirement: 39 b. Direct binary code requirement for the 13 pixels on the contour at 8 b/pixel = 104 b; at 3 b/pixel = 39 b.

It is evident that higher advantages may be gained if a number of long contours with simple patterns are present in the image.

11.4 Application: Source Coding of Digitized Mammograms

Kuduvalli et al. [174, 338] applied several coding techniques for the compression of digitized mammograms. In their work, film mammograms were scanned using an Eikonix-1412 digitizing camera, with a linear CCD array to provide a horizontal scan line of 4,096 pixels. A vertical array size of 4,096 pixels was achieved by stepping the array over 4,096 scan lines. A Gordon Instruments Plannar-1417 light box was used to illuminate the X-ray film being digitized. The gain and offset variations between the CCD elements were corrected for in the camera, and the data were transferred to the host computer over an IEEE-488 bus. Corrections were applied for the light-intensity variations in the Plannar-1417 light source used to illuminate the films, and the digitized image was stored in a Megascan FDP-2111 frame buffer with a capacity of $4,096 \times 4,096 \times 12$ b. Several $12 : 8$ b and $12 : 10$ b transformations were developed to test the dynamic range, light intensity, and focus settings.

The effective dynamic range of an imaging system depends upon the scaling factors used for correcting light-intensity variations and the SNR of the imaging system. Kuduvalli et al. analyzed the intensity profiles of the Plannar-1417 light box, and observed that a scaling factor of about 1.6 was required

to correct the values of the pixels at the edges of the image. Furthermore, the local standard deviation of the intensity levels measured by the camera with respect to a moving-average window of 10 counts was estimated to be about 5.0 counts. The effective noise level at the edge of the imaging field was estimated at 8 counts. For these reasons, it was concluded that two of the least-significant bits in the 12 b data would only contribute to noise and affect the performance of data compression algorithms [968]. In addition, by scanning a standard, calibrated, gray-scale step pattern, the effective dynamic range of the digitizer was observed to be about 2.5 OD . In consideration of all of the above factors, it was determined that truncating the 12 b pixel values to 10 b values would be adequate for representing the digitized image. This procedure reduced the effective noise level by a factor of four, to about 2 counts.

Kuduvalli et al. also estimated the MTF of the digitization system from measurements of the ESF (see Sections 2.9 and 2.12), with a view to demonstrate the resolving capability of the system to capture sub-millimeter details on X-ray films, such as microcalcifications on mammograms. The normalized value of the MTF at one-half of the sampling frequency was estimated to be 0.1, which is considered to be adequate for resolving objects and details at the same frequency (which is the highest frequency component retained in the digitized image) [969].

The average number of bits per pixel obtained for ten X-ray images using the Huffman, arithmetic, and LZW coding techniques are listed in Table 11.5; the zeroth-order entropy values of the images are also listed. The high values of the zeroth-order entropy indicate limits on the performance of the Huffman coding technique. The arithmetic coding method has given bit rates comparable with those provided by the Huffman code, but at considerably higher level of complexity. Figure 11.10 shows plots of the average bit rate as a function of the buffer length in LZW coding, for four of the ten images listed in Table 11.5. The maximum length of the symbol strings scanned was fixed at $L_s = 256$. The LZW code provided the best compression rates among the three methods considered, to the extent of about 50% of the initial number of bits per pixel in the images. The average bit rate provided by LZW coding is well below the zeroth-order entropy values of the images, indicating that efficient encoding of strings of pixel values can exploit the redundancy and correlation present in the data without performing explicit decorrelation.

11.5 The Need for Decorrelation

The results of direct source encoding, discussed in Section 11.4, indicate the limitations of direct encoding of the symbols generated by an image source.

TABLE 11.5

Average Number of Bits Per Pixel with Direct Source Coding for Data Compression of Digitized Mammograms and Chest X-ray Images [174, 338].

Image	Type	Size (pixels)	Entropy	Huffman	Arith.	LZW
1	Mammo.	4,096 × 1,990	7.26	8.20	8.09	5.34
2	Mammo.	4,096 × 1,800	7.61	8.59	8.50	5.76
3	Mammo.	3,596 × 1,632	6.68	6.96	6.88	4.98
4	Mammo.	3,580 × 1,696	7.21	7.80	7.71	4.68
5	Chest	3,536 × 3,184	8.92	9.62	9.43	6.11
6	Chest	3,904 × 3,648	9.43	9.83	9.81	6.27
7	Chest	3,264 × 3,616	6.26	7.20	7.12	4.61
8	Chest	4,096 × 4,096	8.65	9.39	9.35	5.83
9	Mammo.	4,096 × 2,304	8.83	9.71	9.57	6.13
10	Chest	4,096 × 3,800	8.57	9.42	9.33	5.99
Average			7.94	8.67	8.58	5.57

The entropy listed is the zeroth-order entropy. Pixel values in the original images were quantized at 10 *b/pxel*. See also Table 11.7. *Note:* Arith. = Arithmetic coding; LZW: Lempel–Ziv–Welch coding; Mammo. = Mammogram. Reproduced with permission from G.R. Kuduvali and R.M. Rangayyan, “Performance analysis of reversible image compression techniques for high-resolution digital teleradiology”, *IEEE Transactions on Medical Imaging*, 11(3): 430 – 445, 1992. © IEEE.




FIGURE 11.10

Average bit rate as a function of the buffer length, using Lempel-Ziv-Welch coding, for four of the ten images (number 1, 3, 4, and 6) listed in Table 11.5. The abscissa indicates the value of B , with the buffer length given by $n = 2^B$. The maximum length of the symbol strings scanned was fixed at $L_s = 256$. Reproduced with permission from G.R. Kuduvali and R.M. Rangayyan, "Performance analysis of reversible image compression techniques for high-resolution digital teleradiology", *IEEE Transactions on Medical Imaging*, 11(3): 430 - 445, 1992. © IEEE.

Although some of the methods described, such as the Lempel–Ziv and run-length coding methods, have the potential to exploit the redundancy present in images, their efficiency in this task is limited.

The term “decorrelation” indicates a procedure that can remove or reduce the redundancy or correlation present between the elements of a data stream, such as the pixels in an image. The most commonly used decorrelation techniques are the following:

- differentiation, which can remove the commonality present between adjacent elements;
- transformation to another domain where the energy of the image is confined to a narrow range, such as the Fourier, Karhunen–Loève, discrete cosine, or Walsh–Hadamard (orthogonal) transform domains;
- model-based prediction, where the error of prediction would have reduced information content; and
- interpolation, where a subsampled image is transmitted, the pixels in between the preceding data are obtained by interpolation, and the error of interpolation, which has reduced information content, is transmitted.

Observe that the decorrelated data (transform coefficients, prediction error, etc.) need to be encoded and transmitted; the techniques described in Section 11.3 for direct source encoding may also be applied to decorrelated data. In addition, further information regarding initial values and the procedures for the management of the transform coefficients or the model parameters will also have to be sent to facilitate complete reconstruction of the original image.

The advantages of decorrelating image data by differentiation are demonstrated by the following simple example. The use of transforms for image data compression is discussed in Section 11.6. Interpolative coding is briefly described in Section 11.7. Methods for prediction-based data compression are described in Section 11.8. Techniques based upon different scanning strategies to improve the performance of decorrelation by differentiation are discussed in Sections 11.9 and 11.11. Strategies for combining several decorrelation steps are discussed in Section 11.12.

Example: Consider the image in Figure 11.11 (a); the histogram of the image is shown in part (b) of the figure. The image has a good spread of gray levels over its spatial extent, and the histogram, while not uniform, does exhibit a good spread over the dynamic range of the image. The zeroth-order entropy, at $7.41 b$, is close to the maximum possible value of $8 b$ for the image with 256 gray levels. These characteristics suggest limited potential for direct encoding methods.

The image in Figure 11.11 (a) was subjected to a simple first-order partial differentiation procedure, given by

$$f'(m, n) = f(m, n) - f(m - 1, n). \quad (11.13)$$

The result, shown in Figure 11.12 (a), has an extremely limited range of details; the histogram of the image, shown in part (b) of the figure, indicates that, although the image has more gray levels than the original image, most of the gray levels occur with negligible probability. The concentrated histogram leads to a lower value of entropy, at $5.52 b$. Observe that the histogram of the difference image is close to a Laplacian PDF (see Section 3.1.2 and Figure 3.9). The simple operation of differentiation has reduced the entropy of the image by about 25%. The reduced entropy suggests that the coding requirement may be reduced significantly. Observe that the additional information required to recover the original image from its derivative as above is just the first row of pixels in the original image. Data compression techniques based upon differentiation are referred to as differential pulse code modulation (DPCM) techniques. DPCM techniques vary in terms of the reference value used for subtraction (in the differentiation process). The reference value may be derived as a combination of a few neighboring pixels, in which case the method approaches linear prediction in concept.


11.6 Transform Coding

The main premise of transform-domain coding of images is that, when orthogonal transforms are used, the related coefficients represent elements that are mutually uncorrelated. (See Section 3.5.2 for the basics of orthogonal transforms.) Furthermore, because most natural images have limitations on the rate of change of their elemental values (that is, they are generally smooth), their energy is confined to a narrow low-frequency range in the transform domain. These properties lead to two characteristics of orthogonal transforms that are of relevance and importance in data compression:


- orthogonal transforms perform decorrelation, and
- orthogonal transforms compress the energy of the given image into a narrow region.

The second property listed above is commonly referred to as “energy compaction”.

Example: The log-magnitude Fourier spectrum of the image in Figure 11.11 (a) is shown in Figure 11.13 (a). It is evident that most of the energy of the image is concentrated in a small number of DFT coefficients around the origin in the 2D Fourier plane (at the center of the spectrum displayed). In order to demonstrate the energy-compacting nature of the DFT, the (cumulative) percentage of the total energy of the image present at the $(0, 0)$ coordinate of the DFT, and contained within concentric square regions of



(a)




(b)

FIGURE 11.11

(a) A test image of size 225×250 pixels with 256 gray levels. (b) Gray-level histogram of the test image. Dynamic range [18, 255]. Zeroth-order entropy 7.41 *b*.



(a)



(b)

FIGURE 11.12

(a) Result of differentiation of the test image in Figure 11.11 (a). (b) Gray-level histogram of the image in (a). Dynamic range $[-148, 180]$. Zeroth-order entropy 5.52 *b*.

half-width 1, 2, 3, ..., 8 DFT coefficients were computed. The result is plotted in Figure 11.13 (b), which shows that 74% of the energy of the image is present in the DC component, and 90% of the energy is contained within the central 121 DFT components around the DC point; only 7.2% of the energy lies in the high-frequency region beyond the central 17×17 region of the 256×256 DFT array. (Regardless of the small fraction of the total energy present at higher frequencies, it should be observed that high-frequency components bear important information related to the edges and sharpness of the image; see Sections 2.11.1 and 3.4.1).

The DFT is the most commonly used orthogonal transform in the analysis of systems, signals, and images. However, due to the complex nature of the basis functions, the DFT has high computational requirements. In spite of its symmetry, the DFT could lead to increased direct coding requirements due to the need for large numbers of bits for quantization of the transform coefficients. Regardless, the discrete nature of most images at the outset could be used to advantage in lossless recovery from transform coefficients that have been quantized to low levels of accuracy; see Section 11.6.3.

We have already studied the DFT (Sections 2.11 and 3.5.2) and the WHT (Section 3.5.2). The WHT has a major computational advantage due to the fact that its basis functions are composed of only $+1$ and -1 , and has been advantageously applied in data compression. In the following sections, we shall study two other transforms that are popular and relevant to data compression: the discrete cosine transform (DCT) and the Karhunen–Loève transform (KLT).

11.6.1 The discrete cosine transform

The DCT is a modification of the DFT that overcomes the effects of discontinuities at the edges of the latter [970], and is defined as

$$F(k, l) = \frac{a(k, l)}{N} \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} f(m, n) \cos \left[\frac{\pi m}{N} (2k+1) \right] \cos \left[\frac{\pi n}{N} (2l+1) \right], \quad (11.14)$$


for $k = 0, 1, 2, \dots, N-1$, and $l = 0, 1, 2, \dots, N-1$, where

$$a(k, l) = \begin{cases} 1 & \text{if } (k, l) = (0, 0) \\ \frac{1}{2} & \text{otherwise.} \end{cases} \quad (11.15)$$


The inverse transformation is given by

$$F(k, l) = \frac{1}{N} \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} a(k, l) F(k, l) \cos \left[\frac{\pi m}{N} (2k+1) \right] \cos \left[\frac{\pi n}{N} (2l+1) \right], \quad (11.16)$$

for $m = 0, 1, 2, \dots, N-1$, and $n = 0, 1, 2, \dots, N-1$. The basis vectors of the DCT closely approximate the eigenvectors of a Toeplitz matrix (see



(a)



(b)

FIGURE 11.13

(a) Log-magnitude Fourier spectrum of the image in Figure 11.11 (a), computed over a 256×256 DFT array. (b) Distribution of the total image energy. The 10 values represent the (cumulative) percentage of the energy of the image present at the $(0,0)$ or DC position; contained within square boxes of half-width 1, 2, 3, \dots , 8 pixels centered in the DFT array; and in the entire 256×256 DFT array. The numbers of DFT coefficients corresponding to the 10 regions are 1, 9, 25, 49, 81, 121, 169, 225, 289, and 65, 536.

Section 3.5.3) whose elements can be expressed by increasing powers of a constant $\gamma < 1$. The ACF matrices of most natural images can be closely modeled by such a matrix [971]. An $N \times N$ DCT may be computed from the results of a $2N \times 2N$ DFT; thus, FFT algorithms may be used for efficient computation of the DCT.

11.6.2 The Karhunen–Loève transform

The KLT, also known as the principal-component transform, the Hotelling transform, or the eigenvector transform [9, 589], is a data-dependent transform based upon the statistical properties of the given image. The image is treated as a vector \mathbf{f} that is a realization of an image-generating stochastic process. If the image is of size $M \times N$, the vector is of size $P = MN$ (see Section 3.5).

The image \mathbf{f} may be represented without error by a deterministic linear transformation of the form

$$\mathbf{f} = \mathbf{A} \mathbf{g} = \sum_{m=1}^P g_m \mathbf{A}_m, \quad (11.17)$$

$$\mathbf{A} = [\mathbf{A}_1; \mathbf{A}_2; \cdots; \mathbf{A}_P], \quad (11.18)$$

where $|\mathbf{A}| \neq 0$, and \mathbf{A}_m are row vectors that make up the $P \times P$ matrix \mathbf{A} . The matrix \mathbf{A} needs to be formulated such that the vector \mathbf{g} leads to an efficient representation of the original image \mathbf{f} .

The matrix \mathbf{A} may be considered to be made up of P linearly independent row vectors that span the P -dimensional space containing \mathbf{f} . Let \mathbf{A} be orthonormal, that is,

$$\mathbf{A}_m^T \mathbf{A}_n = \begin{cases} 1 & m = n \\ 0 & m \neq n \end{cases}. \quad (11.19)$$

It follows that

$$\mathbf{A}^T \mathbf{A} = \mathbf{I} \text{ or } \mathbf{A}^{-1} = \mathbf{A}^T. \quad (11.20)$$

Then, the row vectors of \mathbf{A} may be considered to form the set of orthonormal basis vectors of a linear transformation. This formulation leads also to the inverse relationship

$$\mathbf{g} = \mathbf{A}^T \mathbf{f} = \sum_{m=1}^P \mathbf{A}_m^T f_m. \quad (11.21)$$

In the procedure described above, each component of \mathbf{g} contributes to the representation of \mathbf{f} . Given the formulation of \mathbf{A} as a reversible linear transformation, \mathbf{g} provides a complete (lossless) representation of \mathbf{f} if all of its $P = MN$ elements are made available.

Suppose that, in the interest of efficient representation of images via the extraction of the most significant information contained, we wish to use only

$Q < P$ components of \mathbf{g} . The omitted components of \mathbf{g} may be replaced with other values b_m , $m = Q + 1, \dots, P$. Then, we have an approximate representation of \mathbf{f} , given as

$$\tilde{\mathbf{f}} = \sum_{m=1}^Q g_m \mathbf{A}_m + \sum_{m=Q+1}^P b_m \mathbf{A}_m. \quad (11.22)$$

The error in the approximate representation as above is

$$\boldsymbol{\varepsilon} = \mathbf{f} - \tilde{\mathbf{f}} = \sum_{m=Q+1}^P (g_m - b_m) \mathbf{A}_m. \quad (11.23)$$

The MSE is given by

$$\begin{aligned} \overline{\boldsymbol{\varepsilon}^2} &= E[\boldsymbol{\varepsilon}^T \boldsymbol{\varepsilon}] \\ &= E[\sum_{m=Q+1}^P \sum_{n=Q+1}^P (g_m - b_m)(g_n - b_n) \mathbf{A}_m^T \mathbf{A}_n] \\ &= \sum_{m=Q+1}^P E[(g_m - b_m)^2]. \end{aligned} \quad (11.24)$$

The last step above follows from the orthonormality of \mathbf{A} .

Taking the derivative of the MSE with respect to b_m and setting the result to zero, we get

$$\frac{\partial \overline{\boldsymbol{\varepsilon}^2}}{\partial b_m} = -2 E[(g_m - b_m)] = 0. \quad (11.25)$$

The optimal (MMSE) choice for b_m is, therefore, given by

$$b_m = E[g_m] = \bar{g}_m = \mathbf{A}_m^T E[\mathbf{f}], \quad m = Q + 1, \dots, P; \quad (11.26)$$

that is, the omitted components are replaced by their means. The MMSE is given by

$$\begin{aligned} \overline{\boldsymbol{\varepsilon}^2}_{\min} &= \sum_{m=Q+1}^P E[(g_m - \bar{g}_m)^2] \\ &= \sum_{m=Q+1}^P E[\mathbf{A}_m^T (\mathbf{f} - \bar{\mathbf{f}}) (\mathbf{f} - \bar{\mathbf{f}})^T \mathbf{A}_m] \\ &= \sum_{m=Q+1}^P \mathbf{A}_m^T \boldsymbol{\sigma}_f \mathbf{A}_m, \end{aligned} \quad (11.27)$$

where $\boldsymbol{\sigma}_f$ is the covariance matrix of \mathbf{f} .

Now, if the basis vectors \mathbf{A}_m are selected as the eigenvectors of $\boldsymbol{\sigma}_f$, that is,

$$\boldsymbol{\sigma}_f \mathbf{A}_m = \lambda_m \mathbf{A}_m, \quad (11.28)$$

and

$$\lambda_m = \mathbf{A}_m^T \boldsymbol{\sigma}_f \mathbf{A}_m \quad (11.29)$$

because $\mathbf{A}_m^T \mathbf{A}_m = \mathbf{I}$, where λ_m are the corresponding eigenvalues, then, we have

$$\overline{\epsilon^2}_{\min} = \sum_{m=Q+1}^P \lambda_m. \quad (11.30)$$

Therefore, the MSE may be minimized by ordering the eigenvectors (the rows of \mathbf{A}) such that the corresponding eigenvalues are arranged in decreasing order, that is, $\lambda_1 > \lambda_2 > \dots > \lambda_P$. Then, if a component g_m of \mathbf{g} is replaced by $b_m = \bar{g}_m$, the MSE increases by λ_m . By replacing the components of \mathbf{g} corresponding to the eigenvalues at the lower end of the list, the MSE is kept at its lowest-possible value for a chosen number of components Q .

From the above formulation and properties, it follows that the components of \mathbf{g} are mutually uncorrelated:

$$\sigma_g = \mathbf{A}^T \sigma_f \mathbf{A} = \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_P \end{bmatrix} = \mathbf{\Lambda}, \quad (11.31)$$

where $\mathbf{\Lambda}$ is a diagonal matrix with the eigenvalues λ_m placed along its diagonal. Because the eigenvalues λ_m are equal to the variances of g_m , a selection of the larger eigenvalues implies the selection of the transform components with the higher variance or information content across the ensemble of the images considered.

The KLT has major applications in principal-component analysis (PCA), image coding, data compression, and feature extraction for pattern classification. Difficulties could exist in the computation of the eigenvectors and eigenvalues of the large covariance matrices of even reasonably sized images. It should be noted that a KLT is optimal only for the images represented by the statistical parameters used to derive the transformation. New transformations will need to be derived if changes occur in the statistics of the image-generating process being considered, or if images of different statistical characteristics need to be analyzed.

Because the KLT is a data-dependent transform, the transformation vectors (the matrix \mathbf{A}) need to be transmitted; however, if a large number of images generated by the same underlying process are to be transmitted, the same optimal transform is applicable, and the transformation vectors need to be transmitted only once. Note that the error between the original image and the image reconstituted from the KLT components needs to be transmitted in order to facilitate lossless recovery of the image.

See Section 8.9.5 for a discussion on the application of the KLT for the selection of the principal components in multiscale directional filtering.

11.6.3 Encoding of transform coefficients

Regardless of the transform used (such as the DFT, DCT, or KLT), the transform coefficients form a set of continuous random variables, and have to be quantized for encoding. This introduces quantization errors in the transform coefficients that are transmitted, and hence, errors arise in the image reconstructed from the transform-coded image. In the following paragraphs, a relationship is derived between the quantization error in the transform domain and the error in the reconstructed image. Kuduvalli and Rangayyan [174, 338, 972] used such a relationship to develop a method for error-free transform coding of images.

Consider the general 2D linear transformation, with the forward and inverse transform kernels consisting of orthogonal basis functions $a(m, n, k, l)$ and $b(m, n, k, l)$, respectively, such that the forward and inverse transforms are given by

$$F(k, l) = \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} f(m, n) a(m, n, k, l), \quad (11.32)$$

$k, l = 0, 1, \dots, N - 1$, and

$$f(m, n) = \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} F(k, l) b(m, n, k, l), \quad (11.33)$$

$m, n = 0, 1, \dots, N - 1$. (See Section 3.5.2.) Now, let the transform coefficient $F(k, l)$ be quantized to $\tilde{F}(k, l)$, with a quantization error $q_F(k, l)$ such that

$$F(k, l) = \tilde{F}(k, l) + q_F(k, l). \quad (11.34)$$

The reconstructed image from the quantized transform coefficients is given by

$$\tilde{f}(m, n) = \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} \tilde{F}(k, l) b(m, n, k, l). \quad (11.35)$$

The error in the reconstructed image is

$$\begin{aligned} q_f(m, n) &= f(m, n) - \tilde{f}(m, n) \\ &= \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} [F(k, l) - \tilde{F}(k, l)] b(m, n, k, l) \\ &= \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} q_F(k, l) b(m, n, k, l). \end{aligned} \quad (11.36)$$

The sum of the squared errors in the reconstructed image is

$$Q_f = \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} |q_f(m, n)|^2$$

$$\begin{aligned}
&= \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} \left\{ \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} q_F(k, l) b(m, n, k, l) \right\} \\
&\quad \times \left\{ \sum_{k'=0}^{N-1} \sum_{l'=0}^{N-1} q_F(k', l') b(m, n, k', l') \right\}^* \\
&= \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} \sum_{k'=0}^{N-1} \sum_{l'=0}^{N-1} q_F(k, l) q_F^*(k', l') \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} b(m, n, k, l) b^*(m, n, k', l') \\
&= \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} |q_F(k, l)|^2 = Q_F, \tag{11.37}
\end{aligned}$$

where the last line follows from the orthogonality of the basis functions $b(m, n, k, l)$, and Q_F is the sum of the squared quantization errors in the transform domain. This result is related to Parseval's theorem in 2D; see Equation 2.55. Applying the expectation operator to the first and the last expressions above, we get

$$\sum_{m=0}^{N-1} \sum_{n=0}^{N-1} E[|q_f(m, n)|^2] = \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} E[|q_F(k, l)|^2], \tag{11.38}$$

or

$$\sum_{m=0}^{N-1} \sum_{n=0}^{N-1} \overline{q_f^2(m, n)} = \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} \overline{q_F^2(k, l)} = \overline{Q^2}, \tag{11.39}$$

where the $\overline{\quad}$ symbol indicates the expected (average) values of the corresponding variables, and $\overline{Q^2}$ is the expected total squared error of quantization (in either the image domain or the transform domain).

It is possible to derive a condition for the minimum average number of bits required for encoding the transform coefficients for a given total distortion in the image domain. Let us assume that the transform coefficients are normally distributed. If the variance of the transform coefficient $F(k, l)$ is $\sigma_F^2(k, l)$, the average number of bits required to encode the coefficient $F(k, l)$ with the MSE $\overline{q_F^2(k, l)}$ is given by its rate-distortion function [973]

$$R(k, l) = \frac{1}{2} \log_2 \left[\frac{\sigma_F^2(k, l)}{\overline{q_F^2(k, l)}} \right]. \tag{11.40}$$

The overall average number of bits required to encode the transform coefficients with a total squared error $\overline{Q^2}$ is

$$\begin{aligned}
R_{\text{av}} &= \frac{1}{N^2} \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} R(k, l) \\
&= \frac{1}{2N^2} \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} \log_2 \left[\frac{\sigma_F^2(k, l)}{\overline{q_F^2(k, l)}} \right]. \tag{11.41}
\end{aligned}$$

We now need to minimize R_{av} subject to the condition given by Equation 11.39. Using the method of Lagrange's multiplier, the minimum occurs when

$$\frac{\partial}{\partial q_F^2(k,l)} \left\{ \frac{1}{2N^2} \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} \log_2 \left[\frac{\sigma_F^2(k,l)}{q_F^2(k,l)} \right] - \lambda \left[Q^2 - \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} \overline{q_F^2(k,l)} \right] \right\} = 0, \quad (11.42)$$

$k, l = 0, 1, \dots, N-1$, where λ is the Lagrange multiplier. It follows that

$$-\frac{1}{2N^2 \ln(2) \overline{q_F^2(k,l)}} + \lambda = 0, \quad (11.43)$$

or

$$\overline{q_F^2(k,l)} = \frac{1}{2N^2 \lambda \ln(2)} = \overline{q^2}, \quad (11.44)$$


$k, l = 0, 1, \dots, N-1$, where $\overline{q^2}$ is the average MSE, which is a constant for all of the transform coefficients. Thus, the average number of bits required to encode the transform coefficients R_{av} is minimum when the total squared error is equally distributed among all of the transform coefficients.

Maximum-error-limited encoding of transform coefficients: Kuduvali and Rangayyan [174, 338, 972] derived the following condition for encoding transform coefficients subject to a maximum error limit. Consider a uniform quantizer with a quantization step size S for encoding the transform coefficients, such that the maximum quantization error is limited to $\pm \frac{S}{2}$. It may be assumed that the quantization error is uniformly distributed over the set of transform coefficients in the range $[-\frac{S}{2}, +\frac{S}{2}]$ (see Figure 11.14). Then, the average squared error in the transform domain is

$$\overline{q^2} = \frac{S^2}{12}. \quad (11.45)$$

From the result in Equation 11.39, it is seen that the errors in the reconstructed image will also have a variance equal to $\overline{q^2}$. We now wish to estimate the fraction of the total number of pixels in the reconstructed image that are in error by more than $\pm S$. This is given by the area under the tail of the PDF of the reconstruction error, shown in Figure 11.14. The worst case occurs when the entropy of the reconstruction errors is at its maximum, under the constraint that the variance of the reconstruction errors is bounded by $\overline{q^2}$; this occurs when the error is normally distributed [126]. Therefore, the upper bound on the estimated fraction of the pixels in error by more than $\pm S$ is

$$E(S) = 2 \int_S^\infty \frac{1}{\sqrt{2\pi q^2}} \exp\left(-\frac{x^2}{2q^2}\right) dx = 2 \operatorname{erfc}(\sqrt{12}) = 5.46 \times 10^{-4}, \quad (11.46)$$

**FIGURE 11.14**

Schematic PDFs of the transform-coefficient quantization error (uniform PDF, solid line) and the image reconstruction error (Gaussian PDF, dashed line). The figure represents the case with the quantization step size $S = 1$.

where $\operatorname{erfc}(x)$ is the error function integral, defined as [974]

$$\operatorname{erfc}(x) = \int_x^\infty \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right) dx. \quad (11.47)$$

Thus, only a negligible number of pixels in the reconstructed image will be in error by more than the quantization step size S . Conversely, if the maximum error is desired to be $\leq S$, a quantization step size of S could be used to encode the transform coefficients, with only a negligibly small number of the reconstructed pixels exceeding the error limit. The pixels in error by more than the specified maximum could be encoded separately with a small overhead. When the maximum allowed error is ≤ 0.5 , error-free reconstruction of the image is possible by simply rounding off the reconstructed values to integers.

Variable-length encoding and bit allocation: The lower bound on the average number of bits required for encoding normally distributed transform coefficients $F(k, l)$ with the MSE $\frac{\sigma_F^2}{q_F^2}(k, l)$ is given by Equation 11.40. Goodness-of-fit studies of PDFs of transform coefficients [975] have shown that transform coefficients tend to follow the Laplacian PDF. The PDF of a transform coefficient $F(k, l)$ may be modeled as

$$p(F(k, l)) = \frac{\gamma(k, l)}{2} \exp(-\gamma(k, l) |F(k, l)|), \quad (11.48)$$

where $\gamma(k, l) = \sqrt{2}/\sigma_F(k, l)$ is the constant parameter of the Laplacian PDF. A shift encoder could be used to encode the transform coefficients such that the maximum quantization error is $\leq S$. The shift-encoding procedure is shown in Figure 11.15. In a shift encoder, $2\nu(k, l)$ levels are nominally allocated to encode a transform coefficient $F(k, l)$, covering the range $[-\{\nu(k, l) - 1\} S, \{\nu(k, l) - 1\} S]$ with the codes $0, 1, 2, \dots, 2\nu(k, l) - 2$. The code $2\nu(k, l) - 1$ indicates that the coefficient is out of the range $[-\{\nu(k, l) - 1\} S, \{\nu(k, l) - 1\} S]$. For the out-of-range coefficients, an additional $2\nu(k, l)$ levels are allocated to cover the ranges $[-\{2\nu(k, l) - 1\} S, -\nu(k, l) S]$, and $[\nu(k, l) S, \{2\nu(k, l) - 1\} S]$. The process is repeated with the allocation of additional levels until the actual value of the transform coefficient to be encoded is reached. If the code value is represented by a simple binary code at each level, the average number of bits required to encode the transform coefficient $F(k, l)$ is given by [338]

$$R(k, l) = \frac{1 + \log_2 \nu(k, l)}{1 - \exp[-\gamma(k, l) \nu(k, l) S]}, \quad (11.49)$$

and the nominal number of bits allocated to encode the transform coefficient is

$$b(k, l) = \lceil \log_2 [\nu(k, l)] \rceil. \quad (11.50)$$

It is now required to find the $b(k, l)$ that minimizes the average number of bits $R(k, l)$ required to encode $F(k, l)$. This can be done by using a nonlinear

optimization technique such as the Newton-Raphson method. However, because only integral values of $b(k, l)$ need to be searched, it is computationally less expensive to search the space of $b(k, l) \in [1, 31]$ for the corresponding minimum value of $R(k, l)$, which is the range of the nominal number of bits allocated to encode the transform coefficient $F(k, l)$.




FIGURE 11.15

Schematic representation of shift coding of transform coefficients with a Laplacian PDF. With reference to the discussion in the text, the figure represents $\nu(k, l) S = 0.5$ and $\gamma(k, l) = 3.0$.

This allocation requires an estimate of the variance $\sigma_F^2(k, l)$ of the transform coefficients, or a model of the energy-compacting property of the transform used. Most of the linear orthogonal transforms used in practice result in the concentration of the variance in the lower-order transform coefficients. The variance of the transform coefficients may be modeled as

$$\sigma_F^2(k, l) = F(0, 0) \exp [-(\alpha k^2 + \beta l^2)], \quad (11.51)$$

where $F(0, 0)$ is the lowest-order transform coefficient, and α and β are the constants of the model. For most transforms (except the KLT), $F(0, 0)$ is the

average value or a scaled version of the average value of the image pixels. The parameters α and β may be estimated using a least-squares fit to the first few transform coefficients of the image.

In an alternative coding procedure, a fixed total number of bits may be allocated for encoding the transform coefficients, and the difference image encoded by a lossless encoding method. In such a procedure, no attempt is made to allocate additional bits for transform coefficients that result in errors that fall out of the quantization range. The total number of bits allocated to encode the transform coefficients is varied until the total average bit rate is at its minimum. Using such a procedure, Cox et al. [976] found an optimal combination of bit rates between the error image and the transform coefficients. Wang and Goldberg [977] used a method of requantization of the quantization errors in addition to encoding the error image; they too observed a minimum total average bit rate after a number of iterations of requantization. Experiments conducted by Kuduvalli and Rangayyan [174, 338, 972] showed that the lowest bit rates obtained by such methods for reversible compression can also be obtained by allocating additional bits to quantize the out-of-range transform coefficients as described earlier. This is due to the fact that a large quantization error in the transform domain, while needing only a few additional bits for encoding, will get redistributed over a large number of pixels in the image domain, thereby increasing the entropy of the error image.

The large sizes of image arrays used for high-resolution representation of medical images preclude the use of full-frame transforms. Partitioning an image into blocks not only leads to computational advantages, but also permits adaptation to the changing statistics of the image. In the coding procedure used by Kuduvalli and Rangayyan [174, 338, 972], the images listed in Table 11.5 were partitioned into blocks of size 256×256 pixels. The model parameters α and β in Equation 11.51 were computed for each block by using a least-squares fit to the corresponding set of transform coefficients. The parameters were stored or transmitted along with the encoded transform coefficients in order to allow the decoder to reconstruct the model and the bit-allocation table. Blocks at the boundaries of the images that were not squares were encoded with 2D DPCM and the Huffman code.

Figure 11.16 shows the average bit rate for one of the images listed in Table 11.5, as a function of the maximum allowed error, using four transforms. The KLT (with the ACF estimated from the image) and the the DCT show the best performance among the four transforms. The performance of the KLT is only slightly superior to, and in some cases slightly worse than that of the DCT; this is to be expected because of the general nonstationarity of medical images, and due to the problems associated with the estimation of the ACF matrix from a finite image.

Figure 11.17 shows the average bit rate, obtained by using the DCT, as a function of the maximum allowed error, for four of the ten images listed in Table 11.5. When the maximum allowed error is ≤ 0.5 , error-free reconstruction of the original image is possible; otherwise, the compression is irreversible.

The average bit rate for error-free reconstruction is seen to be in the range of 5–6 *b/pixel* for the images considered (down from 10 *b/pixel* in their original format).




FIGURE 11.16

Average bit rate as a function of the maximum allowed error, using four transforms (KLT, DCT, DFT, and WHT) with the first image listed in Table 11.5. A block size of 256×256 pixels was used for each transform. The compression is lossless if the maximum allowed error is ≤ 0.5 ; otherwise, it is irreversible or lossy. See also Table 11.7. Reproduced with permission from G.R. Kuduvali and R.M. Rangayyan, "Performance analysis of reversible image compression techniques for high-resolution digital teleradiology", *IEEE Transactions on Medical Imaging*, 11(3): 430–445, 1992. © IEEE.




FIGURE 11.17

Average bit rate as a function of the maximum allowed error, using the DCT, for the first four images listed in Table 11.5. A block size of 256×256 pixels was used. The compression is lossless if the maximum allowed error is ≤ 0.5 ; otherwise, it is irreversible or lossy. See also Table 11.7. Reproduced with permission from G.R. Kuduvali and R.M. Rangayyan, "Performance analysis of reversible image compression techniques for high-resolution digital teleradiology", *IEEE Transactions on Medical Imaging*, 11(3): 430 - 445, 1992. © IEEE.

11.7 Interpolative Coding

Interpolative coding consists of encoding a subsampled image using a reversible compression technique, deriving the values of the remaining pixels via interpolation with respect to their neighboring pixels that have already been processed, and then encoding the difference between the actual pixels and the interpolated pixels in successive stages using discrete symbol coding techniques. This technique is also referred to as hierarchical interpolation (HINT) [978], and is illustrated in Figure 11.18. In the 9×9 image shown in the figure, the pixels marked “1” correspond to the original image decimated by a factor of 4. The decimated image could be encoded using any coding technique. The pixels marked “2” are estimated from those marked “1” by bilinear interpolation, and rounded to ensure reversibility. This completes one iteration of interpolation. Next, the pixels marked “3” are interpolated from the pixels marked “1” and “2”, and the process is repeated to interpolate the pixels marked “4” and “5”. The differences between the actual pixel values marked “2” – “5” and the corresponding interpolated values form a discrete symbol set with a small dynamic range, and may be encoded efficiently using Huffman, arithmetic, or LZW coding.

The illustration in Figure 11.18 corresponds to interpolation of order four; higher-order interpolation may also be used. In general, interpolative coding of order 2^P , where P is an integer, involves 2^P iterations of interpolation.

In the work of Kuduvalli and Rangayyan [174, 338], the digitized radiographic images listed in Table 11.5 were partitioned into blocks of size 256×256 pixels for interpolative coding. The initial subsampled images were decorrelated using 2D DPCM of order 1×1 ; the difference data were compressed by Huffman coding. It was observed that the differences between the interpolated and actual pixel values could be modeled by Laplacian PDFs (see Figure 11.12). The variance of the interpolation errors was seen to decrease with increasing resolution, because pixels are correlated more to their immediate neighbors than to pixels farther away. The interpolation errors at different iterations were modeled by Laplacian PDFs with the variance equal to the corresponding mean-squared interpolation errors. The PDFs were then used in compressing the interpolation errors via arithmetic or Huffman coding. LZW coding does not need modeling of the error distribution, but performed considerably worse than the Huffman and arithmetic codes. Figure 11.19 shows the average bit rate, for eight images, as a function of the order of interpolation using Huffman coding as the post-encoding technique. It is observed that increasing the order of interpolation has only a small effect on the overall average bit rate.

For examples on the performance of HINT, see Tables 11.7, 11.11, 11.15, and 11.16.

1	5	3	5	1	5	3	5	1
5	4	5	4	5	4	5	4	5
3	5	2	5	3	5	2	5	3
5	4	5	4	5	4	5	4	5
1	5	3	5	1	5	3	5	1
5	4	5	4	5	4	5	4	5
3	5	2	5	3	5	2	5	3
5	4	5	4	5	4	5	4	5
1	5	3	5	1	5	3	5	1

FIGURE 11.18

Stages of interpolative coding. The pixels marked “1” are coded and transmitted first. Next, the pixels marked “2” are estimated from those marked “1”, and the differences are transmitted. The procedure continues iteratively with the pixels marked “3”, “4”, and “5”. Reproduced with permission from G.R. Kuduvalli and R.M. Rangayyan, “Performance analysis of reversible image compression techniques for high-resolution digital teleradiology”, *IEEE Transactions on Medical Imaging*, 11(3): 430 – 445, 1992. © IEEE.




FIGURE 11.19

Results of compression of eight of the images listed in Table 11.5 by interpolative and Huffman coding. Order zero corresponds to 1×1 2D DPCM coding. Reproduced with permission from G.R. Kuduvali and R.M. Rangayyan, "Performance analysis of reversible image compression techniques for high-resolution digital teleradiology", *IEEE Transactions on Medical Imaging*, 11(3): 430 - 445, 1992. © IEEE.

11.8 Predictive Coding

Samples of real-life signals and images bear a high degree of correlation, especially over small intervals of time or space. The correlation between samples may also be viewed as statistical redundancy. An outcome of these characteristics is that a sample of a temporal signal may be predicted from a small number of its preceding samples. When such prediction is performed in a linear manner, we have a linear prediction (LP) model, given by [31, 176, 510, 833, 979]

$$\tilde{f}(n) = - \sum_{p=1}^P a(p) f(n-p) + G d(n), \quad (11.52)$$

where $f(n)$ is the signal being modeled; $\tilde{f}(n)$ is the predicted value of $f(n)$; $a(p)$, $p = 1, 2, \dots, P$, are the coefficients of the LP model; and P is the order of the model. The signal $f(n)$ is considered to be the output of a linear system or filter with $d(n)$ as the input or driving signal; G is the gain of the system. Because the prediction model uses only the past values of the output signal $f(n)$, it is known as an autoregressive (AR) model. The need for causality in physically realizable filters and signal processing systems dictates the requirement to use only the past samples of f (and the present value of d) in predicting the current value $f(n)$. If the past values of d are also used, the model will include a moving-average or MA component.

The model represented in Equation 11.52 indicates that, given the initial set of P values of the signal f and the input or driving signal d , any future value of f may be (approximately) computed with a knowledge of the set of coefficients $a(p)$. Therefore, the model coefficients $a(p)$, $p = 1, 2, \dots, P$, represent the signal-generating process. The model coefficients may be used to predict the values of f or to analyze the signal-generating process. Several methods exist to derive the LP model coefficients for a given signal, subject to certain conditions on the error of prediction [31, 176, 510, 833, 979].

In the context of image data compression, we have a few considerations that differ from the temporal signal application described above. In most cases, the input or driving signal d is not known; the omission of the related component in the model represented in Equation 11.52 will cause only a small change in the error of prediction. Furthermore, causality is not a matter of concern in image processing; however, a certain sequence of accessing or processing of the pixels in the given image needs to be defined, which could imply “past” and “future” samples of the image; see Figure 10.19. In the context of image processing, the samples used to predict the current pixel could be labeled as the ROS of the model. Then, we could express the basic 2D LP model as

$$\tilde{f}(m, n) = - \sum_{(p,q) \in ROS} a(p, q) f(m-p, n-q). \quad (11.53)$$

In the application to image coding, the ROS needs to be defined such that, in the decoding process, only those pixels that have already been decoded are included in the ROS for the current pixel being processed.

The error of prediction is given by

$$e(m, n) = f(m, n) - \tilde{f}(m, n), \quad (11.54)$$

and the MSE between the original image and the predicted image is given by

$$\epsilon^2 = E[e^2(m, n)]. \quad (11.55)$$

The coefficients $a(p, q)$ need to be chosen or derived so as to minimize the MSE between the original image and the predicted image. Several approaches are available for the derivation of optimal predictor coefficients [31, 176, 510, 833, 979, 980].

Exact reconstruction of the image requires that, in addition to the initial conditions and the model coefficients, the error of prediction be also transmitted and made available to the decoder. The advantage of LP-based image compression lies in the fact that the error image tends to have a more concentrated PDF than the original image [close to a Laplacian PDF in most cases; see Figure 11.12 (b)], which lends well to efficient data compression.

In the simplest model of prediction, the current pixel $f(m, n)$ may be modeled as being equal to the preceding pixel $f(m, n - 1)$ or $f(m - 1, n)$; let

$$\tilde{f}(m, n) = f(m - 1, n). \quad (11.56)$$

Then, the error of prediction is given by

$$e(m, n) = f(m, n) - f(m - 1, n), \quad (11.57)$$

which represents simple DPCM. [Note: Any data coding method based upon the difference between a data sample and a predicted value of the same using any scheme is referred to as DPCM; hence, all LP-based methods fall under the category of DPCM. The differentiation procedure given by Equation 11.13 and illustrated in Figure 11.12 is equivalent to LP with $\tilde{f}(m, n) = f(m - 1, n)$ as above.] Several simple combinations of the immediate neighbors of the current pixel may also be used in the prediction model; see Section 11.10.2.

Efficient modeling requires the use of the optimal model order (ROS size) and the derivation of the optimal coefficients, subject to conditions related to the minimization of the MSE. Several methods for the derivation of the model coefficients are described in the following sections.

11.8.1 Two-dimensional linear prediction

The error of prediction in the 2D LP model is given by

$$\begin{aligned} e(m, n) &= f(m, n) - \tilde{f}(m, n) \\ &= f(m, n) + \sum_{(p, q) \in ROS} a(p, q) f(m - p, n - q). \end{aligned} \quad (11.58)$$

The squared error is given by

$$e^2(m, n) = f^2(m, n) + 2 \sum_{(p,q) \in ROS} \sum_{(r,s) \in ROS} a(p, q) f(m, n) f(m - p, n - q) \quad (11.59)$$

$$+ \sum_{(p,q) \in ROS} \sum_{(r,s) \in ROS} a(p, q) a(r, s) f(m - p, n - q) f(m - r, n - s).$$

Applying the statistical expectation operator, we get

$$\begin{aligned} \epsilon^2 &= E[e^2(m, n)] \\ &= \phi_f(0, 0) + 2 \sum_{(p,q) \in ROS} \sum_{(r,s) \in ROS} a(p, q) \phi_f(p, q) \\ &+ \sum_{(p,q) \in ROS} \sum_{(r,s) \in ROS} a(p, q) a(r, s) \phi_f(r - p, s - q), \quad (11.60) \end{aligned}$$

where $\phi_f(p, q)$ is the ACF of f , and the image-generating process is assumed to be wide-sense stationary.

The coefficients that minimize the MSE may be derived by setting to zero the derivative of ϵ^2 with respect to $a(p, q)$, $(p, q) \in ROS$, which leads to

$$\phi_f(r, s) + \sum_{(p,q) \in ROS} a(p, q) \phi_f(r - p, s - q) = 0; \quad (r, s) \in ROS. \quad (11.61)$$

Using this result in Equation 11.60, we get

$$\epsilon^2 = \phi_f(0, 0) + \sum_{(p,q) \in ROS} \sum_{(r,s) \in ROS} a(p, q) \phi_f(p, q). \quad (11.62)$$

Combining Equation 11.61 and 11.62, we get

$$\phi_f(r, s) + \sum_{(p,q) \in ROS} a(p, q) \phi_f(r - p, s - q) = \begin{cases} 0 & (r, s) \in ROS \\ \epsilon^2 & (r, s) = (0, 0) \end{cases}. \quad (11.63)$$

The equations represented by the expressions above are known as the 2D normal or Yule–Walker equations, and may be solved to derive the prediction coefficients. Because the method uses the ACF of the image to derive the prediction coefficients, it is known as the autocorrelation method [510]. The ACF may be estimated from the given finite image $f(m, n)$, $m = 0, 1, 2, \dots, M - 1$; $n = 0, 1, 2, \dots, N - 1$, as

$$\tilde{\phi}_f(p, q) = \frac{1}{MN} \sum_{m=p}^{M-1} \sum_{n=q}^{N-1} f(m, n) f(m - p, n - q). \quad (11.64)$$

The prediction coefficients may also be estimated by using least-squares methods to minimize the prediction error averaged over the entire image,

indicated as $(m, n) \in IMG$, as follows:

$$\begin{aligned}
 \varepsilon^2 &= \frac{1}{MN} \sum_{(m,n) \in IMG} \sum_{(m,n) \in IMG} e^2(m, n) \\
 &= \frac{1}{MN} \sum_{(m,n) \in IMG} \sum_{(m,n) \in IMG} \left[f^2(m, n) + 2 \sum_{(p,q) \in ROS} a(p, q) f(m, n) f(m-p, n-q) \right. \\
 &\quad \left. + \sum_{(p,q) \in ROS} \sum_{(r,s) \in ROS} a(p, q) a(r, s) f(m-p, n-q) f(m-r, n-s) \right] \\
 &= \sigma_f(0, 0; 0, 0) + 2 \sum_{(p,q) \in ROS} a(p, q) \sigma_f(0, 0; p, q) \\
 &\quad + \sum_{(p,q) \in ROS} \sum_{(r,s) \in ROS} a(p, q) a(r, s) \sigma_f(p, q; r, s). \tag{11.65}
 \end{aligned}$$

Here, σ_f is the covariance of the image f , defined as

$$\sigma_f(p, q; r, s) = \frac{1}{MN} \sum_{(m,n) \in IMG} \sum_{(m,n) \in IMG} f(m-p, n-q) f(m-r, n-s). \tag{11.66}$$

The coefficients that minimize the averaged error may be derived by setting to zero the derivative of ε^2 with respect to $a(p, q)$, which leads to

$$\sigma_f(0, 0; r, s) + \sum_{(p,q) \in ROS} a(p, q) \sigma_f(p, q; r, s) = 0. \tag{11.67}$$

It follows that

$$\varepsilon^2 = \sigma_f(0, 0; 0, 0) + \sum_{(p,q) \in ROS} a(p, q) \sigma_f(0, 0; p, q). \tag{11.68}$$

The 2D normal equations for this condition are given by

$$\sigma_f(0, 0; r, s) + \sum_{(p,q) \in ROS} a(p, q) \sigma_f(p, q; r, s) = \begin{cases} 0 & (r, s) \in ROS \\ \varepsilon^2 & (r, s) = (0, 0) \end{cases}, \tag{11.69}$$

which may be solved to obtain the prediction coefficients. Because the covariance of the image is used to derive the prediction coefficients, this method is known as the covariance method.

Now, if the region IMG is defined so as to span the entire image array of size $M \times N$, and the image is assumed to be zero outside the range $m = 0, 1, 2, \dots, M-1$; $n = 0, 1, 2, \dots, N-1$, we have

$$\begin{aligned}
 \sigma_f(p, q; r, s) &= \frac{1}{MN} \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} f(m-p, n-q) f(m-r, n-s) \\
 &= \tilde{\phi}_f(r-p, s-q), \tag{11.70}
 \end{aligned}$$

where $\tilde{\phi}_f$ is an estimate of the ACF of the image f . Then, the covariance method yields results that are identical to the results given by the autocorrelation method.

Equation 11.63 may be expressed in matrix form as

$$\Phi_f \mathbf{a} = \boldsymbol{\epsilon}, \tag{11.71}$$

where, for the case of a QP ROS of size $P_1 \times P_2$, the matrices Φ_f , \mathbf{a} , and $\boldsymbol{\epsilon}$ are of size $(P_1 + 1)(P_2 + 1) \times (P_1 + 1)(P_2 + 1)$, $(P_1 + 1)(P_2 + 1) \times 1$, and $(P_1 + 1)(P_2 + 1) \times 1$, respectively. The extended ACF matrix Φ_f is given by

$$\begin{bmatrix} \phi(0, 0) & \cdots & \phi(0, -P_2) & \phi(-1, 0) & \cdots & \phi(-1, -P_2) & \phi(-2, 0) & \cdots & \phi(-P_1, -P_2) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \phi(0, P_2) & \cdots & \phi(0, 0) & \phi(-1, P_2) & \cdots & \phi(-1, 0) & \phi(-2, P_2) & \cdots & \phi(-P_1, 0) \\ \phi(1, 0) & \cdots & \phi(1, -P_2) & \phi(0, 0) & \cdots & \phi(0, -P_2) & \phi(-1, 0) & \cdots & \phi(-P_1 + 1, -P_2) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \phi(1, P_2) & \cdots & \phi(1, 0) & \phi(0, P_2) & \cdots & \phi(0, 0) & \phi(-1, P_2) & \cdots & \phi(-P_1 + 1, 0) \\ \phi(2, 0) & \cdots & \phi(2, -P_2) & \phi(1, 0) & \cdots & \phi(1, -P_2) & \phi(0, 0) & \cdots & \phi(-P_1 + 2, -P_2) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \phi(P_1, P_2) & \cdots & \phi(P_1, 0) & \phi(P_1 - 1, P_2) & \cdots & \phi(P_1 - 1, 0) & \phi(P_1 - 2, P_2) & \cdots & \phi(0, 0) \end{bmatrix}, \tag{11.72}$$

where the subscript f has been dropped from the entries within the matrix for the sake of brevity. The matrices composed by the prediction coefficients and the error are given by

$$\mathbf{a} = \begin{bmatrix} a(0, 0) \\ \vdots \\ a(0, P_2) \\ a(1, 0) \\ \vdots \\ a(1, P_2) \\ a(2, 0) \\ \vdots \\ a(P_1, P_2) \end{bmatrix} \quad \text{and} \quad \boldsymbol{\epsilon} = \begin{bmatrix} \epsilon^2 \\ \vdots \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}. \tag{11.73}$$

The matrix Φ_f is Toeplitz-block-Toeplitz in nature; efficient algorithms are available for the inversion of such matrices [981, 982].

The methods described above to compute the prediction coefficients assume the image to be stationary over the entire frame available. In practice, most images are nonstationary, and may be assumed to be locally stationary only over relatively small segments or ROIs. In order to maintain the optimality of the model over the entire image, and in order to maintain the error of prediction at low levels, we could follow one of two procedures:

- partition the image into small blocks over which stationarity may be assumed, and compute the prediction coefficients independently for each block, or

- adapt the model to the changing statistics of the image.

Encoding the prediction error: In order to facilitate error-free reconstruction of the image, the prediction error has to be transmitted and made available at the decoder (in addition to the prediction coefficients and the initial conditions). For quantized original pixel values, the prediction error may be rounded off to integers that may be encoded without error using a source coding technique such as the Huffman code. The prediction error has been observed to possess a Laplacian PDF [174, 338], which lends well to efficient compression by the Huffman code.

Results of application to medical images: The average bit rates obtained by the application of the autocorrelation method of computing the prediction coefficients, using blocks of size 128×128 pixels, to six of the images listed in Table 11.5, are shown in Figure 11.20, for various model orders. For the images used, comparable performance was obtained using NSHP and QP ROSs of similar extent. Good compression performance was obtained, with average bit rates in the range $2.5 - 3.2$ *b/pixel* for most of the images, with the original pixel values at 10 *b/pixel*, using QP ROSs of size 2×2 or 3×3 .

See Section 5.4.10 for a method for the detection of calcifications based upon the error of prediction.

11.8.2 Multichannel linear prediction

The difference between LP in 1D and 2D may be bridged by multichannel LP, where a certain number of rows of the given image may be viewed as a collection of multichannel signals [337, 338, 979, 983, 984, 985, 986]. Kuduvali and Rangayyan [174, 337, 338] proposed the following procedures for the application of multichannel LP to predictive coding and compression of 2D images.

Consider a multichannel signal with $(P_2 + 1)$ channels, with the channels indexed as $q = 0, 1, 2, \dots, P_2$, and the individual signals labeled as $\mathbf{f}(m)$, $m = 0, 1, 2, \dots, N - 1$. The collection of the signal's values at a position m , given by $\mathbf{f}_q(m)$, $q = 0, 1, 2, \dots, P_2$, may be viewed as a multichannel signal or vector (or a matrix) of size $(P_2 + 1) \times 1$; see Figures 11.21 and 11.22. If we were to use a multichannel linear predictor of order P_1 , we could predict the vector $\mathbf{f}(m)$ as a linear combination of the vectors $\mathbf{f}(m - p)$, $p = 1, 2, \dots, P_1$:

$$\tilde{\mathbf{f}}(m) = - \sum_{p=1}^{P_1} \mathbf{a}(p) \mathbf{f}(m - p), \quad (11.74)$$

where $\mathbf{a}(p)$, $p = 1, 2, \dots, P_1$, are multichannel LP coefficient matrices, each of size $(P_2 + 1) \times (P_2 + 1)$.


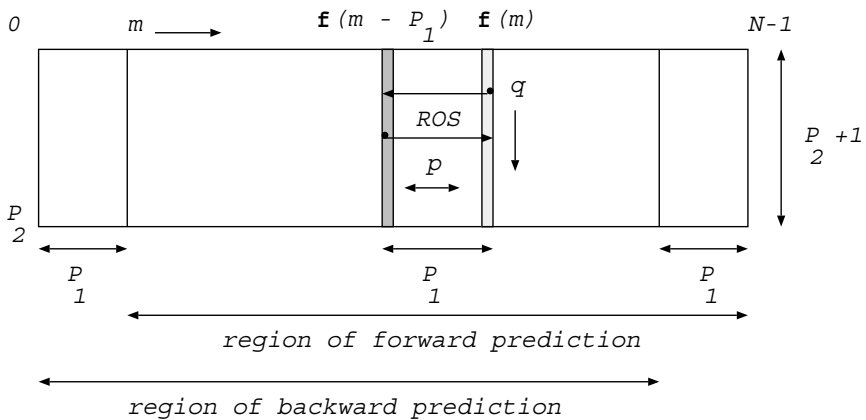


FIGURE 11.20

Results of compression of six of the images listed in Table 11.5 by 2D LP (autocorrelation method) and Huffman coding. The method was applied on a block-by-block basis, using blocks of size 128×128 pixels. In the case of modeling using the NSHP ROS, a model order of 1.5 indicates a $1 \times 1 \times 1$ ROS (see Figure 10.19), 2.5 indicates a $2 \times 2 \times 2$ ROS, etc. The orders of models using the QP ROS are indicated by integers: 2 indicates a 2×2 ROS, 3 indicates a 3×3 ROS, etc. Reproduced with permission from G.R. Kuduvalli and R.M. Rangayyan, "Performance analysis of reversible image compression techniques for high-resolution digital teleradiology", *IEEE Transactions on Medical Imaging*, 11(3): 430 - 445, 1992. © IEEE.

**FIGURE 11.21**

Multichannel linear prediction. Each row of the image is viewed as a channel or component of a multichannel signal or vector. The column index of the image may be considered to be equivalent to a temporal index [174, 337, 338]. The indices shown correspond to Equations 11.74 and 11.88. See also Figure 11.22.

The error of prediction is given by

$$\mathbf{e}(m) = \mathbf{f}(m) + \sum_{p=1}^{P_1} \mathbf{a}(p) \mathbf{f}(m-p). \quad (11.75)$$

The covariance matrix of the error of prediction is given by

$$\boldsymbol{\sigma}_e = E[\mathbf{e}(m) \mathbf{e}^T(m)]. \quad (11.76)$$

For optimal prediction, we need to derive the prediction coefficient matrices $\mathbf{a}(p)$ that minimize the trace of the covariance matrix of the error of prediction. From Equation 11.75 we can write

$$\begin{aligned} \mathbf{e}(m) \mathbf{e}^T(m) &= \mathbf{f}(m) \mathbf{f}^T(m) \\ &+ \sum_{p=1}^{P_1} \mathbf{a}(p) \mathbf{f}(m-p) \mathbf{f}^T(m) + \sum_{q=1}^{P_1} \mathbf{f}(m) \mathbf{f}^T(m-q) \mathbf{a}^T(q) \\ &+ \sum_{p=1}^{P_1} \sum_{q=1}^{P_1} \mathbf{a}(p) \mathbf{f}(m-p) \mathbf{f}^T(m-q) \mathbf{a}^T(q). \end{aligned} \quad (11.77)$$

Applying the statistical expectation operator, and assuming wide-sense stationarity of the multichannel signal-generating process, we get

$$\boldsymbol{\sigma}_e = \boldsymbol{\phi}_c(0) + \sum_{p=1}^{P_1} \mathbf{a}(p) \boldsymbol{\phi}_c(-p)$$




FIGURE 11.22

Multichannel LP applied to a 2D image [337, 338]. See also Figure 11.21. Reproduced with permission from G.R. Kuduvalli and R.M. Rangayyan, "An algorithm for direct computation of 2-D linear prediction coefficients", *IEEE Transactions on Signal Processing*, 41(2): 996–1000, 1993. © IEEE.

$$+ \sum_{q=1}^{P_1} \phi_c(q) \mathbf{a}^T(q) + \sum_{p=1}^{P_1} \sum_{q=1}^{P_1} \mathbf{a}(p) \phi_c(q-p) \mathbf{a}^T(q), \quad (11.78)$$

where ϕ_c is the ACF of the image computed over the set of rows or channels being used in the multichannel prediction model, given by

$$\phi_c(r) = \begin{bmatrix} \phi_{0,0}(r) & \cdots & \phi_{0,P_2}(r) \\ \vdots & \ddots & \vdots \\ \phi_{P_2,0}(r) & \cdots & \phi_{P_2,P_2}(r) \end{bmatrix}, \quad (11.79)$$

and

$$\phi_{p,q}(r) = E[\mathbf{f}_p(m) \mathbf{f}_q(m-r)]. \quad (11.80)$$

In order to minimize the trace of the error covariance matrix σ_e , we could differentiate both the sides of Equation 11.78 with respect to the prediction coefficient matrices $\mathbf{a}(r)$, $r = 1, 2, \dots, P_1$, and equate the result to the null matrix of size $(P_2 + 1) \times (P_2 + 1)$, which leads to

$$\begin{aligned} \mathbf{0} &= \frac{\partial \sigma_e}{\partial \mathbf{a}(r)}; \quad r = 1, 2, \dots, P_1 \\ &= 2 \phi_c(-r) + 2 \sum_{p=1}^{P_1} \phi_c^T(r-p) \mathbf{a}^T(p) \\ &= \phi_c(-r) + \sum_{p=1}^{P_1} \phi_c(p-r) \mathbf{a}^T(p); \quad r = 1, 2, \dots, P_1. \end{aligned} \quad (11.81)$$

[Note: $\phi_c^T(r-p) = \phi_c(p-r)$.]

Now, Equation 11.78 may be rewritten as

$$\begin{aligned} \sigma_e &= \phi_c(0) + \sum_{p=1}^{P_1} \mathbf{a}(p) \phi_c(p) \\ &+ \sum_{q=1}^{P_1} \left[\phi_c(-q) + \sum_{p=1}^{P_1} \mathbf{a}(p) \phi_c(q-p) \right] \mathbf{a}^T(q) \\ &= \phi_c(0) + \sum_{p=1}^{P_1} \mathbf{a}(p) \phi_c^T(p) \\ &= \phi_c(0) + \sum_{p=1}^{P_1} \phi_c(p) \mathbf{a}^T(p). \end{aligned} \quad (11.82)$$

The relationships derived above may be summarized as

$$\Phi_c \mathbf{A} = \sigma, \quad (11.83)$$

where the matrices are given in expanded form as

$$\begin{bmatrix} \phi_c(0) & \phi_c(1) & \cdots & \phi_c(P_1) \\ \phi_c(-1) & \phi_c(0) & \cdots & \phi_c(P_1 - 1) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_c(-P_1) & \phi_c(-P_1 + 1) & \cdots & \phi_c(0) \end{bmatrix} \begin{bmatrix} \mathbf{I} \\ \mathbf{a}^T(1) \\ \vdots \\ \mathbf{a}^T(P_1) \end{bmatrix} = \begin{bmatrix} \boldsymbol{\sigma}_e \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix}. \quad (11.84)$$

In the equation above, the submatrices ϕ_c are as defined in Equation 11.79; \mathbf{I} is the identity matrix of size $(P_2 + 1) \times (P_2 + 1)$; and $\mathbf{0}$ is the null matrix of size $(P_2 + 1) \times (P_2 + 1)$. This system of equations may be referred to as the multichannel version of the Yule-Walker equations. The solution to this set of equations may be used to obtain the 2D LP coefficients by making the following associations:

$$\phi_{p,q}(r) = \phi_f(r, p - q), \quad (11.85)$$

(compare Equations 11.80 and 11.64);

$$\boldsymbol{\epsilon} = \boldsymbol{\sigma}_e \mathbf{a}_0; \quad (11.86)$$

and

$$\mathbf{a}_r = \mathbf{a}^T(r) \mathbf{a}_0, \quad (11.87)$$

where $\mathbf{a}_r = [a(r, 0) \ a(r, 1) \ \cdots \ a(r, P_2)]^T$ is composed by the elements of the r^{th} row of the matrix of prediction coefficients \mathbf{a} given in Equation 11.73 (written as a column matrix).

The Levinson-Wiggins-Robinson algorithm: The multichannel prediction coefficient matrix may be obtained by the application of the algorithms due to Levinson [987] and Wiggins and Robinson [988]. In the multichannel version of this algorithm [337, 338], the prediction coefficients for order $(P_1 + 1)$ are recursively related to those for order P_1 . The prediction model given by Equation 11.74 is known as the forward predictor. Going in the opposite direction, the backward predictor is defined to predict the vector $\mathbf{f}(m)$ in terms of the vectors $\mathbf{f}(m + p)$, $p = 1, 2, \dots, P_1$, as

$$\tilde{\mathbf{f}}(m) = - \sum_{p=1}^{P_1} \mathbf{b}(p) \mathbf{f}(m + p), \quad (11.88)$$

where $\mathbf{b}(p)$, $p = 1, 2, \dots, P_1$, are the multichannel backward prediction coefficient matrices, each of size $(P_2 + 1) \times (P_2 + 1)$; see Figure 11.21.

In order to derive the multichannel version of Levinson's algorithm, let us rewrite the multichannel ACF matrix in Equations 11.83 and 11.84 as follows:

$$\boldsymbol{\Phi}_{P_1+1} = \begin{bmatrix} \phi(0) & \phi(1) & \cdots & \phi(P_1) \\ \phi(-1) & \phi(0) & \cdots & \phi(P_1 - 1) \\ \vdots & \vdots & \ddots & \vdots \\ \phi(-P_1) & \phi(-P_1 + 1) & \cdots & \phi(0) \end{bmatrix}, \quad (11.89)$$

where the subscript $(P_1 + 1)$ is used to indicate the order of the model, and the subscript c has been dropped from the submatrices for compact notation. The matrix Φ_{P_1+1} may be partitioned as

$$\Phi_{P_1+1} = \begin{bmatrix} \phi(0) & \psi_{P_1}^T \\ \psi_{P_1} & \Phi_{P_1} \end{bmatrix} = \begin{bmatrix} \Phi_{P_1} & \psi_{P_1}^\# \\ \psi_{P_1}^{\#T} & \phi(0) \end{bmatrix}, \quad (11.90)$$

where

$$\psi_{P_1} = [\phi(1) \ \phi(2) \ \cdots \ \phi(P_1)]^T, \quad (11.91)$$

$$\psi_{P_1}^\# = [\phi(-P_1) \ \phi(-P_1 + 1) \ \cdots \ \phi(-1)]^T, \quad (11.92)$$

and the property that $\phi(r) = \phi^T(-r)$ has been used. It follows that

$$\psi_{P_1} = \begin{bmatrix} \psi_{P_1-1} \\ \phi(-P_1) \end{bmatrix}, \quad (11.93)$$

and

$$\psi_{P_1}^\# = \begin{bmatrix} \phi(P_1) \\ \psi_{P_1-1}^\# \end{bmatrix}. \quad (11.94)$$

Let us also define partitions of the forward and backward prediction coefficient matrices as follows:

$$\mathbf{A}_{P_1} = \begin{bmatrix} \mathbf{I} \\ \tilde{\mathbf{A}}_{P_1} \end{bmatrix}, \quad (11.95)$$

and

$$\mathbf{B}_{P_1} = \begin{bmatrix} \tilde{\mathbf{B}}_{P_1} \\ \mathbf{I} \end{bmatrix}, \quad (11.96)$$

where \mathbf{A}_{P_1} is the same as \mathbf{A} in Equations 11.83 and 11.84, and \mathbf{B}_{P_1} is formed in a similar manner for the backward predictor.

Using the partitions as defined above, we may rewrite the multichannel Yule-Walker equations, given by Equation 11.84, in two forms for forward and backward prediction, as follows:

$$\begin{bmatrix} \phi(0) & \psi_{P_1}^T \\ \psi_{P_1}^T & \Phi_{P_1} \end{bmatrix} \begin{bmatrix} \mathbf{I} \\ \tilde{\mathbf{A}}_{P_1} \end{bmatrix} = \begin{bmatrix} \sigma_{P_1}^f \\ \mathbf{0}_{P_1} \end{bmatrix}, \quad (11.97)$$

and

$$\begin{bmatrix} \Phi_{P_1} & \psi_{P_1}^\# \\ \psi_{P_1}^{\#T} & \phi(0) \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{B}}_{P_1} \\ \mathbf{I} \end{bmatrix} = \begin{bmatrix} \mathbf{0}_{P_1} \\ \sigma_{P_1}^b \end{bmatrix}, \quad (11.98)$$

where $\mathbf{0}_{P_1}$ is the null matrix of size $(P_1 P_2 + 1) \times (P_2 + 1)$. The matrix $\sigma_{P_1}^f$ is the covariance matrix of the error of forward prediction (the same as σ_e

given by Equation 11.82), with the matrix $\sigma_{P_1}^b$ being the counterpart for the backward predictor. It follows that

$$\tilde{\mathbf{A}}_{P_1} = -\Phi_{P_1}^{-1} \psi_{P_1}, \quad (11.99)$$

$$\tilde{\mathbf{B}}_{P_1} = -\Phi_{P_1}^{-1} \psi_{P_1}^\#, \quad (11.100)$$

$$\sigma_{P_1}^f = \phi(0) + \psi_{P_1}^T \tilde{\mathbf{A}}_{P_1}, \quad (11.101)$$

and

$$\sigma_{P_1}^b = \phi(0) + \psi_{P_1}^{\#T} \tilde{\mathbf{B}}_{P_1}. \quad (11.102)$$

Applying the inversion theorem for partitioned matrices [979], and making use of the preceding six relationships, we get

$$\Phi_{P_1+1}^{-1} = \begin{bmatrix} \mathbf{0} & \mathbf{0}_{P_1}^T \\ \mathbf{0}_{P_1} & \Phi_{P_1}^{-1} \end{bmatrix} + \mathbf{A}_{P_1} [\sigma_{P_1}^f]^{-1} \mathbf{A}_{P_1}^T \quad (11.103)$$

$$= \begin{bmatrix} \Phi_{P_1}^{-1} & \mathbf{0}_{P_1} \\ \mathbf{0}_{P_1}^T & \mathbf{0} \end{bmatrix} + \mathbf{B}_{P_1} [\sigma_{P_1}^b]^{-1} \mathbf{B}_{P_1}^T. \quad (11.104)$$

Multiplying both sides of Equation 11.104 by ψ_{P_1+1} , making use of the partitioned form in Equation 11.93, and using Equation 11.99, we get

$$\tilde{\mathbf{A}}_{P_1+1} = \begin{bmatrix} \tilde{\mathbf{A}}_{P_1} \\ \mathbf{0} \end{bmatrix} - \mathbf{B}_{P_1} [\sigma_{P_1}^b]^{-1} \mathbf{B}_{P_1}^T \psi_{P_1+1}. \quad (11.105)$$

Extracting the lower $(P_2 + 1) \times (P_2 + 1)$ matrix from both sides of Equation 11.105 in its partitioned form, we have

$$\mathbf{a}_{P_1+1}^T(P_1 + 1) = -[\sigma_{P_1}^b]^{-1} \mathbf{B}_{P_1}^T \psi_{P_1+1}, \quad (11.106)$$

which upon transposition yields

$$\begin{aligned} \mathbf{a}_{P_1+1}(P_1 + 1) &= -\psi_{P_1+1}^T \mathbf{B}_{P_1} [\sigma_{P_1}^b]^{-1} \\ &= -[\Delta_{P_1+1}^b]^T [\sigma_{P_1}^b]^{-1}, \end{aligned} \quad (11.107)$$

where

$$\begin{aligned} \Delta_{P_1+1}^b &= \mathbf{B}_{P_1}^T \psi_{P_1+1} \\ &= \phi(-P_1 - 1) + \sum_{p=1}^{P_1} \mathbf{b}_{P_1}(p) \phi(-P_1 - 1 + p). \end{aligned} \quad (11.108)$$

Using Equation 11.107 in Equation 11.105, we get

$$\tilde{\mathbf{A}}_{P_1+1} = \begin{bmatrix} \tilde{\mathbf{A}}_{P_1} \\ \mathbf{0} \end{bmatrix} + \mathbf{B}_{P_1} \mathbf{a}_{P_1+1}^T(P_1 + 1), \quad (11.109)$$

which upon transposition and expansion of the matrix notation yields

$$\mathbf{a}_{P_1+1}(p) = \mathbf{a}_{P_1}(p) + \mathbf{a}_{P_1+1}(P_1+1) \mathbf{b}_{P_1}(P_1+1-p); \quad p = 1, 2, \dots, P_1. \quad (11.110)$$

Similarly, multiplying both sides of Equation 11.103 by $\boldsymbol{\psi}_{P_1+1}^\#$ and using Equation 11.100, we get

$$\mathbf{b}_{P_1+1}(P_1+1) = -[\boldsymbol{\Delta}_{P_1+1}^f]^{-1} [\boldsymbol{\sigma}_{P_1}^f]^{-1}, \quad (11.111)$$

where

$$\boldsymbol{\Delta}_{P_1+1}^f = \boldsymbol{\phi}(P_1+1) + \sum_{p=1}^{P_1} \mathbf{a}_{P_1}(p) \boldsymbol{\phi}(P_1+1-p), \quad (11.112)$$

and

$$\mathbf{b}_{P_1+1}(p) = \mathbf{b}_{P_1}(p) + \mathbf{b}_{P_1+1}(P_1+1) \mathbf{a}_{P_1}(P_1+1-p); \quad p = 1, 2, \dots, P_1. \quad (11.113)$$

Substituting Equation 11.109 in Equation 11.101 and using the partitioned form of $\boldsymbol{\psi}_{P_1+1}$ in Equation 11.93, we get

$$\begin{aligned} \boldsymbol{\sigma}_{P_1+1}^f &= \boldsymbol{\phi}(0) + \boldsymbol{\psi}_{P_1}^T \tilde{\mathbf{A}}_{P_1} + \boldsymbol{\psi}_{P_1+1}^T \mathbf{B}_{P_1} \mathbf{a}_{P_1+1}^T(P_1+1) \\ &= \boldsymbol{\sigma}_{P_1}^f + [\boldsymbol{\Delta}_{P_1+1}^b]^{-1} \mathbf{a}_{P_1+1}^T(P_1+1) \\ &= \boldsymbol{\sigma}_{P_1}^f - \mathbf{a}_{P_1+1}(P_1+1) \boldsymbol{\sigma}_{P_1}^b \mathbf{a}_{P_1+1}^T(P_1+1), \end{aligned} \quad (11.114)$$

where Equation 11.107 has been used in the last step. Similarly, we can obtain an expression for the covariance matrix of the backward prediction error as

$$\boldsymbol{\sigma}_{P_1+1}^b = \boldsymbol{\sigma}_{P_1}^b - \mathbf{b}_{P_1+1}(P_1+1) \boldsymbol{\sigma}_{P_1}^f \mathbf{b}_{P_1+1}^T(P_1+1). \quad (11.115)$$

Now, consider the matrix product

$$\begin{aligned} \begin{bmatrix} \mathbf{0} & \mathbf{B}_{P_1}^T \end{bmatrix} \boldsymbol{\Phi}_{P_1+2} \begin{bmatrix} \mathbf{A}_{P_1} \\ \mathbf{0} \end{bmatrix} &= \begin{bmatrix} \mathbf{0} & \mathbf{B}_{P_1}^T \end{bmatrix} \begin{bmatrix} \boldsymbol{\phi}(0) & \boldsymbol{\psi}_{P_1+1}^T \\ \boldsymbol{\psi}_{P_1+1} & \boldsymbol{\Phi}_{P_1+1} \end{bmatrix} \begin{bmatrix} \mathbf{A}_{P_1} \\ \mathbf{0} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{B}_{P_1}^T \boldsymbol{\psi}_{P_1+1} & \mathbf{B}_{P_1}^T \boldsymbol{\Phi}_{P_1+1} \end{bmatrix} \begin{bmatrix} \mathbf{A}_{P_1} \\ \mathbf{0} \end{bmatrix} \\ &= \begin{bmatrix} \boldsymbol{\Delta}_{P_1+1}^b & \mathbf{0} & \cdots & \boldsymbol{\sigma}_{P_1}^b \end{bmatrix} \begin{bmatrix} \mathbf{I} \\ \tilde{\mathbf{A}}_{P_1} \\ \mathbf{0} \end{bmatrix} \\ &= \boldsymbol{\Delta}_{P_1+1}^b. \end{aligned} \quad (11.116)$$

Taking the transpose of the expression above, and noting that $\boldsymbol{\Phi}_{P_1+2}$ is symmetric, we get

$$[\boldsymbol{\Delta}_{P_1+1}^b]^{-1} = \begin{bmatrix} \mathbf{A}_{P_1}^T & \mathbf{0} \end{bmatrix} \boldsymbol{\Phi}_{P_1+2} \begin{bmatrix} \mathbf{0} \\ \mathbf{B}_{P_1} \end{bmatrix} = \boldsymbol{\Delta}_{P_1+1}^f. \quad (11.117)$$

Equations 11.105, 11.107, 11.109 – 11.112, 11.114, 11.115, and 11.117 constitute the Levinson-Wiggins-Robinson algorithm, with the initialization

$$\sigma_0^f = \sigma_0^b = \phi_c(0). \quad (11.118)$$

With the autocorrelation matrices $\phi(p)$ defined by the association given in Equation 11.85, the matrix Φ_{P_1+1} is a Toeplitz-block-Toeplitz matrix: the block elements (submatrices) $\phi(p)$ along the diagonals of Φ_{P_1+1} are mutually identical, and furthermore, the elements along the diagonals of each submatrix $\phi(p)$ are mutually identical. Thus, Φ_{P_1+1} and $\phi(p)$ are symmetrical about their cross diagonals (that is, they are per-symmetric). A property of Toeplitz-block-Toeplitz matrices that is of interest here is defined in terms of the exchange matrices

$$\mathbf{J} = \begin{bmatrix} 0 & 0 & \cdots & 0 & 1 \\ 0 & 0 & \cdots & 1 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 1 & \cdots & 0 & 0 & 0 \end{bmatrix} \quad (11.119)$$

of size $(P_2 + 1) \times (P_2 + 1)$, and

$$\mathbf{J}_{P_1} = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{J} \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{J} & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{J} & \cdots & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix} \quad (11.120)$$

of size $(P_1 + 1)(P_2 + 1) \times (P_1 + 1)(P_2 + 1)$, such that $\mathbf{J}\mathbf{J} = \mathbf{I}$ and $\mathbf{J}_{P_1+1}\mathbf{J}_{P_1+1} = \mathbf{I}_{P_1+1}$. With these definitions, we have

$$\mathbf{J}\phi(p)\mathbf{J} = \phi^T(p), \quad (11.121)$$

and

$$\mathbf{J}_{P_1+1}\Phi_{P_1+1}\mathbf{J}_{P_1+1} = \Phi_{P_1+1}. \quad (11.122)$$

Now, premultiplying both sides of Equation 11.84 by \mathbf{J}_{P_1+1} and post-multiplying both sides by \mathbf{J} , we get

$$\begin{bmatrix} \phi(0) & \phi(-1) & \cdots & \phi(-P_1) \\ \phi(1) & \phi(0) & \cdots & \phi(-P_1 + 1) \\ \vdots & \vdots & \ddots & \vdots \\ \phi(P_1) & \phi(P_1 - 1) & \cdots & \phi(0) \end{bmatrix} \begin{bmatrix} \mathbf{J} \mathbf{a}_{P_1}^T(P_1) \mathbf{J} \\ \vdots \\ \mathbf{J} \mathbf{a}_{P_1}^T(1) \mathbf{J} \\ \mathbf{I} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \mathbf{J} \sigma_{P_1}^f \mathbf{J} \end{bmatrix}. \quad (11.123)$$

Equation 11.123 is identical to the modified Yule-Walker equations for computing the matrices $\mathbf{b}_{P_1}(p)$ and $\sigma_{P_1}^b$ in Equation 11.98. Comparing the terms in the two equations, we get

$$\mathbf{b}_{P_1}(p) = \mathbf{J} \mathbf{a}_{P_1}^T(p) \mathbf{J}; \quad p = 1, 2, \dots, P_1, \quad (11.124)$$

and

$$\sigma_{P_1}^b = \mathbf{J} \sigma_{P_1}^f \mathbf{J} = \sigma_{P_1}. \quad (11.125)$$

With these simplifications, the recursive procedures in the multichannel Levinson algorithm may be modified for the computation of 2D LP coefficients as follows:

$$\Delta_{P_1+1} = \phi_f(P_1 + 1) + \sum_{p=1}^{P_1} \mathbf{a}_{P_1}(p) \phi_f(P_1 + 1 - p), \quad (11.126)$$

$$\mathbf{a}_{P_1+1}(P_1 + 1) = -\Delta_{P_1+1} \mathbf{J} \sigma_{P_1}^{-1} \mathbf{J}, \quad (11.127)$$

$$\mathbf{a}_{P_1+1}(p) = \mathbf{a}_{P_1}(p) + \mathbf{a}_{P_1+1}(P_1 + 1) \mathbf{J} \mathbf{a}_{P_1}(P_1 + 1 - p) \mathbf{J}; \quad p = 1, 2, \dots, P_1, \quad (11.128)$$

and

$$\sigma_{P_1+1} = \sigma_{P_1} - \mathbf{a}_{P_1+1}(P_1 + 1) \mathbf{J} \sigma_{P_1} \mathbf{J} \mathbf{a}_{P_1+1}^T(P_1 + 1), \quad (11.129)$$

with the initialization

$$\sigma_0 = \phi_f(0), \quad (11.130)$$

where

$$\phi_f(p) = \begin{bmatrix} \phi_f(p, 0) & \cdots & \phi_f(p, -P_2) \\ \vdots & \ddots & \vdots \\ \phi_f(p, P_2) & \cdots & \phi_f(p, 0) \end{bmatrix}. \quad (11.131)$$

Equations 11.126 – 11.131 constitute the 2D Levinson algorithm for solving the 2D Yule–Walker equations for the case of a QP ROS. The Levinson algorithm provides results that are identical to those obtained by direct inversion of Φ_f .

Computation of the 2D LP coefficients directly from the image data: The multichannel version of the Levinson algorithm may be used to derive the multichannel version of the Burg algorithm [986, 337, 338], as follows. Equation 11.109 may be augmented using the partition shown in Equation 11.95 as

$$\mathbf{A}_{P_1+1} = \begin{bmatrix} \mathbf{A}_{P_1} \\ \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{B}_{P_1} \end{bmatrix} \mathbf{a}_{P_1+1}^T(P_1 + 1). \quad (11.132)$$

From Equation 11.75, the forward prediction error vector for order $(P_1 + 1)$ is

$$\begin{aligned} \mathbf{e}_{P_1+1}^f(m) &= \mathbf{f}(m) + \sum_{p=1}^{P_1+1} \mathbf{a}_{P_1+1}(p) \mathbf{f}(m - p) \\ &= \mathbf{A}_{P_1+1}^T \mathbf{F}_{P_1+1}(m), \end{aligned} \quad (11.133)$$

where $\mathbf{F}_{P_1+1}(m)$ is the multichannel data matrix of order $(P_1 + 1)$ at (m) , which may be partitioned as

$$\mathbf{F}_{P_1+1}(m) = \begin{bmatrix} \mathbf{f}(m) \\ \mathbf{F}_{P_1}(m-1) \end{bmatrix} = \begin{bmatrix} \mathbf{F}_{P_1}(m) \\ \mathbf{f}(m - P_1 - 1) \end{bmatrix}. \quad (11.134)$$

Similarly, the backward prediction error vector may be expressed as

$$\begin{aligned} \mathbf{e}_{P_1+1}^b(m) &= \mathbf{f}(m - P_1 - 1) + \sum_{p=1}^{P_1+1} \mathbf{b}_{P_1+1}(p) \mathbf{f}(m - P_1 + p) \\ &= \mathbf{B}_{P_1+1}^T \mathbf{F}_{P_1+1}(m). \end{aligned} \quad (11.135)$$

Transposing both sides of Equation 11.132 and multiplying by $\mathbf{F}_{P_1+1}(m)$, using the partitioned forms shown on the right-hand side of Equation 11.134, as well as using Equations 11.133 and 11.135, we get

$$\mathbf{e}_{P_1+1}^f(m) = \mathbf{e}_{P_1}^f(m) + \mathbf{a}_{P_1+1}(P_1 + 1) \mathbf{e}_{P_1}^b(m - 1). \quad (11.136)$$

Similarly, the backward prediction error vector is given by

$$\mathbf{e}_{P_1+1}^b(m) = \mathbf{e}_{P_1}^b(m - 1) + \mathbf{b}_{P_1+1}(P_1 + 1) \mathbf{e}_{P_1}^f(m). \quad (11.137)$$

The matrices $\mathbf{a}_{P_1+1}(P_1 + 1)$ and $\mathbf{b}_{P_1+1}(P_1 + 1)$ — known as the reflection coefficient matrices [510] — that minimize the the sum of the squared forward and backward prediction errors over the entire multichannel set of data points N given by

$$\epsilon_c^2 = Tr \left[\sum_{m=P_1+1}^{N-1} \left\{ \mathbf{e}_{P_1+1}^f(m) [\mathbf{e}_{P_1+1}^f(m)]^T + \mathbf{e}_{P_1+1}^b(m) [\mathbf{e}_{P_1+1}^b(m)]^T \right\} \right] \quad (11.138)$$

are obtained by solving [986]

$$\begin{aligned} [\boldsymbol{\sigma}_{P_1}^b]^{-1} [\boldsymbol{\sigma}_{P_1}^b]^{-1} \mathbf{E}_{P_1}^b \mathbf{a}_{P_1+1}(P_1 + 1) + \mathbf{a}_{P_1+1}(P_1 + 1) [\boldsymbol{\sigma}_{P_1}^f]^{-1} \mathbf{E}_{P_1}^f [\boldsymbol{\sigma}_{P_1}^f]^{-1} = \\ -[\boldsymbol{\sigma}_{P_1}^b]^{-1} [\boldsymbol{\sigma}_{P_1}^b]^{-1} \mathbf{E}_{P_1}^{fb} - [\boldsymbol{\sigma}_{P_1}^b]^{-1} \mathbf{E}_{P_1}^{fb} [\boldsymbol{\sigma}_{P_1}^f]^{-1}, \end{aligned} \quad (11.139)$$

where

$$\mathbf{E}_{P_1}^f = \sum_{m=P_1}^{N-1} \mathbf{e}_{P_1}^f(m) [\mathbf{e}_{P_1}^f(m)]^T, \quad (11.140)$$

$$\mathbf{E}_{P_1}^b = \sum_{m=P_1}^{N-1} \mathbf{e}_{P_1}^b(m) [\mathbf{e}_{P_1}^b(m)]^T, \quad (11.141)$$

and

$$\mathbf{E}_{P_1}^{fb} = \sum_{m=P_1}^{N-1} \mathbf{e}_{P_1}^f(m) [\mathbf{e}_{P_1}^b(m)]^T. \quad (11.142)$$

Equations 11.136, 11.137, and 11.139 – 11.142 may be used to compute the multichannel reflection coefficients directly from the image data without computing the ACF.

In order to adapt the multichannel version of the Burg algorithm to the 2D image case, we could force the structure obtained by relating the 2D

and multichannel ACFs in Equation 11.85 on to the expressions in Equations 11.137 and 11.139, and redefine the error covariance matrices $\mathbf{E}_{P_1}^f$, $\mathbf{E}_{P_1}^b$, and $\mathbf{E}_{P_1}^{fb}$ to span the entire $M \times N$ image. Then, the 2D counterpart of the reflection coefficient matrix $\mathbf{a}_{P_1+1}(P_1+1)$ is obtained by solving the following equation [338, 337]:

$$\begin{aligned} \sigma_{P_1}^{-1} \sigma_{P_1}^{-1} \mathbf{E}_{P_1}^b \mathbf{a}_{P_1+1}(P_1+1) + \mathbf{a}_{P_1+1}(P_1+1) \sigma_{P_1}^{-1} \mathbf{E}_{P_1}^f \sigma_{P_1}^{-1} = \\ -\sigma_{P_1}^{-1} \sigma_{P_1}^{-1} \mathbf{E}_{P_1}^{fb} - \sigma_{P_1}^{-1} \mathbf{E}_{P_1}^{fb} \sigma_{P_1}^{-1}. \end{aligned} \quad (11.143)$$

In order to compute the error covariance matrices $\mathbf{E}_{P_1}^f$, $\mathbf{E}_{P_1}^b$, and $\mathbf{E}_{P_1}^{fb}$, a strip of width (P_2+1) is defined so as to span the top (P_2+1) rows of the image, as shown in Figure 11.22, and the strip is moved down one row at a time. The region over which the summations are performed includes only those parts of the strip for which the forward and backward prediction operators do not run out of data. At the beginning of the recursive procedure, the error values are initialized to the actual values of the corresponding pixels. Furthermore, the forward and backward prediction error vectors are computed by forcing the relationship in Equation 11.124 on to Equation 11.137, resulting in

$$\mathbf{e}_{P_1+1}^b(m) = \mathbf{e}_{P_1}^b(m-1) + \mathbf{J} \mathbf{a}_{P_1+1}(P_1+1) \mathbf{J} \mathbf{e}_{P_1}^f(m). \quad (11.144)$$

The 2D Burg algorithm for computing the 2D LP coefficients directly from the image data may be summarized as follows:

1. The prediction error covariance matrix σ_0 is initialized to $\phi_f(0)$.
2. The prediction error vectors are computed using Equations 11.136 and 11.144.
3. The prediction error covariance matrices $\mathbf{E}_{P_1}^f$, $\mathbf{E}_{P_1}^b$, and $\mathbf{E}_{P_1}^{fb}$ are computed from the prediction error vectors using Equations 11.140 – 11.142 and summing over strips of width (P_2+1) rows of the image.
4. The reflection coefficient matrix $\mathbf{a}_{P_1+1}(P_1+1)$ is obtained from the prediction error covariance matrices by solving Equation 11.143, which is of the form $AX + XB = C$ and can be solved by using Kronecker products [986, 989].
5. The remaining prediction coefficient matrices $\mathbf{a}_{P_1+1}(p)$ are computed by using Equation 11.128, and the expected value of the prediction error covariance matrix σ_{P_1+1} is updated using Equation 11.129.
6. When the recursive procedure reaches the desired order P_1 , the 2D LP coefficients are computed by solving Equations 11.86 and 11.87.

TABLE 11.6

Variables in the 2D LP, Burg, and Levinson Algorithms for LP [338].

Variable	Size	Description
$f(m, n)$	$M \times N$	2D image array
$\mathbf{f}(m)$	$(P_2 + 1) \times 1$	Multichannel vector at column m spanning $P_2 + 1$ rows
$\phi(r)$	$(P_2 + 1) \times (P_2 + 1)$	Autocorrelation submatrix related to $\mathbf{f}(m)$
Φ	$(P_1 + 1)(P_2 + 1) \times (P_1 + 1)(P_2 + 1)$	Extended 2D autocorrelation matrix
\mathbf{a}	$(P_1 + 1)(P_2 + 1) \times 1$	2D LP coefficient matrix
$\mathbf{a}_{P_1}(p)$	$(P_2 + 1) \times (P_2 + 1)$	Multichannel-equivalent prediction coefficient matrix
$\mathbf{a}_{P_1}(P_1)$	$(P_2 + 1) \times (P_2 + 1)$	Multichannel-equivalent reflection coefficient matrix
σ_{P_1}	$(P_2 + 1) \times (P_2 + 1)$	Multichannel prediction error covariance matrix
$e^f(m, n)$	$M \times N$	Forward prediction error array
$e^b(m, n)$	$M \times N$	Backward prediction error array
$\mathbf{e}^f(m)$	$(P_2 + 1) \times 1$	Multichannel forward prediction error vector
$\mathbf{e}^b(m)$	$(P_2 + 1) \times 1$	Multichannel backward prediction error vector
$\mathbf{e}(m)(n)$	scalar	n^{th} element of the vector $\mathbf{e}(m)$
$\mathbf{E}_{P_1}^f$	$(P_2 + 1) \times (P_2 + 1)$	Forward prediction error covariance matrix
$\mathbf{E}_{P_1}^b$	$(P_2 + 1) \times (P_2 + 1)$	Backward prediction error covariance matrix
$\mathbf{E}_{P_1}^{fb}$	$(P_2 + 1) \times (P_2 + 1)$	Forward-backward prediction error covariance matrix

Bold characters represent vectors or matrices. A QP ROS of size $P_1 \times P_2$ is assumed.

The variables involved in the 2D Burg and Levinson algorithms are summarized in Table 11.6.

The modified multichannel version of the Burg algorithm offers advantages similar to those of its 1D counterpart, over the direct inversion method: it is a fast and efficient procedure to compute the prediction coefficients and prediction errors without computing the autocorrelation function. The optimization of the prediction coefficients does not make any assumptions about the image outside its finite dimensions, and hence, should result in lower prediction errors and efficient coding. Furthermore, the forced 2D structure makes the algorithm computationally more efficient than the direct application of the multichannel Burg procedure.

Computation of the prediction error: In order to compute the prediction error for coding and transmission, the trace of the covariance matrix in Equation 11.138 may be minimized, using Equations 11.136 and 11.144, and eliminating the covariance matrix σ_{P_1+1} , as follows. From Equation 11.138, the squared forward and backward prediction error vectors in the 2D Burg algorithm are given as [338]

$$\begin{aligned}
 \mathbf{E}_{P_1+1} &= \sum_{m=P_1+1}^{N-1} \left[\mathbf{e}_{P_1+1}^f(m) [\mathbf{e}_{P_1+1}^f(m)]^T + \mathbf{e}_{P_1+1}^b(m) [\mathbf{e}_{P_1+1}^b(m)]^T \right] \\
 &= \sum_{m=P_1+1}^{N-1} \left\{ \left[\mathbf{e}_{P_1}^f(m) + \mathbf{a}_{P_1+1}(P_1+1) \mathbf{e}_{P_1}^b(m-1) \right] \right. \\
 &\quad \times \left[\mathbf{e}_{P_1}^f(m) + \mathbf{a}_{P_1+1}(P_1+1) \mathbf{e}_{P_1}^b(m-1) \right]^T \\
 &\quad + \left[\mathbf{e}_{P_1}^b(m-1) + \mathbf{J} \mathbf{a}_{P_1+1}(P_1+1) \mathbf{J} \mathbf{e}_{P_1}^f(m) \right] \\
 &\quad \times \left. \left[\mathbf{e}_{P_1}^b(m-1) + \mathbf{J} \mathbf{a}_{P_1+1}(P_1+1) \mathbf{J} \mathbf{e}_{P_1}^f(m) \right]^T \right\} \\
 &= \mathbf{E}_{P_1}^f + \mathbf{E}_{P_1}^b + \mathbf{E}_{P_1}^{fb} \mathbf{a}_{P_1+1}^T(P_1+1) + \mathbf{a}_{P_1+1}(P_1+1) [\mathbf{E}_{P_1}^{fb}]^T \\
 &\quad + \mathbf{a}_{P_1+1}(P_1+1) \mathbf{E}_{P_1}^b \mathbf{a}_{P_1+1}^T(P_1+1) + [\mathbf{E}_{P_1}^{fb}]^T \mathbf{J} \mathbf{a}_{P_1+1}^T(P_1+1) \mathbf{J} \\
 &\quad + \mathbf{J} \mathbf{a}_{P_1+1}(P_1+1) \mathbf{J} \mathbf{E}_{P_1}^{fb} + \mathbf{J} \mathbf{a}_{P_1+1}(P_1+1) \mathbf{J} \mathbf{E}_{P_1}^f \mathbf{J} \mathbf{a}_{P_1+1}^T(P_1+1) \mathbf{J}.
 \end{aligned} \tag{11.145}$$

The 2D Burg algorithm for the purpose of image compression consists of determining the reflection coefficient matrix $\mathbf{a}_{P_1+1}(P_1+1)$ that minimizes the trace of the error covariance matrix \mathbf{E}_{P_1+1} . This is achieved by differentiating Equation 11.145 with respect to $\mathbf{a}_{P_1+1}(P_1+1)$ and equating the result to the null matrix, which yields

$$2 [\mathbf{E}_{P_1}^{fb}]^T + 2 \mathbf{E}_{P_1}^b \mathbf{a}_{P_1+1}(P_1+1) + 2 \mathbf{E}_{P_1}^{fb} + 2 \mathbf{J} \mathbf{a}_{P_1+1}(P_1+1) \mathbf{J} \mathbf{E}_{P_1}^f = \mathbf{0}, \tag{11.146}$$

or

$$\mathbf{J} \mathbf{E}_{P_1}^b \mathbf{a}_{P_1+1}(P_1+1) + \mathbf{a}_{P_1+1}(P_1+1) \mathbf{J} \mathbf{E}_{P_1}^f = -\mathbf{J} \left\{ \mathbf{E}_{P_1}^{fb} + [\mathbf{E}_{P_1}^{fb}]^T \right\}. \quad (11.147)$$

If the 2D autocorrelation matrices $\phi_f(r)$ are symmetric, the matrix $\mathbf{a}_{P_1+1}(P_1+1)$ will also be symmetric, which reduces Equation 11.147 to

$$\mathbf{E}_{P_1}^b \mathbf{a}_{P_1+1}(P_1+1) + \mathbf{a}_{P_1+1}(P_1+1) \mathbf{E}_{P_1}^f = -\left\{ \mathbf{E}_{P_1}^{fb} + [\mathbf{E}_{P_1}^{fb}]^T \right\} = -2 \mathbf{E}_{P_1}^{fb}. \quad (11.148)$$

When the recursive procedure reaches the desired order P_1 , the multichannel-equivalent 2D prediction error image is obtained as

$$e_0[m(P_2+1) + p, q] = \mathbf{e}_{P_1}^f(mN + q)(p), \quad (11.149)$$

$$p = 0, 1, 2, \dots, P_2; \quad q = 0, 1, 2, \dots, N-1; \quad m = 0, 1, 2, \dots, \frac{M}{P_2+1} - 1.$$

Equation 11.86 is in the form of the normal equations for 1D LP [510]. This suggests that the 1D Burg algorithm for LP [990] may be applied to the multichannel-equivalent 2D prediction error image to obtain the final prediction error image, in a recursive manner, as follows [338]:

1. Compute the sum of the squared forward and backward prediction errors as

$$\epsilon_f^2 = \sum_{m=P_2}^{N-1} \sum_{n=0}^{M-1} |e_{P_2}(m, n)|^2, \quad (11.150)$$

$$\epsilon_b^2 = \sum_{m=P_2}^{N-1} \sum_{n=0}^{M-1} |c_{P_2}(m, n)|^2, \quad (11.151)$$

and

$$\epsilon_{fb}^2 = \sum_{m=P_2}^{N-1} \sum_{n=0}^{M-1} e_{P_2}(m, n) c_{P_2}(m, n), \quad (11.152)$$

where $c_{P_2}(m, n)$ is the $M \times N$ backward prediction error array, initialized as

$$c_0(m, n) = e_0(m, n); \quad m = 0, 1, 2, \dots, M-1; \quad n = 0, 1, 2, \dots, N-1. \quad (11.153)$$

2. Compute the coefficient $a(0, P_2+1)$, known as the reflection coefficient, as

$$a(0, P_2+1) = -\frac{\epsilon_{fb}^2}{\epsilon_f^2 + \epsilon_b^2}. \quad (11.154)$$

3. Obtain the prediction errors at higher orders as

$$e_{P_2+1}(m, n) = e_{P_2}(m, n) + a(0, P_2 + 1) c_{P_2+1}(m, n - 1), \quad (11.155)$$

and

$$c_{P_2+1}(m, n) = a(0, P_2 + 1) e_{P_2+1}(m, n) + e_{P_2}(m, n). \quad (11.156)$$

When the desired order P_2 is reached, the prediction errors $e_{P_2}(m, n)$ are encoded using a method such as the Huffman code. The reflection coefficient matrices $\mathbf{a}_{P_1+1}(P_1 + 1)$ and the 1D reflection coefficients $a(0, P_2)$ are also encoded and transmitted as overhead information.

Error-free reconstruction of the image from the forward and backward prediction errors: In order to reconstruct the original image at the decoder without any error, the prediction coefficients need to be re-computed from the reflection coefficients. The prediction coefficients $a(0, p)$, $p = 0, 1, 2, \dots, P_2$, may be computed recursively using the Burg algorithm as

$$a(0, p) \Leftarrow a(0, p) + a(0, q + 1) a(0, q + 1 - p); \quad p = 1, 2, \dots, q; \quad q = 1, 2, \dots, P_2. \quad (11.157)$$

The multichannel-equivalent 2D prediction error image is given by

$$e_0(m, n) = \sum_{p=1}^{P_2} a(0, p) e_0(m - p, n) + e_{P_2+1}(m, n), \quad (11.158)$$

$$n = 0, 1, 2, \dots, N - 1; \quad m = P_2 + 1, P_2 + 2, \dots, M - 1.$$

The multichannel prediction error vectors are related to the error data defined above as

$$\mathbf{e}_{P_1}^f(mN + q)(p) = e_0[m(P_2 + 1) + p, q], \quad (11.159)$$

$$p = 0, 1, 2, \dots, P_2; \quad q = 0, 1, 2, \dots, N - 1; \quad m = 0, 1, 2, \dots, \frac{M}{P_2 + 1} - 1.$$

The multichannel signal vectors may be reconstructed from the error vectors via multichannel prediction as

$$\tilde{\mathbf{f}}(m) = - \sum_{p=1}^{P_1} \mathbf{a}(p) \mathbf{f}(m - p) + \mathbf{e}_{P_1}^f(m); \quad m = P_1 + 1, P_1 + 2, \dots, N - 1. \quad (11.160)$$

Finally, the original image is recovered from the multichannel signal vectors as

$$\mathbf{f}[m(P_2 + 1) + p, q] = \mathbf{f}(mN + q)(p), \quad (11.161)$$

$$p = 0, 1, 2, \dots, P_2; \quad q = 0, 1, 2, \dots, N - 1; \quad m = 0, 1, 2, \dots, \frac{M}{P_2 + 1} - 1,$$

with rounding of the results to integers. In a practical implementation, for values of $e_{P_2+1}(m, n)$ exceeding a preset limit, the true image pixel values would be transmitted and made available directly at the decoder.

Results of application to medical images: Kuduvali and Rangayyan [174, 337, 338] applied the 2D Levinson and Burg algorithms described above to the 10 high-resolution digitized medical images listed in Table 11.5. The average bit rate with lossless compression of the 10 test images using the 2D block-wise LP method described in Section 11.8.1, the 2D Levinson algorithm, and the 2D Burg algorithm were, respectively, 3.15, 3.02, and 2.81 *b/pixel*, with the original images having 10 *b/pixel* (see also Table 11.7). The multi-channel LP algorithms, in particular, the 2D Burg algorithm, provided better compression than the other methods described in the preceding sections in this chapter. The LP models described in this section are related to AR modeling for spectral estimation; Kuduvali and Rangayyan [337] found the 2D Burg algorithm to provide good 2D spectral estimates that were comparable to those provided by other AR models.

11.8.3 Adaptive 2D recursive least-squares prediction

The LP model with constant prediction coefficients given by Equation 11.53 is based on an inherent assumption of stationarity of the image-generating process. The multichannel-based prediction methods described in Section 11.8.2 are two-pass methods, where an estimation of the statistical parameters of the image is performed in the first pass (such as, for example, the autocorrelation matrix of the image in the 2D Levinson method), and the parameters are then used to estimate the prediction coefficients in the second pass. Once computed, the same prediction coefficients are used for prediction over all of the image data from which the coefficients were estimated. However, this assumption of stationarity is rarely valid in the case of natural images as well as biomedical images. To overcome this problem, in the case of the multichannel-based methods, the approach taken was that of partitioning the image into blocks, and computing the prediction coefficients independently for each block. Another possible approach is to adapt the coefficients recursively to the changing statistical characteristics of the image. In this section, the basis for such adaptive algorithms is described, and a 2D recursive least-squares (2D RLS) algorithm for adaptive computation of the LP coefficients is formulated [338]. The procedures are based upon adaptive filter theory in 1D [833, 979] and in multichannel signal filtering [991, 992, 993].

With reference to the basic 2D LP model given in Equation 11.53, several approaches are available for adaptive computation of the coefficients $a(p, q)$ for each pixel being predicted at the location (m, n) [833]. The approach based on Wiener filter theory [198] (see Section 3.6.1), leading to the 2D LMS algorithm [206, 994] (see Section 3.7.3), although applicable to image compression [995], suffers from the fact that the estimation of the coefficients $a(p, q)$ does not make use of all the image data available up to the current

location. Adaptive estimation of the coefficients based upon the Kalman filter [833, 887, 888, 891, 892, 893] (see Section 10.4.3), where the prediction coefficients are represented as the state vector describing the current state of the image-generating process, has not been explored much. However, this approach depends upon the statistics of the image represented in terms of ensemble averages; because only estimates of the ensemble averages can be obtained, this approach is likely to be suboptimal.

The approach that is described in this section for adaptive prediction, based upon the work of Kuduvalli [338], is founded upon the method of least squares. This approach is deterministic in its formulation, and involves the minimization of a weighted sum of prediction errors. In Section 11.8.2, it was observed that the estimation of the prediction coefficients based on the direct minimization of the actual prediction errors (the 2D Burg method) yielded better results in image compression than the method based on the estimation of an ensemble image statistic (the 2D ACF) from the image data (the 2D Levinson method). This result suggests that a deterministic approach could also be appropriate for the adaptive computation of prediction coefficients.

In 2D RLS prediction, the aim is to minimize a weighted sum of the squared prediction errors, computed up to the present location, given by

$$\epsilon^2(m, n) = \sum_{(p,q) \in ROS} \sum w(m, n, p, q) [e(p, q)]^2, \quad (11.162)$$

where $e(p, q)$ is the prediction error at (p, q) , and $w(m, n, p, q)$ is a weighting factor chosen to selectively “forget” the errors from the preceding pixel locations (“the past”) in order for the prediction coefficients to adapt to the changing statistical nature of the image at the current location. Boutalis et al. [992] used an exponential weighting factor whose magnitude reduces in the direction opposite to the scanning model used in the generation of the image. With this weighting-factor model and special ROSs, Boutalis et al. used the multichannel version of the RLS algorithm directly for adaptive estimation of images; however, their weighting-factor model does not take into account the 2D nature of images: the weight assigned to the error at a location adjacent in the row direction to the current location is higher than the weight assigned to the error at a location adjacent in the column direction. Kuduvalli [338] proposed a weighting-factor model that is truly 2D in its formulation. In this method, using a rectangular region spanning the image up to the current location for minimizing the sum of the prediction errors as shown in Figure 11.23, and an exponential weighting factor defined as $w(m, n, p, q) = \lambda^{(m-p+n-q)}$, where $0 < \lambda \leq 1$ is a forgetting factor, the weighted squared error is defined as

$$\epsilon^2(m, n) = \sum_{p=0}^m \sum_{q=0}^n \lambda^{(m-p+n-q)} [e(p, q)]^2. \quad (11.163)$$

Let us consider a QP ROS of order $P \times Q$ for prediction, as shown in Figure 11.23, and use the following notation for representing the prediction




FIGURE 11.23

ROSs in adaptive LP by the 2D RLS method [338]. While the image is scanned from the position $(m, n-1)$ to (m, n) , a column of m pixels becomes available as new information that may be used to update the forward and backward predictors. Observe that a part of the column of new information is hidden by the ROS for both forward and backward prediction in the figure.

coefficients and the image data spanning the current ROS as vectors:

$$\mathbf{a}(m, n) = [a_0^T(m, n) \ a_1^T(m, n) \ \cdots \ a_P^T(m, n)]^T, \quad (11.164)$$

where

$$a_p(m, n) = [a(m, n)(p, 0) \ a(m, n)(p, 1) \ \cdots \ a(m, n)(p, Q)]^T, \quad (11.165)$$

with $a(m, n)(0, 0) = 1$, and

$$\mathbf{F}_{P+1}(m, n) = [\mathbf{f}_m^T(n) \ \mathbf{f}_{m-1}^T(n) \ \cdots \ \mathbf{f}_{m-P}^T(n)]^T, \quad (11.166)$$

with

$$\mathbf{f}_{m-p}(n) = [f(m-p, n) \ f(m-p, n-1) \ \cdots \ f(m-p, n-Q)]^T. \quad (11.167)$$

Here, the subscripts P and $P+1$ represent the order (size) of the matrices and vectors, and the indices (m, n) indicate that the values of the parameters corresponding to the pixel location (m, n) . Observe that $\mathbf{a}(m, n)$ is a $P \times Q$ matrix, with $a(m, n)(p, q)$ representing its element at (p, q) . With this notation, the prediction error may be written as

$$e(m, n) = \mathbf{a}^T(m, n) \mathbf{F}_{P+1}(m, n). \quad (11.168)$$

The 2D RLS normal equations: The coefficients that minimize the weighted sum of the squared prediction errors $\epsilon^2(m, n)$ given in Equation 11.163 are obtained as the solution to the 2D RLS normal equations, which are obtained as follows [338]. Let us perform partitioning of the matrices $\mathbf{a}(m, n)$ and $\mathbf{F}_{P+1}(m, n)$ as

$$\mathbf{a}(m, n) = \begin{bmatrix} 1 \\ \tilde{\mathbf{a}}(m, n) \end{bmatrix} \quad (11.169)$$

and

$$\mathbf{F}_{P+1}(m, n) = \begin{bmatrix} f(m, n) \\ \tilde{\mathbf{F}}_{P+1}(m, n) \end{bmatrix} = \begin{bmatrix} \mathbf{F}_{P+1}^\#(m, n) \\ f(m-P, n-Q) \end{bmatrix}. \quad (11.170)$$

Observe that the coefficient matrix $\tilde{\mathbf{a}}(m, n)$ and the data matrix $\tilde{\mathbf{F}}_{P+1}(m, n)$ consist of all of the 2D RLS coefficients $a(m, n)(p, q)$ and all of the image pixels $f(m-p, n-q)$ such that $(p, q) \in \text{QP ROS}$ for a forward predictor. With partitioning as above, the prediction error in Equation 11.168 may be written as

$$e(m, n) = f(m, n) + \tilde{\mathbf{a}}^T(m, n) \tilde{\mathbf{F}}_{P+1}(m, n). \quad (11.171)$$

The sum of the squared prediction errors in Equation 11.163 may now be expressed as

$$\epsilon^2(m, n) = \sum_{p=0}^m \sum_{q=0}^n \lambda^{(m-p+n-q)} [e(p, q)]^2,$$

$$\begin{aligned}
 &= \sum_{p=0}^m \sum_{q=0}^n \lambda^{(m-p+n-q)} \left[f(p, q) + \tilde{\mathbf{a}}^T(m, n) \tilde{\mathbf{F}}_{P+1}(p, q) \right] \\
 &\times \left[f(p, q) + \tilde{\mathbf{a}}^T(m, n) \tilde{\mathbf{F}}_{P+1}(p, q) \right]^T \\
 &= \sum_{p=0}^m \sum_{q=0}^n \lambda^{(m-p+n-q)} \left[f^2(p, q) + 2 \tilde{\mathbf{a}}^T(m, n) \tilde{\mathbf{F}}_{P+1}(p, q) f(p, q) \right. \\
 &\left. + \tilde{\mathbf{a}}^T(m, n) \tilde{\mathbf{F}}_{P+1}(p, q) \tilde{\mathbf{F}}_{P+1}^T(p, q) \tilde{\mathbf{a}}(m, n) \right]. \tag{11.172}
 \end{aligned}$$

In order to determine the coefficients $a(m, n)(p, q)$ that minimize $\epsilon^2(m, n)$, we could differentiate the expression above for $\epsilon^2(m, n)$ with respect to the coefficient matrix $\tilde{\mathbf{a}}(m, n)$ and equate the result to the null matrix of size $[(P + 1)(Q + 1) - 1] \times 1$, which yields

$$\begin{aligned}
 \mathbf{0} &= \frac{\partial \epsilon^2(m, n)}{\partial \tilde{\mathbf{a}}(m, n)} \\
 &= \sum_{p=0}^m \sum_{q=0}^n \lambda^{(m-p+n-q)} \left[\tilde{\mathbf{F}}_{P+1}(p, q) f(p, q) \right. \\
 &\left. + \tilde{\mathbf{F}}_{P+1}(p, q) \tilde{\mathbf{F}}_{P+1}^T(p, q) \tilde{\mathbf{a}}(m, n) \right]. \tag{11.173}
 \end{aligned}$$

Equation 11.173 may be expressed in matrix notation as

$$\sum_{p=0}^m \sum_{q=0}^n \lambda^{(m-p+n-q)} \tilde{\mathbf{F}}_{P+1}(p, q) \left[f(p, q) \quad \tilde{\mathbf{F}}_{P+1}^T(p, q) \right] \begin{bmatrix} \mathbf{1} \\ \tilde{\mathbf{a}}(m, n) \end{bmatrix} = \mathbf{0}. \tag{11.174}$$

In addition to the above, using Equation 11.173 in Equation 11.172, we have

$$\epsilon^2(m, n) = \sum_{p=0}^m \sum_{q=0}^n \lambda^{(m-p+n-q)} \left[f^2(p, q) + f(p, q) \tilde{\mathbf{F}}_{P+1}^T(p, q) \tilde{\mathbf{a}}(m, n) \right], \tag{11.175}$$

which may be written in matrix form as

$$\epsilon^2(m, n) = \sum_{p=0}^m \sum_{q=0}^n \lambda^{(m-p+n-q)} f(p, q) \left[f(p, q) \quad \tilde{\mathbf{F}}_{P+1}^T(p, q) \right] \begin{bmatrix} \mathbf{1} \\ \tilde{\mathbf{a}}(m, n) \end{bmatrix}. \tag{11.176}$$

Combining Equations 11.174 and 11.176, we get

$$\begin{aligned}
 &\sum_{p=0}^m \sum_{q=0}^n \lambda^{(m-p+n-q)} \begin{bmatrix} f(p, q) \\ \tilde{\mathbf{F}}_{P+1}(p, q) \end{bmatrix} \left[f(p, q) \quad \tilde{\mathbf{F}}_{P+1}^T(p, q) \right] \\
 &\times \begin{bmatrix} \mathbf{1} \\ \tilde{\mathbf{a}}(m, n) \end{bmatrix} = \begin{bmatrix} \epsilon^2(m, n) \\ \mathbf{0} \end{bmatrix}, \tag{11.177}
 \end{aligned}$$

or

$$\sum_{p=0}^m \sum_{q=0}^n \lambda^{(m-p+n-q)} \mathbf{F}_{P+1}(p, q) \mathbf{F}_{P+1}^T(p, q) \mathbf{a}(m, n) = \begin{bmatrix} \epsilon^2(m, n) \\ \mathbf{0} \end{bmatrix}, \quad (11.178)$$

which may be expressed as

$$\Phi_{P+1}(m, n) \mathbf{a}(m, n) = \boldsymbol{\rho}(m, n), \quad (11.179)$$

where

$$\boldsymbol{\rho}(m, n) = [\epsilon^2(m, n) \ 0 \ 0 \ \cdots \ 0]^T, \quad (11.180)$$

and $\Phi_{P+1}(m, n)$ is the deterministic autocorrelation matrix of the weighted image given by

$$\Phi_{P+1}(m, n) = \sum_{p=0}^m \sum_{q=0}^n \lambda^{(m-p+n-q)} \mathbf{F}_{P+1}(p, q) \mathbf{F}_{P+1}^T(p, q). \quad (11.181)$$

Equation 11.179 represents the 2D RLS normal equations, solving which we can obtain the prediction coefficients $a(m, n)(p, q)$ that adapt to the statistics of the image at the location (m, n) .

Solving the 2D RLS normal equations: Direct inversion of the autocorrelation matrix in Equation 11.179 gives the desired matrix of prediction coefficients as

$$\mathbf{a}(m, n) = \Phi_{P+1}^{-1}(m, n) \boldsymbol{\rho}(m, n). \quad (11.182)$$

The matrix $\Phi_{P+1}(m, n)$ is of size $(P+1)(Q+1) \times (P+1)(Q+1)$; the inversion of such a matrix at every pixel (m, n) of the image could be computationally intensive. Kuduvalli [338] developed the following procedure to reduce the size of the matrix to be inverted to $(Q+1) \times (Q+1)$. The procedure starts with a recursive relationship expressing the solution for the normal equations at the pixel location (m, n) in terms of that at $(m-1, n)$. Consider the expression

$$\begin{aligned} \Phi_{P+1}(m, n) &= \sum_{p=0}^m \sum_{q=0}^n \lambda^{(m-p+n-q)} \mathbf{F}_{P+1}(p, q) \mathbf{F}_{P+1}^T(p, q) \\ &= \begin{bmatrix} \phi_{00}(m, n) & \phi_{01}^T(m, n) & \cdots & \phi_{0P}^T(m, n) \\ \phi_{01}(m, n) & \phi_{11}(m, n) & \cdots & \phi_{1P}^T(m, n) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_{0P}(m, n) & \phi_{1P}(m, n) & \cdots & \phi_{PP}(m, n) \end{bmatrix}, \end{aligned} \quad (11.183)$$

where

$$\phi_{rs}(m, n) = \sum_{p=0}^m \sum_{q=0}^n \lambda^{(m-p+n-q)} \mathbf{f}_{p-r}(q) \mathbf{f}_{p-s}^T(q). \quad (11.184)$$

Observe that

$$\phi_{rs}(m, n) = \phi_{0(s-r)}(m - r, s), \tag{11.185}$$

which follows from the assumption that the image data have been windowed such that $f(m, n) = 0$ for $m < 0$ or $n < 0$.

The normal equations may now be expressed as

$$\begin{bmatrix} \phi_{00}(m, n) & \phi_{01}^T(m, n) & \cdots & \phi_{0P}^T(m, n) \\ \phi_{01}(m, n) & \phi_{11}(m, n) & \cdots & \phi_{1P}^T(m, n) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_{0P}(m, n) & \phi_{1P}(m, n) & \cdots & \phi_{PP}(m, n) \end{bmatrix} \begin{bmatrix} a_0(m, n) \\ a_1(m, n) \\ \vdots \\ a_P(m, n) \end{bmatrix} = \begin{bmatrix} \rho(m, n) \\ \mathbf{0}_{Q+1} \\ \vdots \\ \mathbf{0}_{Q+1} \end{bmatrix}, \tag{11.186}$$

where $\mathbf{0}_{Q+1}$ is the null matrix of size $(Q + 1) \times 1$.

Equation 11.186 may be solved in two steps: First, solve

$$\begin{bmatrix} \phi_{00}(m, n) & \phi_{01}^T(m, n) & \cdots & \phi_{0P}^T(m, n) \\ \phi_{01}(m, n) & \phi_{11}(m, n) & \cdots & \phi_{1P}^T(m, n) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_{0P}(m, n) & \phi_{1P}(m, n) & \cdots & \phi_{PP}(m, n) \end{bmatrix} \begin{bmatrix} \mathbf{I}_{Q+1} \\ A_1(m, n) \\ \vdots \\ A_P(m, n) \end{bmatrix} = \begin{bmatrix} \mathbf{F}_{Q+1}(m, n) \\ \mathbf{0}_{Q+1} \\ \vdots \\ \mathbf{0}_{Q+1} \end{bmatrix}, \tag{11.187}$$

for the $(Q + 1) \times (Q + 1)$ matrices $A_p(m, n)$, $p = 1, 2, \dots, P$, and $\mathbf{F}_{Q+1}(m, n)$. Here, \mathbf{I}_{Q+1} is the identity matrix of size $(Q + 1) \times (Q + 1)$, and $\mathbf{0}_{Q+1}$ is the null matrix of size $(Q + 1) \times (Q + 1)$. Then, obtain the solution to Equation 11.186 by solving

$$\mathbf{F}_{Q+1}(m, n) a_0(m, n) = \rho(m, n), \tag{11.188}$$

and using the relationship

$$A_p(m, n) a_0(m, n) = a_p(m, n); \quad p = 1, 2, \dots, P. \tag{11.189}$$

This approach is similar to the approach taken to solve the 2D Yule–Walker equations by the 2D Levinson method, described in Section 11.8.2, and leads to a recursive algorithm that is computationally efficient; the details of the algorithm are given by Kuduvalli [338].

Results of application to medical images: Kuduvalli [338] conducted preliminary studies on the application of the 2D RLS algorithm to predictive coding and compression of medical images. In the application to coding, the value of the pixel at the current location (m, n) is not available at the decoder before the prediction coefficient matrix $\mathbf{a}(m, n)$ is computed. However, the prediction coefficient matrix $\mathbf{a}(m - 1, n)$ is available. Thus, for error-free decoding, the a priori prediction error computed using the prediction coefficient matrix $\mathbf{a}(m - 1, n)$ is encoded. These error values, which have a PDF that is close to a Laplacian PDF, may be efficiently encoded using methods such as the Huffman code. Using a QP ROS of size 4×4 and a forgetting factor of $\lambda = 0.95$, Kuduvalli [338] obtained an average bit rate of 2.71 *b/pixel* for two of the images listed in Table 11.5; this rate, however, is only marginally

lower than the bit rate of 2.77 b/pixel for the same two images obtained by using the 2D Burg algorithm described in Section 11.8.2. Although the 2D RLS algorithm has the elegance of being a truly 2D algorithm that adapts to the changing statistics of the image on a pixel-by-pixel basis, the method did not yield appreciable advantages in image data compression. Regardless, the method has applications in other areas, such as spectrum estimation and filtering.

Kuduvalli and Rangayyan [174] performed a comparative analysis of several image compression techniques, including direct source coding, transform coding, interpolative coding, and predictive coding, applied to the high-resolution digitized medical images listed in Table 11.5. The average bit rates obtained using several coding and compression techniques are listed in Table 11.7. It should be observed that decorrelation can provide significant advantages over direct source encoding of the original pixel data. The adaptive predictive coding techniques have performed better than the transform and interpolative coding techniques tested.

In a study of the effect of sampling resolution on image data compression, Kuduvalli and Rangayyan [174] prepared low-resolution versions of the images listed in Table 11.5 by smoothing and downsampling. The results of the application of the 2D Levinson predictive coding algorithm yielded average bit rates of $3.5 - 5 \text{ b/pixel}$ for 512×512 images, and $2.5 - 3 \text{ b/pixel}$ for $4,096 \times 4,096$ images (with the original images at 10 b/pixel). This result indicates that high-resolution images possess more redundancy, and hence may be compressed by larger extents than their low-resolution counterparts. Therefore, increasing the resolution of medical images does not increase the amount of the related compressed data in direct proportion to the increase in matrix size, but by a lower factor. This result could be a motivating factor supporting the use of high resolution in medical imaging, without undue concerns related to significant increases in data-handling requirements.

See Aiazzi et al. [996] for a description of other methods for adaptive prediction and a comparative analysis of several methods for lossless image data compression.

11.9 Image Scanning Using the Peano-Hilbert Curve

Peano scanning is a method of scanning an image by following the path described by a space-filling curve [997, 998, 999, 1000, 1001, 1002, 1003, 1004]. Giuseppe Peano, an Italian mathematician, described the first space-filling curve in an attempt to map a line into a 2D space [997]. The term “Peano scanning” is used to refer to such a scanning scheme irrespective of the space-filling curve used to define the scan path. Peano’s curve was modified by

TABLE 11.7
 Average Bit Rates Obtained in
 the Lossless Compression of
 the Medical Images Listed in
 Table 11.5 Using Several Image
 Coding Techniques
 [174, 337, 338].

Coding method	Bits/ pixel
Original	10.00
Entropy H_0	7.94
Huffman	8.67
Arithmetic	8.58
LZW	5.57
DCT	5.26
Interpolative	3.45
2D LP	3.15
2D Levinson	3.02
2D Burg	2.81
2D RLS*	2.71

The Huffman code was used to encode the results of the transform, interpolative, and predictive coding methods. *Only two images were compressed with the 2D RLS method.

Hilbert [998], and the modified curve came to be known as the “Peano-Hilbert” curve.

Moore [1005] studied the geometric and analytical interpretation of continuous space-filling curves. Space-filling curves have aided in the development of fractals [462] (see Section 7.5 for a discussion on fractals). The Peano-Hilbert curve has been applied to display continuous-tone images [1006] in order to eliminate deficiencies of the ordered dithered technique, such as Moiré fringes. Lempel and Ziv [1003] used the Peano-Hilbert curve to scan images and define the lowest bound of compressibility. Zhang et al. [1001, 1002] explored the statistical characteristics of medical images using Peano scanning.

Provine and Rangayyan [1000, 999] studied the application of Peano scanning for image data compression, with an additional step of decorrelation using differentiation, orthogonal transforms, or LP; the following paragraphs describe the basics of the methods involved and the results obtained in their work.

11.9.1 Definition of the Peano-scan path

If a physical scanner that can scan an image by following the Peano curve is not available, Peano scanning may be simulated by selecting pixels from a raster-scanned image by traversing the 2D data along the path described by the Peano-Hilbert curve. The reordered pixel data so obtained (in a 1D stream) may be subjected to decorrelation and encoding operations as desired. An inverse Peano-scanning operation would be required at the receiving end to reconstruct the original image. A general image compression scheme as above is summarized in Figure 11.24.

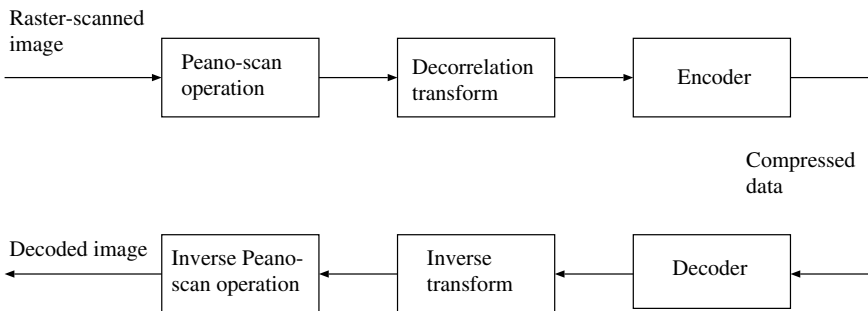


FIGURE 11.24

Image compression using Peano scanning.

The Peano-scan operation is recursive in nature, and spans a 2D space encountering a total of $2^i \times 2^i$ points (where i is an integer and $i > 1$). From the

perspective of processing a 2D array containing the pixel values of an image, this would require that the dimensions of the array be an integral power of 2. In the scanning procedure, the given image of size $2^n \times 2^n$ is divided into four quadrants, each of them forming a subimage; see Figure 11.25. Each of the subimages is further divided into four quadrants, and the procedure continues. The original image is divided into a total of $T_i = 2^{2(n-i+1)}$ subimages, each of size $2^{i-1} \times 2^{i-1}$, where $i = 1, 2, \dots, n$. In the following discussion, the subimages formed by the recursive subdivision procedure as above will be referred to as $s_{i-1}(k)$, $k = 1, 2, \dots, T_i$, where k increases along the direction of the scan path; the entire image will be referred to as s_n . Thus, each of the four quadrants formed by partitioning a subimage s_i , for any i , is of size $2^{i-1} \times 2^{i-1}$. The division of a given image into subimages is shown in Figure 11.25: the recursive division of subimages is performed until the s_1 subimages are formed. The four pixels within the smallest 2×2 subimage are denoted as p_1, p_2, p_3 , and p_4 respectively, in the order of being scanned. As the scan path builds recursively, the path definition is based on the basic definitions for a 2×2 subimage, as well as the recursive definitions for subimages of larger size, until the entire image is scanned. The four basic definitions of the Peano-scanning operation are given in Figure 11.26.

The recursive definitions of the Peano-scanning operation are given in Figure 11.27, which inherently use the basic definitions (shown in Figure 11.26) to obtain further pixels from the subimages. The definitions go down recursively from $i = n$ to $i = 1$, that is, from the image s_n down to the subimages s_1 ; the basic definitions are used to obtain the pixels from the s_1 subimages.

The scan-path definition for an image or a subimage depends on i . At a higher level, that is, for an s_i , $i > 1$, the recursive definition is as follows:

- If i is odd, the recursive definition is “R” (see Figure 11.27).
- If i is even, the recursive definition is “D”.

From the recursive definitions shown in Figure 11.27, the definitions of the scan pattern in each of the subimages is obtained as follows:

- If $s_i(k)$ has the recursive definition “R”, $s_{i-1}(4k-3)$ will follow the path given by “D”; $s_{i-1}(4k-2)$ and $s_{i-1}(4k-1)$ the path “R”; and $s_{i-1}(4k)$ the path “U”.
- If $s_i(k)$ has the recursive definition “D”, $s_{i-1}(4k-3)$ will follow the path given by “R”; $s_{i-1}(4k-2)$ and $s_{i-1}(4k-1)$ the path “D”; and $s_{i-1}(4k)$ the path “L”.
- If $s_i(k)$ has the recursive definition “L”, $s_{i-1}(4k-3)$ will follow the path given by “U”; $s_{i-1}(4k-2)$ and $s_{i-1}(4k-1)$ the path “L”; and $s_{i-1}(4k)$ the path “D”.
- If $s_i(k)$ has the recursive definition “U”, $s_{i-1}(4k-3)$ will follow the path given by “L”; $s_{i-1}(4k-2)$ and $s_{i-1}(4k-1)$ the path “U”; and $s_{i-1}(4k)$ the path “R”.




FIGURE 11.25

Division of a 16×16 image into subimages during Peano scanning. Reproduced with permission from J.A. Provine and R.M. Rangayyan, "Lossless compression of Peanoscaned images", *Journal of Electronic Imaging*, 3(2): 176 - 181, 1994. © SPIE and IS&T.




FIGURE 11.26

Basic definitions of the Peano-scanning operation. The points marked $p1 - p4$ represent the four pixels in a 2×2 subimage. Each scan pattern shown visits four pixels in the order indicated by the arrows. R: right. L: left. D: down. U: up. Reproduced with permission from J.A. Provine and R.M. Rangayyan, "Lossless compression of Peanoscanned images", *Journal of Electronic Imaging*, 3(2): 176 - 181, 1994. © SPIE and IS&T.




FIGURE 11.27

Recursive definitions of the Peano-scanning operation. Reproduced with permission from J.A. Provine and R.M. Rangayyan, "Lossless compression of Peanoscanned images", *Journal of Electronic Imaging*, 3(2): 176 - 181, 1994. © SPIE and IS&T.

The index k can take any value in the range $1, 2, \dots, T_i$ for $i = 1, 2, \dots, n$. From the recursive definitions, it is evident that k cannot continuously increase horizontally or vertically, due to the nature of the scan. Furthermore, because the scan path takes its course recursively, except for $i = n$, the recursive definitions “L” and “U” (see Figure 11.27) are also possible for lower values of i . All the subimages are divided and defined recursively until all the subimages s_1 are defined. The basic definitions of the Peano scan are then followed for each of the s_1 subimages. The Peano-scan pattern for a 16×16 image is illustrated in Figure 11.28, where the heavy dots indicate the positions of the first 16 pixels scanned. Understanding the scan pattern is facilitated by viewing Figure 11.28 along with Figure 11.25.




FIGURE 11.28

Peano-scan pattern for a 16×16 image. The positions of the pixels on the scan pattern are shown by heavy dots in the first 4×4 subimage. Reproduced with permission from J.A. Provine and R.M. Rangayyan, “Lossless compression of Peanoscanned images”, *Journal of Electronic Imaging*, 3(2): 176 – 181, 1994. © SPIE and IS&T.

Implementation of Peano scanning in software could use the recursive nature of the scan path efficiently to obtain the pixel stream. Recursive functions may call themselves within their body as the image is divided progressively into subimages until the s_1 subimages are formed, and the pixels are obtained recursively as the function builds back from the s_1 subimages to the full image s_n . Thus, the 2D image is unwrapped into a 1D data stream by following a continuous scan path.

The inverse Peano-scanning operation accomplishes the task of filling up the 2D array with the 1D data stream. This operation corresponds to the original works of Peano [997] and Hilbert [998], where the continuous mapping of a straight line into a 2D plane was described. Because the Peano-scanning operation is reversible, no loss of information is incurred.

11.9.2 Properties of the Peano-Hilbert curve

The Peano-Hilbert curve has several interesting and useful properties. The curve is continuous but not differentiable: it does not have a tangent at any point. Moore [1005] gave an explanation of the Peano-Hilbert curve adhering to this property. This property motivated the development of several other curves with the same property, which are used in the domain of fractals.

The Peano-Hilbert curve fills the 2D space continuously without passing through any point more than once. This feature enables the mapping of a 2D array into a 1D data stream. The recursive nature of the curve is useful in efficient implementation of the path of the curve. These two properties aid in scanning an image recursively quadrant by quadrant, leaving each quadrant only after having obtained every pixel within that quadrant, with each pixel visited only once in the process; see Figures 11.28 and 11.25. Preservation of the local 2D context in the scanning path could be expected to increase the correlation between successive elements in the 1D data stream. This aspect could facilitate improved image data compression.

Two other aspects of the Peano-Hilbert curve have proven to be useful in the bilevel display of continuous-tone images [1006, 1007]. Linearizing a 2D array along the path described by the Peano-Hilbert curve reduces the error between the sum of the bilevel values and the sum of the continuous-tone values of the original image because 2D locality is maintained by the scan path, unlike the 1D vector formed by concatenating the horizontal raster-scan lines of the image. The problem of long sections of scan lines running adjacent to one another is eliminated by following the Peano-scan path instead of the raster scan. Thus, Moiré patterns can be eliminated in regions of uniform intensity when presenting gray-level images on a bilevel display.

11.9.3 Implementation of Peano scanning

A practical problem that could arise in implementing the Peano-scanning operation on large images is the difficulty in allocating memory for the long linear

array used within the body of the recursive function for storing the scanned data. Provine [999] suggested the following approach to address this problem, by using a symmetrical pattern exhibited by the Peano-Hilbert curve.


The Peano-Hilbert curve exhibits a symmetrical pattern, which may be described as follows: For any subimage $s_i(k)$, the scan paths for the subimages $s_{i-1}(4k-3)$ and $s_{i-1}(4k-2)$ are the mirror reflections of the paths for the subimages $s_{i-1}(4k)$ and $s_{i-1}(4k-1)$, respectively. Two types of symmetry exist in the Peano-scan paths for a $2^i \times 2^i$ subimage depending on whether i is odd or even: If i is odd, the pattern for the upper half of the 2D space is reflected in the lower half; if i is even, the pattern in the left-hand half of the 2D space is reflected in the right-hand half; see Figure 11.29.

A symmetrical scan pattern exists for any subimage formed by the recursive division process. Hence, for the smallest subimage, the symmetrical pattern suggests that the basic definitions effectively obtain only two pixels, one after the other, either horizontally or vertically. In other words, the basic scan pattern effectively obtains only two pixels, p_1 and p_2 , out of the four pixels in a 2×2 subimage; see Figure 11.30. The manner in which the remaining two pixels p_3 and p_4 are obtained follows the symmetry property stated above (substituting $i = 1$ and $k = 1$). The sequence in which the two pixels p_3 and p_4 are obtained is shown in Figure 11.30 in dashed lines.


In scanning large images, for any subimage $s_i(k)$, the pattern of the Peano-scan path from the first pixel of $s_{i-1}(4k-3)$ to the last pixel of $s_{i-1}(4k-2)$ is the same as that from the last pixel of $s_{i-1}(4k)$ to the first pixel of $s_{i-1}(4k-1)$. Because the Peano-scan path does not leave any quadrant without visiting all the pixels within the quadrant, two equal sections of the image can be unwrapped independently, without affecting each other, into individual linear arrays by following the same scan path but in opposite directions. The resulting linear arrays, when concatenated appropriately, give the required 1D data stream. For the case illustrated in Figure 11.29 (a), the 64 pixels can be obtained as a linear sequence by tracing the Peano-scan path on pixels 1–32 in the upper half, followed by the pixels 32–1 in the lower half. Thus, several 1D arrays of a reasonable size may be used to hold the pixels obtained from different sections of the image. After all the subimages have been scanned (in parallel, if desired), the arrays may be concatenated appropriately to form the long sequence containing the entire image data.

11.9.4 Decorrelation of Peano-scanned data

The Peano-scanning operation scans the given picture recursively, quadrant by quadrant. Therefore, we could expect the 2D local statistics to be preserved in the resulting 1D data stream. Furthermore, we could also expect a higher correlation between pixels for larger lags in the Peano-scanned data than between the pixels obtained by concatenating the raster-scanned lines into a 1D array of pixels.



(a)




B

(b)

FIGURE 11.29

Symmetrical patterns exhibited by the Peano-Hilbert curve for (a) an 8×8 subimage and (b) a 16×16 subimage. The line AB indicates the axis of symmetry in each case. The 64 pixels in the image in (a) are labeled in the order of being scanned [the pixels 1 – 32 in the upper half, followed by the pixels 32 – 1 in the lower half of the subimage in (a)]. Figure courtesy of J.A. Provine [999].

**FIGURE 11.30**


Symmetrical patterns in the basic definitions of the Peano scan. Figure courtesy of J.A. Provine [999].

Figure 11.31 shows the ACFs obtained for the Lenna test image and a mamrogram for raster-scanned and Peano-scanned pixel streams. As expected, Peano scanning has maintained higher inter-pixel correlation than raster scanning. Similar observations have been made by Zhang et al. [1001, 1002] in their study on the stochastic properties of medical images and data compression with Peano scanning.


The simplest method for decorrelating pixel data is to produce a data sequence containing the differences between successive pixels. As shown earlier in Section 11.5 and Figure 11.12, the differentiated data may be expected to have a Laplacian PDF, which is useful in compressing the data. Provine and Rangayyan [999, 1000] applied a simple first-difference operation to decorrelate Peano-scanned image data, and encoded the resulting values using the Huffman, arithmetic, and LZW coding schemes. In addition, they applied the 1D DCT and LP modeling procedures to raster-scanned and Peano-scanned data streams, as well as the 2D DCT and LP modeling procedures to the original image data. Some of the results obtained by Provine and Rangayyan are summarized in Tables 11.8, 11.9, and 11.10. The application of either the Huffman or the arithmetic code to the differentiated Peano-scanned data stream resulted in the lowest average bit rate in the study.

11.10 Image Coding and Compression Standards

Two highly recognized international standards for the compression of still images are the Joint Bi-level Image experts Group (JBIG) standard [1008, 1009, 1010, 1011] and the Joint Photographic Experts Group (JPEG) standard [1012, 1011]. JBIG and JPEG are sanctioned by the International Organization for Standardization (ISO) and the Comité Consultatif International Téléphonique et Télégraphique (CCITT). Although JBIG was initially proposed for bilevel image compression, it may also be applied to continuous-



(a)



(b)

FIGURE 11.31

ACF of raster-scanned and Peano-scanned pixels plotted as a function of the distance (lag) between the scanned pixels: (a) for the Lenna image and (b) for a mammogram. Figure courtesy of J.A. Provine [999].

TABLE 11.8

Average Bit Rate with the Application of the Huffman, Arithmetic, and LZW Coding Schemes to Raster-scanned and Peano-scanned Data Obtained from Eight Test Images [999].

Image (8 <i>b/pixel</i>)	Size (pixels)	Entropy	Average number of bits/pixel			
		H_0 (bits)	LZW			
			Huffman	Arith.	Raster	Peano
Airplane	512×512	6.49	6.84	6.65	7.47	9.06
Baboon	512×512	7.14	9.40	7.30	9.64	9.63
Cameraman	256×256	7.04	7.39	7.40	8.99	8.09
Lenna-256	256×256	7.57	8.12	7.95	9.10	8.97
Lenna-512	512×512	7.45	10.38	7.61	9.00	8.85
Peppers	512×512	7.37	10.89	7.54	7.94	8.06
Sailboat	512×512	7.27	9.35	7.43	8.69	9.56
Tiffany	512×512	6.38	7.74	6.54	7.51	9.51
Mean	—	7.09	8.76	7.30	8.54	8.97
SD	—	0.44	1.46	0.48	0.8	0.62

See also Tables 11.9 and 11.10. *Note:* Arith. = arithmetic coding.

tone images by treating bit planes as independent bilevel images. (*Note:* The term “continuous-tone” images is used to represent gray-level images, color images, and multicomponent images; whereas some authors use the term “m-ary” for the same purpose, the former is preferred as it is used by JPEG.) The efficiency of such an application depends upon preprocessing for bit-plane decorrelation. The Moving Picture Experts Group (MPEG) standard [1013, 1014] applies to the compression of video images. In the context of medical image data handling and PACS, the ACR and the US National Electrical Manufacturers Association (NEMA) proposed standards known as the ACR/ NEMA and DICOM (Digital Imaging and Communications in Medicine) standards [1015, 1016, 1017, 1018]. The following sections provide brief reviews of the standards mentioned above.

TABLE 11.9

Average Bit Rate with Differentiated Peano-scanned Data (PD), Compared with the Results of 1D and 2D DPCM Encoding of Raster-scanned Data from Eight Test Images [999].

Image	Average number of bits/pixel								
	Huffman			Arithmetic			LZW		
	PD	DPCM		PD	DPCM		PD	DPCM	
		1D	2D		1D	2D		1D	2D
Airplane	4.52	5.49	4.60	4.48	5.50	4.58	5.13	5.33	5.09
Baboon	6.46	7.14	7.34	6.46	7.29	7.39	7.67	7.48	7.66
Cameraman	5.38	5.82	5.57	5.41	5.81	5.56	5.88	5.97	5.94
Lenna-256	5.66	6.35	5.74	5.64	6.40	5.71	6.24	6.60	6.02
Lenna-512	5.05	5.93	5.37	5.00	5.82	5.29	5.80	6.43	5.82
Peppers	5.06	5.82	6.20	4.97	5.73	6.09	5.63	5.72	5.88
Sailboat	5.55	6.34	6.61	5.54	6.39	6.59	6.50	6.51	6.65
Tiffany	4.63	5.59	5.02	4.60	5.52	4.91	5.26	5.83	5.52
Mean	5.29	6.06	5.81	5.26	6.06	5.77	6.01	6.23	6.07
SD	0.62	0.54	0.89	0.64	0.61	0.91	0.81	0.67	0.78

See also Tables 11.8 and 11.10.

11.10.1 The JBIG Standard

JBIG is a standard for progressive coding of bilevel images that supports three coding modes: progressive coding, compatible progressive/ sequential coding, and single-layer coding. A review of the single-layer coding mode is presented in the following paragraphs.

For a bilevel image $b(m, n)$, $0 \leq m < M$, $0 \leq n < N$, a typical JBIG coding scheme includes four main functional blocks as shown in Figure 11.32. The following items form important steps in the JBIG coding procedure:

- The *typical prediction* step is a line-skipping algorithm. A given line is marked as “typical” if it is identical to the preceding line. The encoder adds a special label to the encoded data stream for each typical line instead of encoding the line. The decoder generates the pixels of typical lines by line duplication.

TABLE 11.10

Average bit rates for eight test images with several decorrelation and encoding methods [999, 1000].

Image	Average bit rate (bits/pixel)				
	Direct arith. (Peano or raster)	PD (Huffman) encoding	PD (arith.) encoding	2D linear predictive coding (raster)	2D DCT coding (raster)
Airplane	6.65	4.52	4.48	4.50	5.94
Baboon	7.30	6.46	6.46	6.59	9.47
Cameraman	7.40	5.38	5.41	5.48	8.42
Lenna-256	7.95	5.66	5.64	5.49	8.00
Lenna-512	7.61	5.05	5.00	4.92	6.65
Peppers	7.54	5.06	4.97	5.30	7.12
Sailboat	7.43	5.55	5.54	5.72	7.70
Tiffany	6.54	4.63	4.60	4.76	6.55
Mean	7.30	5.29	5.26	5.35	7.48
SD	0.48	0.62	0.64	0.65	1.15

See also Tables 11.8 and 11.9. *Note:* arith. = arithmetic coding; PD = Differentiated Peano-scanned data.

- The *adaptive templates* block provides substantial coding gain by looking for horizontal periodicity in the bilevel image. When a periodic template is changed, the encoder multiplexes a control sequence into the output data stream.
- The *model templates* block is a context arithmetic coder. The context is determined by ten particular neighboring pixels that are defined by two model templates: the three-line template and the two-line template as shown in Figure 11.33 (labeled as 'X' and 'A', where 'A' is the adaptive pixel whose position could be varied during the coding process [1008, 1009]).
- The *adaptive arithmetic encoder* is an entropy coder that determines the necessity of coding a given pixel based upon the outputs of the typical prediction block and the model templates block. If necessary,

the encoder notes the context and uses its internal probability estimator to estimate the conditional probability that the current pixel will be of a given value.

It should be noted that the JBIG coding algorithm includes at least three decorrelation steps for a given bilevel image (or bit plane).




FIGURE 11.32

The single-layer JBIG encoder. Reproduced with permission from L. Shen and R.M. Rangayyan, “Lossless compression of continuous-tone images by combined inter-bit-plane decorrelation and JBIG coding”, *Journal of Electronic Imaging*, 6(2): 198 – 207, 1997. © SPIE and IS&T.




FIGURE 11.33

Two context-model templates used in the single-layer JBIG encoder. ?: Pixel being encoded. X: Pixels in the context model. A: Adaptive pixel. Reproduced with permission from L. Shen and R.M. Rangayyan, “Lossless compression of continuous-tone images by combined inter-bit-plane decorrelation and JBIG coding”, *Journal of Electronic Imaging*, 6(2): 198 – 207, 1997. © SPIE and IS&T.

In a 1D form of JBIG coding, run-length coding is used to encode each line in the image by using a variable-length coding scheme; see Gonzalez and Woods [8] for further details of this and other JBIG procedures. See Sections 11.11, 11.12, and 11.13 for discussions on the results of application of JBIG.

11.10.2 The JPEG Standard

JPEG is a continuous-tone image compression standard that supports both lossless and lossy coding [8, 1011, 1012]. The standard includes three systems [8]:

- A lossy baseline coding system based upon block-wise application of the DCT.
- An extended coding system for greater compression, higher precision, and progressive transmission and recovery.
- A lossless independent coding system for reversible compression.

In the baseline coding system, the pixel data are limited to a precision of $8b$. The image is broken into 8×8 blocks, shifted in gray level, and transformed using the DCT. The transform coefficients are quantized with variable-length code assignment (to a maximum of $11b$). Due to the presence of block artifacts and other errors in lossy compression, this mode of JPEG would not be suitable for the compression of medical images.

The JPEG 2000 standard is based upon the wavelet transform [8, 1019]. JPEG 2000 offers the option of progressive coding (from lossy toward lossless), as well as the option of coding ROIs with higher quality than that for the other regions in the given image [1019], which may be of interest in some applications.

The lossless mode of JPEG uses a form of predictive (DPCM) coding. A linear combination of each pixel's neighbors at the left, upper, and upper-left positions is employed to predict the pixel's value, and then the difference between the true value of the pixel and its predicted value is coded through an entropy coder, such as the Huffman or arithmetic coder. Lossless JPEG defines seven linear combinations known as prediction selection values (PSV).

For a continuous-tone image $f(m, n)$, $0 \leq m < M$, $0 \leq n < N$, the predictors used for an interior pixel $f(m, n)$, $0 < m < M$, $0 < n < N$ in lossless JPEG coding are as follows:

- PSV=0: no prediction [or $\tilde{f}(m, n) = 0$], which indicates entropy encoding of the original image directly;
- PSV=1: $\tilde{f}(m, n) = f(m - 1, n)$;
- PSV=2: $\tilde{f}(m, n) = f(m, n - 1)$;
- PSV=3: $\tilde{f}(m, n) = f(m - 1, n - 1)$;
- PSV=4: $\tilde{f}(m, n) = f(m - 1, n) + f(m, n - 1) - f(m - 1, n - 1)$;
- PSV=5: $\tilde{f}(m, n) = f(m - 1, n) + [f(m, n - 1) - f(m - 1, n - 1)]/2$;
- PSV=6: $\tilde{f}(m, n) = f(m, n - 1) + [f(m - 1, n) - f(m - 1, n - 1)]/2$;

- PSV=7: $\tilde{f}(m, n) = [f(m-1, n) + f(m, n-1)]/2$;

where $\tilde{f}(m, n)$ is the predicted value for the pixel $f(m, n)$. The boundary pixels could be treated through various possible ways, which will not make a significant difference in the final compression ratio due to their small population. A typical method is to use PSV=1 for the first row and PSV=2 for the first column, with special treatment of the first pixel at position (0,0) using the value of $2^K - 1$, where K is the number of precision bits of the pixels.

Sung et al. [1020] evaluated the application of JPEG 2000 for the compression of mammograms, and suggested that compression ratios of up to 15 : 1 were possible without visual loss, “preserving significant medical information at a confidence level of 99%”. It was also suggested that compression of up to 80 : 1 could be achieved “without affecting clinical diagnostic performance”. See Sections 11.11, 11.12, and 11.13 for discussions on the results of application of lossless JPEG to several test images.

11.10.3 The MPEG Standard

The MPEG standard includes several schemes for the compression of video images for various applications, based upon combinations of the DCT and DPCM, including motion compensation [1013, 8, 1021]. The techniques exploit data redundancy and correlation within each frame as well as between frames; furthermore, they take advantage of certain psychovisual properties of the human visual system. Recent versions of MPEG include special features suitable for video-conference, multimedia, streaming media, and video-game systems [1021]. Most of such special features are not of relevance in lossless compression of biomedical image data.

11.10.4 The ACR/ NEMA and DICOM Standards

The proliferation of medical imaging technology and devices of several types in the 1970s and 1980s led to a situation where, due to the lack of standards, interconnection and communication of data between imaging and computing devices was not possible. In order to rectify this situation, the ACR and NEMA established the ACR/ NEMA 300 standard [1017, 1018] on digital imaging and communication, specifying the desired hardware interface, a minimum set of software commands, and a consistent set of data formats to facilitate communication between imaging devices and computers across networks. This was followed by another standard on data compression — the ACR/ NEMA PS 2 [1016] — specifying the manner in which header data were to be provided such that a recipient of the compressed data could identify the data compression method and parameters used, and reconstruct the image data. The standard permits the use of several image decorrelation and data compression techniques, including transform (DCT), predictive (DPCM), Huffman, and Lempel–Ziv coding techniques. The DICOM standard [1015]

includes a number of enhancements to the ACR/ NEMA standard, including conformance levels and applicability to a networked environment.

11.11 Segmentation-based Adaptive Scanning

Shen and Rangayyan [321, 320] proposed a segmentation-based lossless image coding (SLIC) method based on a simple but efficient region-growing procedure. An embedded region-growing procedure was used to produce an adaptive scanning pattern for the given image with the help of a discontinuity-index map that required a small number of bits for encoding. The JBIG method was used for encoding both the error-image data and the discontinuity-index map data. The details of the SLIC method and the results obtained are described in the following paragraphs.

11.11.1 Segmentation-based coding

Kunt et al. [566] proposed a contour-texture approach to picture coding; they called such approaches “second-generation” image coding techniques. The main idea behind such techniques is to first segment the image into nearly homogeneous regions surrounded by contours such that the contours correspond, as much as possible, to those of the objects in the image, and then to encode the contour and texture information separately. Because contours can be represented as 1D signals and the pixels within a region are highly correlated, such methods are expected to lead to high compression ratios. Although the idea appears to be promising, its implementation meets with a series of difficulties. A major problem exists at its very important first step — segmentation — which determines the final performance of the segmentation-based coding method. It is well recognized that there are no satisfactory segmentation algorithms for application to a wide variety of general images. Most of the available segmentation algorithms are sophisticated and give good performance only for specific types of images.

In order to overcome the problem mentioned above in relation to segmentation, Shen and Rangayyan [321, 320] proposed a simple region-growing method. In this procedure, instead of generating a contour set, a discontinuity map is obtained during the region-growing procedure. Concurrently, the method also produces a corresponding error image based upon the difference between each pixel and its corresponding “center pixel”. The discontinuity map and the error image are then encoded separately.

11.11.2 Region-growing criteria

The aim of segmentation in image compression is not the identification of objects or the analysis of features; instead, the aim is to group spatially connected pixels lying within a small gray-level dynamic range. The region-growing procedure in SLIC starts with a single pixel, called the seed pixel (#0 in Figure 11.34). Each of the seed's 4-connected neighboring pixels, from #1 to #4 in the order as shown in Figure 11.34, is checked with a region-growing (or inclusion) condition. If the condition is satisfied, the neighboring pixel is included in the region. The four neighbors of the newly added neighboring pixel are then checked for inclusion in the region. This recursive procedure is continued until no spatially connected pixel meets the growing condition. A new region-growing procedure is then started with the next pixel in the image that is not already a member of a region; the procedure ends when every pixel in the image has been included in one of the regions grown.




FIGURE 11.34

Demonstration of a seed pixel (#0) and its 4-connected neighbors (#1 to #4) in region growing. Reproduced with permission from L. Shen and R.M. Rangayyan, “A segmentation-based lossless image coding method for high-resolution medical image compression”, *IEEE Transactions on Medical Imaging*, 16(3): 301 – 306, 1997. © IEEE.

The region-growing conditions used in SLIC are the following:

1. The neighboring pixel is not a member of any of the regions already grown.
2. The absolute difference between the neighboring pixel and the corresponding center pixel is less than the limit *error-level* (to be defined later).

Figure 11.35 demonstrates the relationship between a neighboring pixel and its center pixel: at the specific stage illustrated in the figure, pixel *A* has become a member of the region (3) being grown, and its four neighbors, namely *B*, *C*, *D*, and *E*, are being checked for inclusion (*E* is already a member of the region). Under this circumstance, pixel *A* is the center pixel of the neighboring pixels *B*, *C*, *D*, and *E*. When a new neighboring pixel is included in the region being grown, its *error-level*-shift-up difference with respect to its center pixel is stored as the pixel's "error" value. If only the first of the two region-growing conditions is met, the discontinuity index of the pixel is incremented. By this process, after region growing, a "discontinuity-index image data part" and an "error-image data part" will be obtained. The maximum value of the discontinuity index is 4. Most of the segmentation-based coding algorithms reported in the literature include contour coding and region coding; instead of these steps, the SLIC method uses a discontinuity-index map and an error-image data part.

The *error-level* in SLIC is determined by a preselected parameter *error-bits* as

$$\text{error-level} = 2^{(\text{error-bits} - 1)}. \quad (11.190)$$

For instance, if the error image is allowed to take up to 5 *b/pixel* (*error-bits* = 5), the corresponding *error-level* is 16; the allowed difference range is then $[-16, 15]$. The error value of the seed pixel of each region is defined as the value of its lower *error-bits* bits; the value of the higher ($N - \text{error-bits}$) bits of the pixel is stored in a "high-bits seed-data part", where N is the number of bits per pixel in the original image data.

The three data parts described above are used to fully recover the original image during the decoding process. The region-growing conditions during decoding are that the neighboring pixel under consideration for inclusion be not in any of the previously grown regions, and that its discontinuity index equal 0. When the conditions are met for a pixel, its value is restored as the sum of its *error-level*-shift-down error value and its center pixel value (except for the seed pixels of every region). If only the first of the two conditions is satisfied, the discontinuity index of that pixel is decremented. Thus, the discontinuity index generated during segmentation is used to guide region growing during decoding. The "high-bits seed-data part" is combined with the "error-image data part" to recover the seed pixel value of each region.

Figure 11.36 provides a simple example for illustration of the region-growing procedure and its result. The 8×8 image in the example, shown in Figure

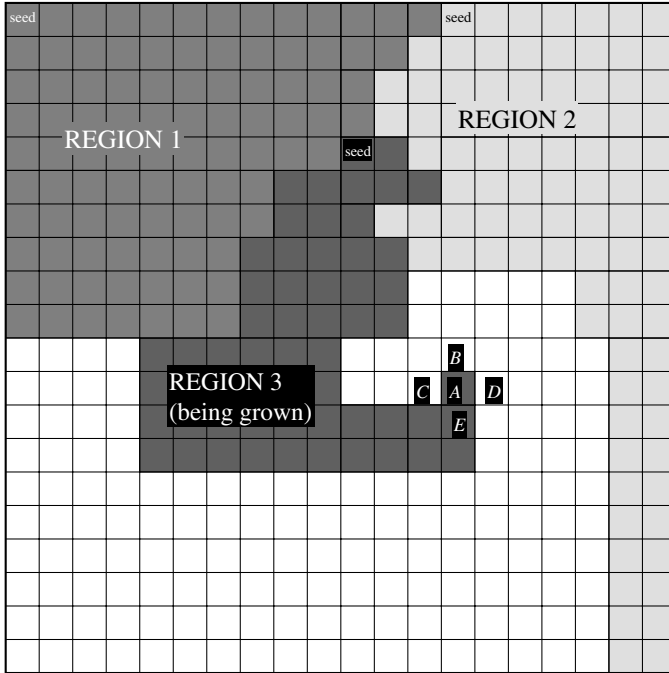


FIGURE 11.35

Demonstration of a neighboring pixel (B , C , D , or E) being checked for inclusion against the current center pixel (A) during region growing. Reproduced with permission from L. Shen and R.M. Rangayyan, "A segmentation-based lossless image coding method for high-resolution medical image compression", *IEEE Transactions on Medical Imaging*, 16(3): 301 – 306, 1997. © IEEE.

11.36 (a), is a section of an eye of the 512×512 Lenna image [shown in Figure 11.37 (a)]. The value of *error-bits* was set to 5 for this example. Figure 11.36 (b) shows the result of region growing. The corresponding three data parts, namely the discontinuity-index image data, error-image data, and high-bits seed data, are shown in Figure 11.36 (c), (d), and (e), respectively. The full 512×512 Lenna image and the corresponding discontinuity-index image data (scaled) and error-image data (scaled) are shown in Figure 11.37.

11.11.3 The SLIC procedure

The complete SLIC procedure is illustrated in Figure 11.38. At the encoding end, the original image is transformed into three parts: discontinuity-index image data, error-image data, and high-bits seed data, by the region-growing procedure. The first two data parts are encoded using the Gray code (see Table 11.1), broken down into bit planes, and finally encoded using JBIG. The last data part is stored or transmitted as is; it needs only $N - \text{error-bits}$ bits per region.

At the decoding end, the JBIG-coded data files are JBIG-decoded first, and then the Gray-coded bit planes are composed back to binary code. Finally, the three parts are combined together by the same region-growing procedures as before to recover the original image.

A lossless compression method basically includes two major stages: one is image transformation, with the purpose of data decorrelation; the other is encoding of the transformed data. However, in the SLIC procedure, image transformation is achieved in both the region-growing procedure, and later, in the JBIG procedure, whereas encoding is accomplished within the JBIG procedure.

11.11.4 Results of image data compression with SLIC

Five mammograms and five chest radiographs were used to evaluate the performance of SLIC [320, 321]. The procedure was tested using 8 *b/pixel* and 10 *b/pixel* versions of the images obtained by direct mapping and by discarding the two least-significant bits, respectively, from the original 12 *b* images. The method was also tested with commonly used 8 *b* nonmedical test images. The performance of SLIC was compared with that of JBIG [1008, 1009], JPEG [1012], adaptive Lempel–Ziv (ALZ) coding [966], HINT [978], and 2D LP coding [174]. In using JBIG, for direct encoding of the image or for encoding the error-image data part and the discontinuity-index data part, parameters were selected so as to use three lines of the image ($NLPS0 = 3$) in the underlying model (“3D”) and no progressive spatial resolution buildup. The lossless JPEG package used in the study includes features of automatic determination of the best prediction pattern and optimal Huffman table generation. The UNIX utility *compress* was used for ALZ compression. For the 10 *b* test images, the 2D Burg LP algorithm (see Section 11.8.1) followed by

84	84	91	83	72	57	66	126
86	90	80	76	55	65	113	173
54	60	57	64	77	107	160	198
63	65	75	88	127	158	188	202
101	102	116	137	163	186	197	198
132	138	146	157	182	187	193	197
149	156	156	158	168	173	176	187
133	142	151	154	151	158	169	167

(a)

84 (seed)	84	91	83	72	57	66	126 (seed)
86	90	80	76	55	65	113 (seed)	173 (seed)
54	60	57	64	77	107 (seed)	160 (seed)	198 (seed)
63	65	75	88	127 (seed)	158 (seed)	188	202
101 (seed)	102	116	137 (seed)	163 (seed)	186	197	198
132	138	146	157	182	187	193	197
149	156	156	158	168	173	176	187
133	142	151	154	151	158	169	167

(b)

0	0	0	0	0	0	0	1
0	0	0	0	2	0	2	2
1	1	1	0	0	2	2	2
0	0	1	1	2	4	1	0
2	2	2	3	4	0	0	0
0	0	0	1	0	0	0	0
0	0	0	0	0	0	1	0
1	0	0	0	1	0	0	1

(c)

91-84+16

20	16	23	8	5	1	25	30
18	22	5	9	6	24	17	13
10	19	9	4	29	11	0	6
25	21	26	29	31	30	2	20
5	17	30	9	3	5	25	12
10	8	5	15	11	17	12	15
9	16	14	6	2	2	5	6
7	2	11	12	9	1	9	14

169-176+16

(d)

2	3	3	5	3	5	6	3	4	3	4	5
---	---	---	---	---	---	---	---	---	---	---	---

(e)

FIGURE 11.36


A simple example of the region-growing procedure and its result with *error-bits* set to be 5. (a) Original image with the size of 8 rows by 8 columns (a section of the 512×512 Lenna image shown in Figure 11.37). (b) The result of region growing. The seed pixel of every region grown is identified; a few regions include only the corresponding seed pixels. (c) The discontinuity-index image data part. (d) The error-image data part. (e) The high-bits seed-data part (from left to right). Reproduced with permission from L. Shen and R.M. Rangayyan, "A segmentation-based lossless image coding method for high-resolution medical image compression", *IEEE Transactions on Medical Imaging*, 16(3): 301 – 306, 1997. © IEEE.



(a)



(b)



(c)

FIGURE 11.37

The 512×512 Lenna image and the results with SLIC (with *error-bits* = 5): (a) Original image. (b) The discontinuity-index image data part (scaled). (c) The error-image data part (scaled). See also Figure 11.36. Reproduced with permission from L. Shen and R.M. Rangayyan, “A segmentation-based lossless image coding method for high-resolution medical image compression”, *IEEE Transactions on Medical Imaging*, 16(3): 301 – 306, 1997. © IEEE.

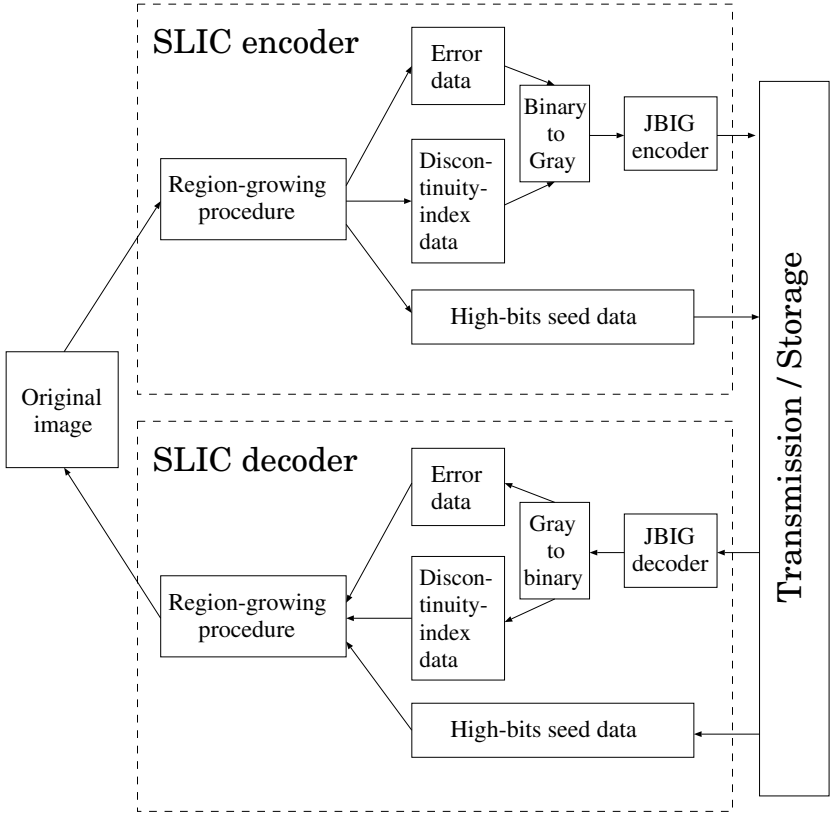


FIGURE 11.38

Illustration of the segmentation-based lossless image coding (SLIC) procedure. Reproduced with permission from L. Shen and R.M. Rangayyan, "A segmentation-based lossless image coding method for high-resolution medical image compression", *IEEE Transactions on Medical Imaging*, 16(3): 301 – 306, 1997. © IEEE.

Huffman error coding (2D-Burg-Huffman) was utilized instead of the JPEG package. (The 2D-Burg-Huffman program was designed specifically for 10 *b* images; the JPEG program permits only 8 *b* images.)

The SLIC method has a tunable parameter, which is *error-bits*. The data compression achieved was found to be almost the same for most of the 8 *b* medical images by using the value 4 or 5 for *error-bits*. Subsequently, *error-bits* = 5 was used in the remaining experiments with the 8 *b* versions of the medical images. The performance of the SLIC method is summarized in Table 11.11, along with the results of ALZ, JBIG, JPEG, and HINT compression. It is seen that SLIC has outperformed all of the other methods studied with the test-image set used, except in the case of one mammogram and one chest radiograph for which JBIG gave negligibly better results. On the average, SLIC improved the bit rate by about 9%, 29%, and 13%, as compared with JBIG, JPEG, and HINT, respectively.

TABLE 11.11

Average Bits Per Pixel with SLIC (*error-bits* = 5), ALZ, JBIG, JPEG (Best Mode), and HINT Using Five 8 *b* Mammograms (m1 to m4 and m9) and Five 8 *b* Chest Radiographs (c5 to c8 and c10) by Bits/Pixel.

Image (8 <i>b/pixel</i>)	Entropy					
	H_0	ALZ	JBIG	JPEG	HINT	SLIC
m1 (4,096 × 1,990)	6.13	2.95	1.82	2.14	1.87	1.73
m2 (4,096 × 1,800)	6.09	2.89	1.84	2.18	1.92	1.78
m3 (3,596 × 1,632)	5.92	3.02	2.40	2.37	2.12	2.12
m4 (3,580 × 1,696)	5.90	2.45	1.96	1.89	1.61	1.44
c5 (3,536 × 3,184)	7.05	3.03	1.60	1.92	1.75	1.42
c6 (3,904 × 3,648)	7.46	3.32	1.88	2.15	2.00	1.76
c7 (3,120 × 3,632)	5.30	1.73	0.95	1.60	1.40	0.99
c8 (3,744 × 3,328)	7.45	3.13	1.50	1.89	1.64	1.35
m9 (4,096 × 2,304)	6.04	2.60	1.67	2.12	1.84	1.67
c10 (3,664 × 3,680)	7.17	2.95	1.63	1.97	1.73	1.50
Average	6.45	2.81	1.72	2.02	1.79	1.58

The lowest bit rate in each case is highlighted. Reproduced with permission from L. Shen and R.M. Rangayyan, "A segmentation-based lossless image coding method for high-resolution medical image compression", *IEEE Transactions on Medical Imaging*, 16(3): 301 – 306, 1997. © IEEE.

Whereas the SLIC procedure performed well with high-resolution medical images, its performance with low-resolution general images was comparable to the performance of JBIG and JPEG, as shown in Table 11.12. When SLIC used an optimized JBIG algorithm with the maximum number of lines, that is, $NLPS0 = \text{row}$, the number of rows of the image (the last column of Table 11.12) in the inherent model instead of only three lines ($NLPS0=3$ in Table 11.12), its performance was better than that of JBIG and comparable to that of JPEG.

TABLE 11.12

Average Bits Per Pixel with SLIC ($\text{error-bits} = 8$), ALZ, JBIG, and JPEG (Best Mode) Using Eight Commonly Used 8 *b* Images.

Image (8 <i>b</i> /pixel)	Entropy (order=0)	SLIC				
		ALZ	JBIG	JPEG	NLPS0	
					3	row
Airplane (512×512)	6.49	5.71	4.24	4.20	4.22	4.11
Baboon (512×512)	7.14	7.84	6.37	6.18	6.55	6.44
Cameraman (256×256)	7.01	6.68	4.92	4.95	5.14	4.91
Lenna (256×256)	7.57	7.91	5.37	5.07	5.26	5.04
Lenna (512×512)	7.45	7.05	4.80	4.69	4.74	4.63
Peppers (512×512)	7.37	6.86	4.82	4.76	4.90	4.80
Sailboat (512×512)	7.27	7.07	5.34	5.26	5.46	5.36
Tiffany (512×512)	6.38	5.88	4.38	4.35	4.43	4.32
Average	7.19	6.87	5.03	4.93	5.09	4.95

Reproduced with permission from L. Shen and R.M. Rangayyan, “A segmentation-based lossless image coding method for high-resolution medical image compression”, *IEEE Transactions on Medical Imaging*, 16(3): 301 – 306, 1997. © IEEE.

In the studies of Shen and Rangayyan [321, 320], setting $\text{error-bits} = 6$ was observed to be a good choice for the compression of the 10 *b* versions of the medical images. Table 11.13 lists the results of compression with the ALZ, JBIG, 2D-Burg-Huffman, HINT, and SLIC methods. The SLIC technique has provided lower bit rates than the other methods. The average bit rate

is 2.92 *b/pixel* with SLIC, whereas HINT, JBIG, and 2D-Burg-Huffman have average bit rates of 3.03, 3.17, and 3.14 *b/pixel*, respectively.

TABLE 11.13

Average Bits Per Pixel with SLIC (*error-bits* = 6), ALZ, JBIG, 2D-Burg-Huffman (2DBH), and HINT Using Five 10 *b* Mammograms (m1 to m4 and m9) and Five 10 *b* Chest Radiographs (c5 to c8 and c10).

Image (10 <i>b/pixel</i>)	Entropy (order=0)	ALZ	JBIG	2DBH	HINT	SLIC
m1 (4,096 × 1,990)	7.27	5.45	2.89	2.92	2.75	2.73
m2 (4,096 × 1,800)	7.61	6.02	3.19	3.19	3.15	3.08
m3 (3,596 × 1,632)	6.68	4.73	3.13	2.97	2.81	2.81
m4 (3,580 × 1,696)	7.21	5.14	3.20	2.76	2.58	2.57
c5 (3,536 × 3,184)	8.92	7.21	3.33	3.31	3.28	2.99
c6 (3,904 × 3,648)	9.43	7.84	3.87	3.81	3.89	3.59
c7 (3,120 × 3,632)	7.07	4.67	2.30	2.58	2.34	2.27
c8 (3,744 × 3,328)	9.29	7.40	3.25	3.16	3.02	2.89
m9 (4,096 × 2,304)	7.81	5.93	3.18	3.41	3.29	3.17
c10 (3,664 × 3,680)	9.05	7.51	3.35	3.27	3.18	3.07
Average	8.03	6.19	3.17	3.14	3.03	2.92

The lowest bit rate in each case is highlighted. Reproduced with permission from L. Shen and R.M. Rangayyan, "A segmentation-based lossless image coding method for high-resolution medical image compression", *IEEE Transactions on Medical Imaging*, 16(3): 301 – 306, 1997. © IEEE.

Most of the previously reported segmentation-based coding techniques [566, 1022, 1023] involve three procedures: segmentation, contour coding, and texture coding. For segmentation, a sophisticated procedure is generally employed for the extraction of closed contours. In the SLIC method, a simple single-scan neighbor-pixel checking algorithm is used. The major problem after image segmentation with the other methods lies in the variance of pixel intensity within the regions, which has resulted in the application of such methods to lossy coding instead of lossless coding. In order to overcome this problem, the SLIC method uses the most-correlated neighboring pixels to generate a low-dynamic-range error image during segmentation. Contour

coding is replaced by a coding step applied to a discontinuity-index map with a maximum of 5 levels, and texture coding is turned into the encoding of a low-dynamic-range error image. Further improvement of the performance of SLIC may be possible by the application of more efficient coding methods to the error image and discontinuity-index data parts, and by modifying the region-growing procedure.

The SLIC method was extended to the compression of 3D biomedical images by Lopes and Rangayyan [1024]. The performance of the method varied depending upon the nature of the image. However, the advantage of implementation of SLIC in 3D versus 2D on a slice-by-slice basis was not significant. For example, both 2D and 3D SLIC with the compression utility *bzip2* [1025] (as the encoding scheme after decomposition of the image by the segmentation-based procedure) reduced the data in a 3D CT head examination from the original 16 *b/voxel* size to 5.4 *b/voxel*; the zeroth-order entropy of the image was 7.9 *b/voxel*. The inclusion of the SLIC procedure improved the performance of the compression utilities *gzip* and *compress* by 15 – 20%.

Acha et al. [1026] extended the SLIC method to the compression of color images in the application of diagnosis of burn injuries. Color images of size 832×624 pixels at 24 *b/pixel* in the RGB format were compressed to the rate of 7.7 *b/pixel*; application of the JPEG lossless method resulted in a rate of 8.2 *b/pixel*. In a further study, Serrano et al. [1027] converted the same images as in the study of Acha et al. [1026] from the RGB format to the YIQ (luminance, in-phase, and quadrature) system in a lossless manner, and showed that the bit rate could be further reduced by the use of SLIC to 4.6 *b/pixel*.

11.12 Enhanced JBIG Coding

The SLIC procedure demonstrates a successful example of the application of the JBIG method for coding the discontinuity-index map and error-data parts. We may, therefore, expect the incorporation of decorrelation procedures, such as predictive algorithms, into JBIG coding to provide better performance. Shen and Rangayyan [1028, 320] proposed a combination of multiple decorrelation procedures, including a lossless JPEG-based predictor, a transform-based inter-bit-plane decorrelator, and a JBIG-based intra-bit-plane decorrelator; the details of their procedure are described in the following paragraphs.

Although JBIG includes an efficient intra-bit-plane decorrelation procedure, it needs an efficient preprocessing algorithm for inter-bit-plane decorrelation in order to achieve good compression efficiency with continuous-tone images. There are several ways in which bit planes may be created from a continuous-tone image. One common choice is to use the bits of a folded-binary or Gray

representation of intensity. The Gray code is the most common alternative to the binary code for representing digital numbers; see Table 11.1. The major advantage of the Gray code is that only one bit changes between each pair of successive code words, which is useful to provide a good bit-plane representation for original image pixels: most neighboring pixels have highly correlated and close values, and thus most neighboring bits within each Gray-coded bit plane may be expected to have the same value. It has been shown that by using the Gray representation, JBIG can obtain compression ratios at least comparable to those of lossless JPEG [1029].

Coding schemes other than the Gray code may be derived based on specific requirements. For instance, for a K -bit image $f(m, n)$, the prediction error $e(m, n)$ could be represented as

$$e(m, n) = f(m, n) - \tilde{f}(m, n), \quad (11.191)$$

where $\tilde{f}(m, n)$ is the predicted value of the original pixel $f(m, n)$. In general, up to $K + 1$ bits could be required to represent the difference between two K -bit numbers. However, because the major concern in compression is to retrieve $f(m, n)$ from $e(m, n)$ and $\tilde{f}(m, n)$, we could make use of the following binary arithmetic operation:

$$\tilde{e}(m, n) = e(m, n) \& \underbrace{\{0\underbrace{11 \cdots 1}_{K \text{ bits}}\}}_b, \quad (11.192)$$

where $\&$ is the bit-wise AND operation, and the subscript b indicates binary representation. Then, only K bits are necessary to represent the prediction error $\tilde{e}(m, n)$. The original pixel value may be retrieved as

$$f(m, n) = [\tilde{e}(m, n) + \tilde{f}(m, n)] \& \underbrace{\{0\underbrace{11 \cdots 1}_{K \text{ bits}}\}}_b. \quad (11.193)$$

With the transformation as above, the value of $(2^K - v)$ for $\tilde{e}(m, n)$ denotes a prediction error $e(m, n)$ of either $(2^K - v)$ or $-v$. In general, the lower and the higher ends of the value of the error $\tilde{e}(m, n)$ appear more frequently than mid-range values.

The F_1 transformation [1030] could make bit-plane coding more efficient by increasing the run length in the most-significant bits:

$$v_1 = F_1(v) = \begin{cases} 0; & v = 0 \\ 2v - 1; & v \leq 2^{K-1} \\ 2(2^K - v); & v > 2^{K-1} \end{cases}, \quad (11.194)$$

with the inverse transform given by

$$v = F_1^{-1}(v_1) = \begin{cases} 0; & v_1 = 0 \\ 2^K - p; & v_1 = 2p \text{ (even)} \\ q; & v_1 = 2q - 1 \text{ (odd)} \end{cases}. \quad (11.195)$$

The F_2 transformation [1030, 1031] has a similar function, but through the reversal of the higher-half of the value range:

$$v_2 = F_2(v) = \begin{cases} v; & v < 2^{K-1} \\ v \oplus \underbrace{\{0 \ 11 \cdots 1\}}_{K-1 \text{ bits}}; & v \geq 2^{K-1} \end{cases}, \quad (11.196)$$

where \oplus is the bit-wise exclusive OR operation; its inverse transform is

$$v = F_2^{-1}(v_2) = F_2(v_2). \quad (11.197)$$

Shen and Rangayyan [1028] investigated the combined use of the PSV system in JPEG (see Section 11.10.2) and bit-plane coding using JBIG, along with one of the Gray, F_1 , and F_2 transforms. In using JBIG (for direct coding of the original image or for coding the prediction error image), the method was parameterized to use the three-line model template and stripe size equal to the number of rows of the image with no progressive spatial resolution buildup (see Section 11.10.1). The lossless JPEG scheme was set to generate the optimal Huffman table for each PSV value.

The methods were tested with a commonly used set of eight images (see Table 11.8). A comparison of the compression efficiencies of lossless JPEG, JBIG, and PSV-incorporated JBIG with one of the Gray, F_1 , or F_2 transforms is shown in Figure 11.39, for one of the test images. Each group of vertical bars in each case consists of five bars, corresponding to the entropy of the prediction error image, followed by the actual bit rates with lossless JPEG coding, and PSV-incorporated JBIG bit-plane coding with the Gray, F_1 , and F_2 transformations, respectively. In the figure, there are three horizontal lines representing the zeroth-order entropy of the original image, the bit rate by direct JBIG coding with the Gray transformation, and the best bit rate among all of the methods tested. The best performance for each test image was achieved with PSV-incorporated JBIG bit-plane coding with the F_1 transformation and PSV=7, except for the 256×256 Lenna image, for which the best rate was given with PSV=6 and JBIG bit-plane coding of the prediction error after F_2 transformation. The F_1 and F_2 transforms provided similar performance and performed better than the Gray transform, with the F_1 transform giving slightly lower bit rates in most of the cases.

In the results obtained by Shen and Rangayyan [1028], it was observed that the zeroth-order entropies of the prediction error images were much lower than those of the original images, that the bit rates by lossless JPEG compression were always higher than the zeroth-order entropies of the prediction error images, and that PSV-incorporated JBIG bit-plane coding provided bit rates lower than the zeroth-order entropies of the prediction error images (with the exceptions being the Baboon and Sailboat images). These observations show that a simple prediction procedure, such as the PSV scheme employed in lossless JPEG, is useful for decorrelation. In particular, prediction with PSV=7 followed by the F_1 transform and JBIG bit-plane coding achieves an

average bit rate that is about 0.2 b/pixel lower than those achieved with direct Gray-coded JBIG compression and the optimal mode of lossless JPEG. This also indicates that the F_1 transform is a better inter-bit-plane decorrelator than the Gray code for prediction error images, and that the intra-bit-plane decorrelation steps within the JBIG algorithm are not redundant with prior decorrelation by the PSV system in JPEG.




FIGURE 11.39

Comparison of image compression efficiency with the enhanced JBIG scheme. The five bars in each case represent the zeroth-order entropy of the prediction error image and actual bit rates with lossless JPEG, and PSV-incorporated JBIG with the Gray, F_1 , and F_2 transformation, respectively (from left to right). Reproduced with permission from L. Shen and R.M. Rangayyan, “Lossless compression of continuous-tone images by combined inter-bit-plane decorrelation and JBIG coding”, *Journal of Electronic Imaging*, 6(2): 198 – 207, 1997. © SPIE and IS&T.

Shen and Rangayyan [1028] applied the best mode of the PSV-incorporated JBIG bit-plane coding scheme (the PSV=7 predictor followed by JBIG bit-plane coding of F_1 -transformed prediction error, denoted as PSV7- F_1 -JBIG) to the JPEG standard set of continuous-tone test images; the results are shown in Table 11.14. The table also shows the zeroth-order entropies of the prediction error images, the results of the best mode of lossless JPEG coding

(for 8 b component images only), the results of direct Gray-transformed JBIG coding, and the results of the best mode of the CREW technique (Compression with Reversible Embedded Wavelets, one of the methods proposed to the Committee of Next Generation Lossless Compression of Continuous-tone Still Pictures) [1032, 1033].

The results in Table 11.14 demonstrate that the enhanced JBIG bit-plane coding scheme for continuous-tone images performs the best among the four algorithms tested. In terms of the average bit rate, the scheme outperforms direct JBIG coding and the best mode of CREW by 0.13 and 0.12 $b/pixel$, respectively, and achieves much lower bit rates than the zeroth-order entropies of the prediction error images by an average of 0.46 $b/pixel$ for the entire test set of images. If only the 8 b component images are considered, the enhanced JBIG technique provides better compression performance by 0.56, 0.08, 0.52, and 0.18 $b/pixel$, in terms of average bit rates when compared with lossless JPEG coding, direct Gray-transform JBIG coding, the zeroth-order entropy of the prediction error image, and the best mode of lossless CREW, respectively.

For comparison with SLIC in the context of radiographic images, the enhanced JBIG (PSV7-F1-JBIG) procedure was applied for compression of the five mammograms and five chest radiographs listed in Tables 11.11 and 11.13. The bit rates achieved for the 8 b and 10 b versions of the images are listed in Tables 11.15 and 11.16, respectively. The results indicate that the enhanced JBIG procedure provides an additional 5% improvement over SLIC on the 8 b image set, and that the method lowers the average bit rates to 2.75 $b/pixel$ from the 2.92 $b/pixel$ rate of SLIC for the 10 b images.

11.13 Lower-limit Analysis of Lossless Data Compression

There is, as yet, no practical technique available for the determination of the lowest limit of bit rate in reversible compression of a given image, although such a number should exist, based upon information theory. It is, therefore, difficult to judge how good a compression algorithm is, other than by comparing its performance with those of other published methods, or with the zeroth-order entropy of the decorrelated data, if available: the latter approach is commonly used by researchers when different test-image sets are involved, and when other compression programs are not available. Either of the two approaches mentioned above can only analyze relatively how well a compression algorithm performs, in comparison with the other techniques available or the zeroth-order entropy. It should be apparent from the results presented in the preceding sections that the usefulness of comparative analysis is limited. For example, SLIC provided the best compression results among the methods

TABLE 11.14

Compression of the JPEG Test Image Set Using Enhanced JBIG (EJBIG), Lossless JPEG (8 *b* Component Images Only), Direct Gray-transformed JBIG Coding, and CREW Coding (Bits/Component) with the Best Bit Rate Highlighted in Each Case.

Image (cols×rows×comp.×bits)	Best	Direct	PSV = 7		Best
	JPEG	JBIG	H_{e0}	EJBIG	CREW
hotel (720×576×3×8)	4.37	4.20	4.26	4.08	4.05
gold (720×576×3×8)	4.33	4.31	4.16	4.13	4.08
bike (2,048×2,560×4×8)	4.33	3.94	4.20	3.80	3.92
woman (2,048×2,560×4×8)	4.84	4.64	4.74	4.37	4.33
cafe (2,048×2,560×4×8)	5.63	5.27	5.60	5.17	5.17
tools (1,524×1,200×4×8)	5.69	5.38	5.60	5.37	5.47
bike3 (781×919×3×8)	5.15	4.58	5.07	4.73	5.11
water (3,072×2,048×3×8)	2.62	1.96	2.50	1.89	1.86
cats (3,072×2,048×3×8)	3.70	2.90	3.64	2.70	2.65
aerial1 (1,024×1,024×3×11)	—	8.96	8.75	8.79	8.71
aerial2 (720×1,024×3×8)	4.93	4.61	5.00	4.44	4.47
cmpnd1 (512×768×3×8)	2.51	1.32	2.24	1.76	2.26
cmpnd2 (1,024×1,400×3×8)	2.50	1.39	2.26	1.72	2.32
finger (512×512×1×8)	5.85	6.49	5.85	5.84	5.84
x_ray (2,048×1,680×1×12)	—	6.55	6.26	6.19	6.10
cr (1,744×2,048×1×10)	—	5.61	5.32	5.43	5.38
ct (512×512×1×12)	—	4.66	5.32	4.01	4.08
us (512×488×1×8)	3.04	2.57	3.61	2.57	2.92
mri (256×256×1×11)	—	6.62	6.45	6.33	6.06
faxball (1,024×512×3×8)	1.50	0.67	1.38	0.51	1.15
graphic (2,644×3,046×3×8)	2.81	2.84	2.79	2.34	2.51
chart (1,752×2,375×3×8)	2.23	1.33	2.26	1.34	1.66
chart_s (1,688×2,347×3×8)	3.86	2.87	3.97	3.02	3.26
Average (all)	—	4.07	4.40	3.94	4.06
Average (8 <i>b</i> / comp. only)	3.88	3.40	3.84	3.32	3.50

Note: H_{e0} = zeroth-order entropy of the PSV prediction error; comp. = number of components; cols = number of columns. Reproduced with permission from L. Shen and R.M. Rangayyan, "Lossless compression of continuous-tone images by combined inter-bit-plane decorrelation and JBIG coding", *Journal of Electronic Imaging*, 6(2): 198 – 207, 1997. © SPIE and IS&T.

TABLE 11.15

Comparison of Enhanced JBIG (PSV7-F1-JBIG or EJBIG) with JBIG, JPEG (Best Lossless Mode), HINT, and SLIC (*error-bits* = 5) Using Five 8 *b* Mammograms (m1 to m4 and m9) and Five 8 *b* Chest Radiographs (c5 to c8 and c10) by Bits/Pixel [320].

Image (8 <i>b/pixel</i>)	H_0	JBIG	JPEG	HINT	SLIC	EJBIG
m1 (4,096×1,990)	6.13	1.82	2.14	1.87	1.73	1.62
m2 (4,096×1,800)	6.09	1.84	2.18	1.92	1.78	1.66
m3 (3,596×1,632)	5.92	2.40	2.37	2.12	2.12	1.93
m4 (3,580×1,696)	5.90	1.96	1.89	1.61	1.44	1.34
c5 (3,536×3,184)	7.05	1.60	1.92	1.75	1.42	1.41
c6 (3,904×3,648)	7.46	1.88	2.15	2.00	1.76	1.70
c7 (3,120×3,632)	5.30	0.95	1.60	1.40	1.00	0.95
c8 (3,744×3,328)	7.45	1.50	1.89	1.64	1.35	1.40
m9 (4,096×2,304)	6.04	1.67	2.12	1.84	1.67	1.54
c10 (3,664×3,680)	7.17	1.63	1.97	1.73	1.50	1.44
Average	6.45	1.72	2.02	1.79	1.58	1.50

The lowest bit rate is highlighted in each case. See also Table 11.11. *Note:* H_0 = zeroth-order entropy.

TABLE 11.16

Comparison of Enhanced JBIG (PSV7-F1-JBIG or EJBIG) with JBIG, 2D-Burg-Huffman (2DBH), HINT, and SLIC (*error-bits* = 6) Using Five 10 *b* Mammograms (m1 to m4 and m9) and Five 10 *b* Chest Radiographs (c5 to c8 and c10) by Bits/Pixel [320].

Image (10 <i>b/pixel</i>)	H_0	JBIG	2DBH	HINT	SLIC	EJBIG
m1 (4,096×1,990)	7.27	2.89	2.92	2.75	2.73	2.59
m2 (4,096×1,800)	7.61	3.19	3.19	3.15	3.08	2.88
m3 (3,596×1,632)	6.68	3.13	2.97	2.81	2.81	2.57
m4 (3,580×1,696)	7.21	3.20	2.76	2.58	2.57	2.39
c5 (3,536×3,184)	8.92	3.33	3.31	3.28	3.00	2.88
c6 (3,904×3,648)	9.43	3.87	3.81	3.89	3.59	3.38
c7 (3,120×3,632)	7.07	2.30	2.58	2.34	2.27	2.14
c8 (3,744×3,328)	9.29	3.25	3.16	3.02	2.89	2.81
m9 (4,096×2,304)	7.81	3.18	3.41	3.29	3.17	2.93
c10 (3,664×3,680)	9.05	3.35	3.27	3.18	3.07	2.90
Average	8.03	3.17	3.14	3.03	2.92	2.75

The lowest bit rate is highlighted in each case. See also Table 11.13. *Note:* H_0 = zeroth-order entropy.

tested in Section 11.11; however, it is seen in Section 11.12 that the enhanced JBIG scheme performs better than SLIC.

Even if it were to be not practical to achieve the lowest possible bit rate in the lossless compression of a given image, it is of interest to estimate an achievable lower-bound bit rate for an image. Information theory indicates that the lossless compression efficiency is bounded by high-order entropy values. However, the accuracy of estimating high-order statistics is limited by the length of the data (the number of samples available), the number of intensity levels, and the order. The highest possible order of entropy that can be estimated with high accuracy is limited due to the finite length of the data available. In spite of this limitation, high-order entropy values can provide a better estimate of the lower-bound bit rate than the commonly used zeroth-order entropy. Shen and Rangayyan [1028, 320] proposed methods for the estimation of high-order entropy and the lower-bound bit rate, which are described in the following paragraphs.

11.13.1 Memoryless entropy

A memoryless source is the simplest form of an information source, in which successive source symbols are statistically independent [1034]. Such a source is completely specified by its source alphabet $A = \{a_0, a_1, a_2, \dots, a_n\}$ and the associated probabilities of occurrence $\{p(a_0), p(a_1), p(a_2), \dots, p(a_n)\}$. The memoryless entropy $H(A)$, which is also known as the average amount of information per source symbol, is defined as

$$H(A) = - \sum_{i=0}^n p(a_i) \log_2 p(a_i). \quad (11.198)$$

The m^{th} -order extension entropy is defined as

$$\begin{aligned} H_m(A) &= H_m(a_{i_0}, a_{i_1}, a_{i_2}, \dots, a_{i_m}) \\ &= - \sum_{A^m} p(a_{i_0}, a_{i_1}, a_{i_2}, \dots, a_{i_m}) \log_2 p(a_{i_0}, a_{i_1}, a_{i_2}, \dots, a_{i_m}), \end{aligned} \quad (11.199)$$

where $p(a_{i_0}, a_{i_1}, a_{i_2}, \dots, a_{i_m})$ is the probability of a symbol string from the m^{th} -order extension of the memoryless source, and A^m represents the set of all possible strings with m symbols following a_{i_0} . For a memoryless source, using the property

$$p(a_{i_0}, a_{i_1}, a_{i_2}, \dots, a_{i_m}) = p(a_{i_0}) p(a_{i_1}) p(a_{i_2}) \dots p(a_{i_m}), \quad (11.200)$$

we get

$$H(A) = \frac{H_m(A)}{m+1}. \quad (11.201)$$

11.13.2 Markov entropy

A memoryless source model could be restrictive in many applications due to the fact that successive source symbols can be significantly interdependent, which means the source has memory. Image sources are such examples, in which there always exists some statistical dependence among neighboring pixels, even after the source symbol stream has been decorrelated. A source possessing dependence or memory as above may be modeled as a Markov source, in which the probability of occurrence of a source symbol a_i depends upon the probabilities of occurrence of a finite number m of the preceding symbols [1034]. The corresponding m^{th} -order Markov entropy $H(a_{i_0} | a_{i_1}, a_{i_2}, \dots, a_{i_m})$ may be computed as

$$H(a_{i_0} | a_{i_1}, a_{i_2}, \dots, a_{i_m}) = - \sum_{A^m} p(a_{i_0}, a_{i_1}, a_{i_2}, \dots, a_{i_m}) \times \log_2 p(a_{i_0} | a_{i_1}, a_{i_2}, \dots, a_{i_m}), \quad (11.202)$$

where $p(a_{i_0}, a_{i_1}, a_{i_2}, \dots, a_{i_m})$ is the probability of a particular state and $p(a_{i_0} | a_{i_1}, a_{i_2}, \dots, a_{i_m})$ is the PDF of a_{i_0} conditioned upon the occurrence of the string $\{a_{i_1}, a_{i_2}, \dots, a_{i_m}\}$. It may be shown that

$$H(a_{i_0} | a_{i_1}, a_{i_2}, \dots, a_{i_m}) \leq H(a_{i_0} | a_{i_1}, a_{i_2}, \dots, a_{i_{m-1}}) \leq \dots \leq H(a_{i_0} | a_{i_1}) \leq H(a_{i_0}) = H(A), \quad (11.203)$$

where the equality is satisfied if and only if the source symbols are statistically independent.

11.13.3 Estimation of the true source entropy

Although a given image or its prediction error image could be modeled with a Markov source, the order of the model and the conditional PDFs will be usually unknown. However, it is seen from Equation 11.203 that the higher the order of the conditional probability, the lower is the resulting entropy, which is closer to the true source entropy. In order to maintain a reasonable level of accuracy in the estimation of the conditional probability, larger numbers of data strings are needed for higher-order functions; the estimation error ϵ is given by [1035]

$$\epsilon = \frac{2^{mK} - 1}{2N \ln 2} \quad (11.204)$$

for an m^{th} -order Markov source model with 2^K intensity levels and N data samples. Thus, for a specific image with known K and N , the highest order of conditional Markov entropy that could be calculated within a practical estimation error of conditional probability is bounded by

$$\max m = \lfloor \frac{\log_2(2\epsilon N \ln 2 + 1)}{K} \rfloor, \quad (11.205)$$

where $\lfloor x \rfloor$ is the floor function that returns the largest integer less than or equal to the argument x . Therefore, given an error limit ϵ , the only way to derive a higher-order parameter is to decrease K by splitting data bytes, because the data length N cannot be extended.

For example, the highest orders of Markov entropy that may be calculated with a probability estimation error of less than 0.05 for a 512×512 , 8 *b/pixel* image are 1, 3, 7, and 14, for the original data (one 8 *b* data part, $K = 8$), split data with two 4 *b* data parts ($K = 4$), split data with four 2 *b* data parts ($K = 2$), and split data with eight 1 *b* data parts ($K = 1$), respectively. Figure 11.40 shows the Markov entropy values up to the maximum possible order of $\max m$ with $\epsilon = 0.05$ for four forms of splitting for the test image Airplane. It is obvious that the entropy values become larger with splitting into more data parts due to the high correlation present among the data parts, although the maximum order could go higher after splitting and the entropy values are decreasing with increasing order for each form of splitting; see also Table 11.17. This indicates that decorrelation of the data bits is needed before splitting in order to get a good estimate of the entropy, because the source entropy is not changed by any reversible transformation.

From the results of the enhanced JBIG algorithm, it is seen that the Gray, F_1 , and F_2 transformations provide good decorrelation among bit planes after PSV-based prediction. This is demonstrated with four plots of the binary, Gray, F_1 , and F_2 representations of PSV=7 prediction error data of the Airplane image in Figure 11.41 as well as in Table 11.17. The binary representation is seen to lead to poor decorrelation among the bits of the prediction error data: it actually makes the maximum-order entropy increase to 6.95 *b/pixel* while splitting the error data into eight 1 *b* data parts (the zeroth-order entropies of the original error data and the original image are 4.18 and 6.49 *b/pixel*, respectively). It is also seen that the highest-order entropy values that could be estimated within the error limit specified are increasing with increasing number of data parts when using the binary representation. On the other hand, with the Gray, F_1 , or F_2 transformation, the situation is different. It is seen that the highest-order entropy values with splitting become smaller than the entropy of the original data when one of the three transformations is used, which shows their efficient decorrelation effect among the bit planes of the prediction error image. Finally, it is seen that the F_1 transform is the best among the four representation schemes, with the lowest estimated entropy value of 3.46 *b/pixel* achieved when the prediction error is split into four 2 *b* data parts.

The lowest estimated entropy value is not guaranteed to occur when the prediction error is split into four 2 *b* data parts. The F_1 transform was observed to always provide the best or near-best performance. Table 11.18 lists the lowest estimated Markov entropy values with PSV=7 prediction for the test-image set used, together with their bit rates obtained with the enhanced JBIG scheme (PSV7- F_1 -JBIG) and the zeroth-order entropies of the PSV=7 prediction error images. It is seen that the higher-order entropy values pro-




FIGURE 11.40

Markov entropy values up to the maximum order possible with error limit $\epsilon = 0.05$ for four forms of splitting, for the 512×512 , 8 *b/pixel* Airplane image; see also Table 11.17. *Note:* Order=0 indicates memoryless entropy. Reproduced with permission from L. Shen and R.M. Rangayyan, “Lossless compression of continuous-tone images by combined inter-bit-plane decorrelation and JBIG coding”, *Journal of Electronic Imaging*, 6(2): 198 – 207, 1997. © SPIE and IS&T.

TABLE 11.17

Estimated Values of the Highest Possible Order of Markov Entropy ($b/pixel$) for the Airplane Image.

Data	Code/ transform	One part 8 b	Two parts 4 $b/part$	Four parts 2 $b/part$	Eight parts 1 $b/part$
Original	Binary	4.06	4.21	4.52	5.10
Airplane image	Gray	4.06	4.00	4.02	4.21
	F_1	4.06	4.26	4.43	4.93
	F_2	4.06	4.21	4.50	4.97
PSV=7 prediction error of Airplane	Binary	3.75	4.21	5.04	6.95
	Gray	3.75	3.63	3.58	3.68
	F_1	3.75	3.61	3.46	3.60
	F_2	3.75	3.63	3.54	3.61

Values are shown with and without prediction, combined with four different code representation (transformation) schemes, and with error limit $\epsilon = 0.05$. Reproduced with permission from L. Shen and R.M. Rangayyan, "Lossless compression of continuous-tone images by combined inter-bit-plane decorrelation and JBIG coding", *Journal of Electronic Imaging*, 6(2): 198 – 207, 1997. © SPIE and IS&T.




Figure 11.41 (a)




Figure 11.41 (b)




Figure 11.41 (c)

vide lower estimates of the bit-rate limit than the zeroth-order entropies. The average entropy value decreases to 4.52 from 4.88 $b/pixel$ with higher-order entropy estimation while using binary representation; by using the F_1 transform instead of binary representation of the error data, the average Markov entropy value further reduces to 4.15 $b/pixel$.

The disadvantage of using the zeroth-order entropy to measure the performance of a data compression algorithm is clearly shown in Table 11.18: the enhanced JBIG (PSV7-F1-JBIG) coding scheme achieves an average bit rate of 4.70 $b/pixel$ compared with the average zeroth-order entropy of 4.88 $b/pixel$ for the prediction error images. Considering the higher-order entropy values shown, it appears that the compression efficiency of the enhanced JBIG technique could be further improved. An important application of high-order entropy estimation could be to provide a potentially achievable lower bound of bit rate for an original or decorrelated image, if the high-order entropy is estimated with adequate accuracy.




FIGURE 11.41

Plots of the Markov entropy values up to the maximum order possible with error limit $\epsilon = 0.05$, with four forms of splitting, for PSV=7 prediction error of the 512×512 , 8 b/pixel Airplane image, with (a) Binary representation; (b) Gray representation; (c) F_1 transformation; and (d) F_2 transformation. *Note:* Order=0 indicates memoryless entropy. Reproduced with permission from L. Shen and R.M. Rangayyan, "Lossless compression of continuous-tone images by combined inter-bit-plane decorrelation and JBIG coding", *Journal of Electronic Imaging*, 6(2): 198 – 207, 1997. © SPIE and IS&T.

TABLE 11.18

Lowest Estimated Markov Entropy Values with PSV=7 Prediction for Eight 8 *b* Test Images.

8 <i>b</i> image (columns × rows)	JPEG with PSV = 7			
	Bit rate		Lowest entropy	
	H_{e0}	(EJBIG)	F_1	Binary
Airplane (512×512)	4.18	3.83	3.46	3.75
Baboon (512×512)	6.06	6.04	5.39	5.77
Cameraman (256×256)	4.90	4.67	3.79	4.28
Lenna-256 (256×256)	5.15	4.94	4.16	4.54
Lenna-512 (512×512)	4.65	4.45	4.09	4.41
Peppers (512×512)	4.66	4.54	4.10	4.43
Sailboat (512×512)	5.14	5.05	4.51	4.91
Tiffany (512×512)	4.27	4.11	3.74	4.05
Average	4.88	4.70	4.15	4.52

Also shown are bit rates via enhanced JBIG bit-plane coding of F_1 -transformed PSV=7 prediction error (PSV7-F1-JBIG or EJBIG) and the zeroth-order entropies (H_{e0}) of the PSV=7 prediction error images (in *b/pixel*). Reproduced with permission from L. Shen and R.M. Rangayyan, “Lossless compression of continuous-tone images by combined inter-bit-plane decorrelation and JBIG coding”, *Journal of Electronic Imaging*, 6(2): 198 – 207, 1997. © SPIE and IS&T.

11.14 Application: Teleradiology

Teleradiology is commonly defined as the practice of radiology at a distance [338, 1036, 1037, 1038, 1039, 1040, 1041]. Teleradiology offers a technological approach to the problem of eliminating the delay in securing the consultation of a radiologist to patients in rural and remote areas. The timely availability of radiological diagnosis via telecommunication could potentially reduce the morbidity, mortality, and costs of transportation to tertiary health-care centers of patients in remotely situated areas, and in certain situations, in developing countries as well.

In the military environment, a study in 1983 [1042] indicated that over 65% of the medical facilities with radiographic equipment had no radiologists assigned to them, and an additional 15% had only one radiologist. In such cases, teleradiology could be a vehicle for redistributing the image-reading workload from under-staffed sites to more adequately staffed central locations [1043].

According to a study conducted in 1989, the province of Alberta, Canada, had a total of 130 health-care centers with radiological imaging facilities, out of which only 30 had resident radiologists [1044]. Sixty-one of the other centers depended upon visiting radiologists. The remaining 39 centers used to send their radiographs to other centers for interpretation, with a delay of 3 – 14 days in receiving the results [1044]. The situation was comparable in the neighboring provinces of Saskatchewan and Manitoba, and it was observed that the three provinces could benefit significantly from teleradiology. Even in the case of areas served by contract radiologists, teleradiology can permit evaluation and consultation by other radiologists at tertiary health-care centers in emergency situations as well as in complicated cases.

Early attempts at teleradiology systems [1045, 1046] consisted of analog transmission of slow-scan TV signals over existing telephone lines, ultra-high-frequency (UHF) radiolinks, and other such analog channels [1037]. Analog transmission and the concomitant slow transmission rates were satisfactory for low-resolution images such as nuclear medicine images. However, the transmission times were prohibitively high for high-resolution images, such as chest radiographs. Furthermore, the quality of the images received via analog transmission is a function of the distance, which could result in an unpredictable performance of radiologists with the received images. Thus, the natural progression of teleradiology systems was toward digital transmission. The initial choice of the transmission medium was the ordinary telephone line, operating at 300–1,200 *bps* (bits per second). Several commercial teleradiology systems were based upon the use of telephone lines for data transmission. Improvements in modem technology allowing transmission speeds of up to 19.2 *Kbps* over standard telephone lines, and the establishment of a number of 56 *Kbps* lines for commercial use by telephone companies made this medium viable for low-resolution images [1047].

The major reason for users' reluctance in accepting early teleradiology systems was the inability to meet the resolution of the original film. Spatial resolution of even $2,048 \times 2,048$ pixels was found to be inadequate to capture the sub-millimeter features found in chest radiographs and mammograms [1048]. It was recommended that spatial resolution of the order of $4,096 \times 4,096$ pixels, with at least 1,024 shades of gray, would be required to capture accurately the diagnostic information on radiographic images of the chest and breast. This demand led to the development of high-resolution laser digitizers capable of digitizing X-ray films to images of the order of $4,096 \times 4,096$ pixels, with 12 *b/pixel*, by the mid 1990s. Imaging equipment capable of direct digital acquisition of radiographic images to the same resolution as above were also developed in the late 1990s. Teleradiology system designers were then faced with the problem of dealing with the immense amount of data involved in such high-resolution digital images. The transmission of such large amounts of data over ordinary telephone lines involved large delays, which could be overcome to some extent by using parallel lines for increased data transfer rate [1037]. The use of satellite channels was also an option to speed up image data transmission [1049], but problems associated with image data management and archival hindered the anticipated widespread acceptance of high-resolution teleradiology systems. Such difficulties motivated advanced research into image data compression and encoding techniques.

The development of PACS and teleradiology systems share some historical common ground. Although delayed beyond initial predictions, both PACS and teleradiology established their presence and value in clinical practice by the late 1990s. The following paragraphs provide a historical review of teleradiology [338, 1041].

11.14.1 Analog teleradiology

The first instance of transmitting pictorial information for medical diagnosis dates back to 1950 when Gershon-Cohen and Cooley used telephone lines, and a facsimile system adapted to convert medical images into video signals, for transmitting images between two hospitals 45 *km* apart in Philadelphia, PA [1050]. In a pioneering project in 1959, Jutras [1051] conducted what is perhaps the first teleradiology trial, by interlinking two hospitals, 8 *km* apart, in Montreal, Québec, Canada, using a coaxial cable to transmit telefluoroscopy examinations. The potential of teleradiology in the provision of the services of a radiologist to remotely situated areas, and in the redistribution of radiologists' workload from under-staffed centers to more adequately staffed centers was immediately recognized, and a number of clinical evaluation projects were conducted [1037, 1045, 1046, 1052, 1053, 1054, 1055, 1056, 1057]. Most of the early attempts consisted of analog transmission of medical images via standard telephone lines, dedicated coaxial cables, UHF radio, microwave, and satellite channels, with display on TV monitors at the receiving terminal. James et al. [1057] give a review of the results of the early experiments. Andrus and

Bird [1036] describe the concept of a teleradiology system in which the radiologist, stationed at a medical center, controls a video camera to zoom in on selected areas of interest of an image at another site located far away, and observes the results in real time on a TV screen. Steckel [1058] conducted experiments with a system using an 875-line closed-circuit TV system for transmitting radiographic images within a hospital for educational purposes, and concluded that the system's utility far outweighed disadvantages such as the inability to view a sequence of images belonging to a single study.

In 1972, Webber and Corbus [1045] used existing telephone lines and slow-scan TV for transmitting radiographs and nuclear medicine images. The resolution achieved was satisfactory for nuclear medicine images, but both the spatial resolution and the gray-scale dynamic range (radiometric resolution) were found to be inadequate for radiographs. A similar experiment using telephone lines and slow-scan TV by Jelasco et al. [1046] resulted in 80% correct interpretation of radiographs. Other experiments with slow-scan TV over telephone lines [1057] demonstrated the inadequacy of this medium, and also that the diagnostic accuracy with such systems varied with the nature of the images.

Webber et al. [1054] used UHF radio transmission, in 1973, for transmitting nuclear medicine images and radiographs. While the system worked satisfactorily for nuclear medicine images, evaluation of chest X-ray images needed zoom and contrast manipulation of the TV monitor. Murphy et al. [1053] used a microwave link for the transmission of images of chest radiographs acquired with a remotely controlled video camera, over a distance of about 4 km, and indicated that it would be an acceptable method for providing health care to people in remote areas.

Andrus et al. [1052] transmitted X-ray images of the abdomen, chest, bone, and skull over a 45 km round loop, using a 4 MHz, 512-line TV channel including three repeater stations. The TV camera was remotely controlled using push buttons and a joystick to vary the zoom, aperture, focus, and direction of the camera. It was concluded that the TV interpretations were of acceptable accuracy. Such real-time operation calls for special skills on the part of the radiologist, requires coordination between the operator at the image acquisition site and the radiologist at the receiving center, and could take up a considerable amount of the radiologist's valuable time. Moreover, practical microwave links exist only between and within major cities, and cannot serve the communication needs of teleradiology terminals in rural and remote areas. In addition, the operating costs over the duration of interactive manipulations could reach high levels, and render such a scheme uneconomical.

In 1973, Lester et al. [1055] used a satellite (ATS-1) for analog transmission of video-taped radiologic information, and concluded that satisfactory radiographic transmission is possible "if a satisfactory sensor of radiographic images were constructed." In 1979, Carey et al. [1056] reported on the results of an analog teleradiology experiment using the Hermes spacecraft. They reported the effectiveness of TV fluoroscopy to be 90% of that with conventional

procedures. Page et al. [1059] used a two-way analog TV network with the Canadian satellite ANIK-B to transmit radiographic images from northern Québec to Montréal, and reported an initial accuracy in TV interpretation of 81% with respect to film reading. The accuracy rose to 94% after a 3-month training of the participant radiologists in the use of the TV system. The noise associated with analog transmission, the low resolution of the TV monitors used, and the requirement on the part of the radiologists to participate in real-time control of the image-acquisition cameras made the concept of TV transmission of radiographic images unacceptable. Furthermore, the noise associated with analog transmission is dependent upon the distance. Not surprisingly, James et al. [1057] reported that their teleradiology system, transmitting emergency department radiographs via a satellite channel from a local TV studio, was unacceptable due to a decrease in the accuracy of image interpretation to about 86% with respect to that with standard protocols.

11.14.2 Digital teleradiology

Given the advantages of digital communication over analog methods [1060], the natural progression of teleradiology was toward the use of digital data transmission techniques. The advent of a number of digital medical imaging modalities facilitated this trend [1061, 39, 1062]. Digital imaging also allowed for image processing, enhancement, contrast scaling, and flexible manipulation of images on the display monitors after acquisition. Many of the initial attempts at digital teleradiology [1047, 1063, 1064, 1065, 1066, 1067] were based upon microcomputers and used low-resolution digitization, display, and printers. The resolution was of the order of 256×256 to 512×512 pixels with 256 shades of gray, mostly because of the unavailability of high-resolution equipment. Gayler et al. [1063] described a laboratory evaluation of such a microcomputer-based teleradiology system, based upon a $512 \times 512 \times 8 b$ format for image acquisition and display, and evaluated radiologists' performance with routine radiographs. They found the diagnostic performance to be significantly worse than that using the original film radiographs. Nevertheless, they concluded that microcomputer-based teleradiology systems "warrant further evaluation in a clinical environment."

In 1982, Rasmussen et al. [1067] compared the performance of radiologists with images transmitted by analog and digital means and light-box viewing of the original films. The resolution of digitization used was 512×256 pixels with $6 b/\text{pixel}$. The digital images were converted to analog signals for analog transmission. It was concluded that the resolution used would provide satisfactory radiographic images for gross pathological disorders, but that subtle features would require higher resolution.

Gitlin [1065], Curtis et al. [1066], and Skinner et al. [1068] followed the laboratory evaluation of Gayler et al. [1063] with field trials using standard telephone lines at 9,600 *bps* for the transmission of $512 \times 512 \times 8 b$ images from five medical-care facilities to a central hospital in Maryland. A relative

accuracy of 97% with video-image readings was reported [1066], as compared to standard film interpretation. This was a substantially higher accuracy than that obtained in a preceding laboratory study [1063]; the improvement was attributed to the larger percentage of normal images used in the field trial, and to the higher experience of the analysts in clinical radiology.

In a field trial in 1984, Gitlin [1065] used a $1,024 \times 1,024$ matrix of pixels, 9,600 *bps* telephone lines, and lossy data compression to bring down the transmission times. A relative accuracy with video interpretation of 87%, with respect to standard film interpretation, was observed. The relative accuracy was observed to be dependent upon the type of data compression used, among other factors.

Gordon et al. [1069] presented an analysis of a number of scenarios and tradeoffs for practical implementation of digital teleradiology. In related papers, Rangayyan and Gordon [1070] and Rangaraj and Gordon [1071] discussed the potential for providing advanced imaging services such as CT through teleradiology.

In 1987, DiSantis et al. [1072] digitized excretory urographs to $1,024 \times 1,024$ matrices, and transmitted the images over standard telephone lines, after data compression, to a receiving unit approximately 3 *km* away. A panel of three radiologists interpreted the images on video monitors, and the results were compared with the original film readings performed about a week earlier. An agreement of 93% was found between the film and video readings in the diagnosis of obstructions. However, only 64% of urethral calculi detected with the original radiographs were also detected with the video images. This result demonstrated clearly that, whereas a resolution of $1,024 \times 1,024$ pixels could be adequate for certain types of diagnosis, higher resolution is required for capturing all of the diagnostic information that could be present on the original film.

In 1987, Kagetsu et al. [1064] reported on the performance of a commercially available teleradiology system using $512 \times 512 \times 8$ *b* images and transmission over 9,600 *bps* standard telephone lines after 2.5:1 data compression. Experiments were conducted with a wide variety of radiographs over a four-month period. An overall relative accuracy of 89% was reported, between the received images on video display and the original films. Based on these results, Kagetsu et al. recommended a review of the original films at some later date because of the superior spatial and contrast resolution of film.

Several commercial systems were released for digital teleradiology in the late 1980s. Although such systems were adequate for handling low-resolution images in CT, MRI, and nuclear medicine, they were not suitable for handling large-format images such as chest radiographs and mammograms. Experiments with such systems demonstrated the inadequacy of low-resolution digital teleradiology systems as an alternative to the physical transportation of the films or the patients to centers with radiological diagnostic facilities. Although higher resolution was required in the digitized images, the substantial increase in the related volume of data and the associated difficulties remained

a serious concern. Furthermore, the use of lossy data compression schemes to remain within the data-rate limitation of telephone lines and other low-speed communication channels was observed to be unacceptable.

11.14.3 High-resolution digital teleradiology

The development of high-resolution image digitizers and display equipment, and the routine utilization of high-data-rate communication media, paved the way for high-resolution digital teleradiology. In 1989, Carey et al. [1049] reported on the performance of the DTR-2000 teleradiology system from DuPont, consisting of a $1,684 \times 2,048$ -pixel laser digitizer with 4,096 quantization levels, a T1 satellite transmission channel (at 1.544 *Mbps*), and a DuPont laser film recorder with 256 possible shades of gray. A nonlinear mapping was performed from the original 4,096 quantization levels to 256 levels on the film copy to make use of the fact that the eye is more sensitive to contrast variations at lower density. With this mapping, at the lower end of the gray scale, small differences in gray values correspond to larger differences in optical densities than at the higher end of the gray scale. Thus, the overall optical density range of the film is much larger than can be obtained by linear mapping. Carey et al. [1049] transmitted radiographic and ultrasonographic images over the system from Seaforth to London, in Ontario, Canada, and reported a relative accuracy of 98% in reading the laser-sensitive film as compared to the original film. It was concluded that the laser-sensitive film "clearly duplicated the original film findings." However, they also reported "contouring" on the laser-sensitive film, which might have been due to the nonlinear mapping of the 4,096 original gray levels to 256 levels on the film. Certain portions of the original gray scale with rapidly changing gray levels could have been mapped into the same optical density on the film, giving rise to contouring artifacts.

Barnes et al. [1047] suggested that the challenge of integrating the increasing number of medical imaging technologies could be met by networked multimodality imaging workstations. Cox et al. [1073] compared images digitized to $2,048 \times 2,048 \times 12$ *b* and displayed on monitors with $2,560 \times 2,048 \times 8$ *b* pixels, digital laser film prints, and conventional film. They reported significant differences in the performance of the three display formats: digital hard copy performed as well as or better than conventional film, whereas the interactive display failed to match the performance of the other two. They suggested that although the differences could be eliminated by training the personnel in reading from displays and by using image enhancement techniques, it was premature to conclude either way.

In 1990, Batnitzky et al. [1074] conducted an assessment of the then-available technologies for film digitization, display, generation of hard copy, and data communication for application in teleradiology systems. They concluded that $2,048 \times 2,048 \times 12$ *b* laser digitizers, displays with scan lines of $1,024 - 2,048, 8 - 12$ *b/pixel*, hard copiers that interpolate $2,048 \times 2,048$

matrices to $4,096 \times 4,096$ matrices, and the merger of computer and communication technologies resulting in flexible wide-area networks, had paved the way for the acceptance of "final interpretation teleradiology," completely eliminating the need to go back to the original films. Gillespy et al. [1075] described the installation of a DuPont Clinical Review System, consisting of a laser film digitizer with $1,680 \times 2,048 \times 12 b$ pixels, and a $1,024 \times 840 \times 12 b$ display unit, and reported that "clinicians were generally satisfied with the unit." Several studies on the contrast and resolution of high-resolution digitizers [1076, 1077, 1078] demonstrated that the resolution of the original film was maintained in the digitized images.

Several systems are now available for digital teleradiology, including high-resolution laser digitizers that can provide images of the order of $4,000 \times 5,000 \times 12 b$ pixels with a spatial resolution of $50 \mu m$ or better; high-luminance monitors that can display up to $2,560 \times 2,048$ pixels at $12 b/pixel$ and non-interlaced refresh rates of $70 - 80 fps$; and laser-film recorders with a spatial resolution of $50 \mu m$ that can print images of size $4,000 \times 5,000 \times 12 b$ pixels. Satellite, cable, or fiber-optic transmission equipment and channels may be leased with transmission rates of several *Mbps*. However, the large amount of data related to high-resolution images can create huge demands in data transmission and archival capacity. Lossless data compression techniques can bring down the amount of data, and have a significant impact on the practical implementation of teleradiology and related technologies.

The introduction of data compression, encoding, and decoding in digital teleradiology systems raises questions on the overall throughput of the system in transmission and reception, storage and retrieval of image data, patient confidentiality, and information security. The compression of image data removes the inherent redundancy in images, and makes the data more sensitive to errors [1079]. In dedicated communication links, appropriate error control should be provided for detecting and correcting such errors. In the case of packet-switched communication links, the removal of redundancy by data compression could result in increased retransmission overhead. However, with sophisticated digital communication links operating at typical bit-error rates of 1 in 10^9 , and channel utilization (throughput) efficiency of about 97% using high-level packet-switched protocols [1080], the advantages of data compression far outweigh the overheads due to the reasons mentioned above.

High-resolution digital teleradiology is now feasible without any sacrifice in image quality, and can serve as an alternative to transporting patients or films. Distance should no longer be a limitation in providing reliable diagnostic service by city-based expert radiologists to patients in remote or rural areas.

11.15 Remarks

Lossless data compression is desirable in medical image archival and transmission. In this chapter, we studied several lossless data compression techniques. A lossless compression scheme generally consists of two steps: decorrelation and encoding. The success of a lossless compression method is mainly based upon the efficiency of the decorrelation procedure used. In practice, a decorrelation procedure could include several cascaded decorrelation blocks, each of which could accomplish a different type of decorrelation and facilitate further decorrelation by the subsequent blocks. Some of the methods described in this chapter illustrate creative ways of combining multiple decorrelators with different characteristics for achieving better compression efficiency.

Several information-theoretic concepts and criteria as applicable to data compression were also discussed in this chapter. A practical method for the estimation of high-order entropy was presented, which could aid in the lower-limit analysis of lossless data compression. High-order entropy estimation could aid in the design, analysis, and evaluation of cascaded decorrelators.

A historical review of selected works in the development of PACS and teleradiology systems was presented in the concluding section, in order to demonstrate the need for image compression and data transmission in a practical medical application. PACS, teleradiology, and telemedicine are now well established areas that are providing advanced technology for improved health care [1039, 1040, 1081, 1082, 1083].

11.16 Study Questions and Problems

Selected data files related to some of the problems and exercises are available at the site

www.enel.ucalgary.ca/People/Ranga/enel697

1. The probabilities of occurrence of eight symbols are given as (0.1, 0.04, 0.3, 0.15, 0.03, 0.2, 0.15, 0.03).
Derive the Huffman code for the source.
2. For a 4-symbol source, derive all possible sets of the Huffman code. (The exact PDF of the source is not relevant.)
Create a few strings of symbols and generate the corresponding Huffman codes. Verify that the result satisfies the basic requirements of a code, including unique and instantaneous decodability.
3. For the 4×4 image shown below, design a Huffman coding scheme. Show all steps of your design.

Compute the entropy of the image and the average bit rates using direct binary coding and Huffman coding.

$$\begin{bmatrix} 1 & 0 & 2 & 2 \\ 0 & 2 & 1 & 1 \\ 3 & 3 & 2 & 2 \\ 2 & 2 & 2 & 2 \end{bmatrix}. \quad (11.206)$$

4. Discuss the similarities and differences between the Karhunen–Loève, discrete Fourier, and the Walsh-Hadamard transforms. Discuss the particular aspects of each transform that are of importance in transform-domain coding for image data compression.
5. For the 5×5 image shown below, compute the bit rate using
 - (a) direct binary coding;
 - (b) horizontal run-length coding, and
 - (c) predictive coding (or DPCM) using the model $\hat{f}(m, n) = f(m, n - 1)$.
 Show and explain all steps. State your assumptions, if any, and explain your procedures.

$$\begin{bmatrix} 121 & 125 & 119 & 121 & 121 \\ 121 & 119 & 125 & 125 & 125 \\ 126 & 126 & 126 & 126 & 126 \\ 130 & 135 & 135 & 135 & 135 \\ 129 & 129 & 129 & 129 & 129 \end{bmatrix}. \quad (11.207)$$

6. For the 3×3 image shown below, prepare the bit planes using the direct binary and Gray codes. Examine the bit planes for the application of run-length coding. Which code can provide better compression? Explain your observations and results.

$$\begin{bmatrix} 0 & 2 & 2 \\ 2 & 1 & 1 \\ 3 & 2 & 2 \end{bmatrix}. \quad (11.208)$$

11.17 Laboratory Exercises and Projects

1. Write a program (in C, C++, or MATLAB) to compute the histogram and the zeroth-order entropy of a given image. Apply the program to a few images in your collection. Study the nature of the histograms and relate their characteristics as well as entropy to the visual features present in the corresponding images.
2. Write a program to compute the zeroth-order and first-order entropy of an image considering pairs of gray-level values. Apply the program to a few images in your collection and analyze the trends in the zeroth-order and first-order entropy values.

What are the considerations, complexities, and limitations involved in computing entropy of higher orders?

3. For the image in the file `RajREye.dat` with 3 b/pixel , create bit planes using (a) the binary code, and (b) the Gray code. Compute the entropy of each bit plane. Compute the average entropy over all of the bit planes for each code. What is the expected trend?

Do your results meet your expectations? Explain.

4. Develop a program to compute the DFT of an image. Write steps to compute the energy contained in concentric circles or squares of several sizes spanning the full spectrum of the image and to plot the results.

Apply the program to a few images in your collection. Relate the spectral energy distribution to the visual characteristics of the corresponding images. Discuss the relevance of your findings in data compression.

5. Develop a program to compute the error of prediction based upon a few simple predictors, such as

(a) $\tilde{f}(m, n) = f(m - 1, n)$.

(b) $\tilde{f}(m, n) = f(m, n - 1)$.

(c) $\tilde{f}(m, n) = [f(m - 1, n) + f(m, n - 1) + f(m - 1, n - 1)]/3$.

Derive the histograms and the entropies of the original image and the error of prediction for a few images with each of the predictors listed above. Evaluate the results and comment upon the relevance of your findings in image coding and data compression.

Pattern Classification and Diagnostic Decision

The final purpose of biomedical image analysis is to classify a given image, or the features that have been detected in the image, into one of a few known categories. In medical applications, a further goal is to arrive at a diagnostic decision regarding the condition of the patient. A physician or medical specialist may achieve this goal via visual analysis of the image and data presented: comparative analysis of the given image with others of known diagnoses or the application of established protocols and sets of rules assist in such a decision-making process. Images taken earlier of the same patient may also be used, when available, for comparative or differential analysis. Some measurements may also be made from the given image to assist in the analysis. The basic knowledge, clinical experience, expertise, and intuition of the physician play significant roles in this process.

When image analysis is performed via the application of computer algorithms, the typical result is the extraction of a number of numerical features. When the numerical features relate directly to measurements of organs or features represented by the image — such as an estimate of the size of the heart or the volume of a tumor — the clinical specialist may be able to use the features directly in his or her diagnostic logic. However, when parameters such as measures of texture and shape complexity are derived, a human analyst is not likely to be able to analyze or comprehend the features. Furthermore, as the number of the computed features increases, the associated diagnostic logic may become too complicated and unwieldy for human analysis. Computer methods would then be desirable to perform the classification and decision process.

At the outset, it should be borne in mind that a biomedical image forms but one piece of information in arriving at a diagnosis: the classification of a given image into one of many categories may assist in the diagnostic procedure, but will almost never be the only factor. Regardless, pattern classification based upon image analysis is indeed an important aspect of biomedical image analysis, and forms the theme of the present chapter. Remaining within the realm of CAD as introduced in Figure 1.33 and Section 1.11, it would be preferable to design methods so as to *aid* a medical specialist in arriving at a diagnosis rather than to provide a decision.

A generic problem statement for pattern classification may be expressed as follows: *A number of measures and features have been derived from a biomedical image. Develop methods to classify the image into one of a few specified categories. Investigate the relevance of the features and the classification methods in arriving at a diagnostic decision about the patient.*

Observe that the features mentioned above may have been derived manually or by computer methods. Recognize the distinction between classifying the given image and arriving at a diagnosis regarding the patient: the connection between the two tasks or steps may not always be direct. In other words, a pattern classification method may facilitate the labeling of a given image as being a member of a particular class; arriving at a diagnosis of the condition of the patient will most likely require the analysis of several other items of clinical information. Although it is common to work with a prespecified number of pattern classes, many problems do exist where the number of classes is not known *a priori*. A special case is screening, where the aim is to simply decide on the presence or absence of a certain type of abnormality or disease. The initial decision in screening may be further focused on whether the subject appears to be free of the specific abnormality of concern or requires further investigation.

The problem statement and description above are rather generic. Several considerations arise in the practical application of the concepts mentioned above to medical images and diagnosis. Using the detection of breast cancer as an example, the following questions illustrate some of the problems encountered in practice.

- Is a mass or tumor present? (Yes/No)
 - If a mass or tumor is present
 - Give or mark its location.
 - Compare the density of the mass to that of the surrounding tissues: hypodense, isodense, hyperdense.
 - Describe the shape of its boundary: round, ovoid, irregular, macrolobulated, microlobulated, spiculated.
 - Describe its texture: homogeneous, heterogeneous, fatty.
 - Describe its edge: sharp (well-circumscribed), ill-defined (fuzzy).
 - Decide if it is a benign mass, a cyst (solid or fluid-filled), or a malignant tumor.
- Are calcifications present? (Yes/No)
 - If calcifications are present:
 - Estimate their number per cm^2 .
 - Describe their shape: round, ovoid, elongated, branching, rough, punctate, irregular, amorphous.

- Describe their spatial distribution or cluster.
- Describe their density: homogeneous, heterogeneous.
- Are there signs of architectural distortion? (Yes/No)
- Are there signs of bilateral asymmetry? (Yes/No)
- Are there major changes compared to the previous mammogram of the patient?
- Is the case normal? (Yes/No)
If the case is abnormal:
 - Is the disease benign or malignant (cancer)?

The items listed above give a selection of the many features of mammograms that a radiologist would investigate; see Ackerman et al. [1084] and the BI-RADSTM manual [403] for more details. Figure 12.1 shows a graphical user interface developed by Alto et al. [528, 1085] for the categorization of breast masses related to some of the questions listed above. Figure 12.2 illustrates four segments of mammograms demonstrating masses and tumors of different characteristics, progressing from a well-circumscribed and homogeneous benign mass to a highly spiculated and heterogeneous tumor.

The subject matter of this book — image analysis and pattern classification — can provide assistance in responding to only some of the questions listed above. Even an entire set of mammograms may not lead to a final decision: other modes of diagnostic imaging and means of investigation may be necessary to arrive at a definite diagnosis.

In the following sections, a number of methods for pattern classification, decision making, and evaluation of the results of classification are reviewed and illustrated.

(*Note:* Parts of this chapter are reproduced, with permission, from R.M. Rangayyan, *Biomedical Signal Analysis: A Case-Study Approach*, IEEE Press and Wiley, New York, NY. 2002, © IEEE.)

12.1 Pattern Classification

Pattern recognition or classification may be defined as the categorization of the input data into identifiable classes via the extraction of significant features or attributes of the data from a background of irrelevant detail [402, 721, 1086, 1087, 1088, 1089, 1090]. In biomedical image analysis, after quantitative features have been extracted from the given images, each image (or ROI) may be represented by a feature vector $\mathbf{x} = [x_1, x_2, \dots, x_n]^T$, which is also known

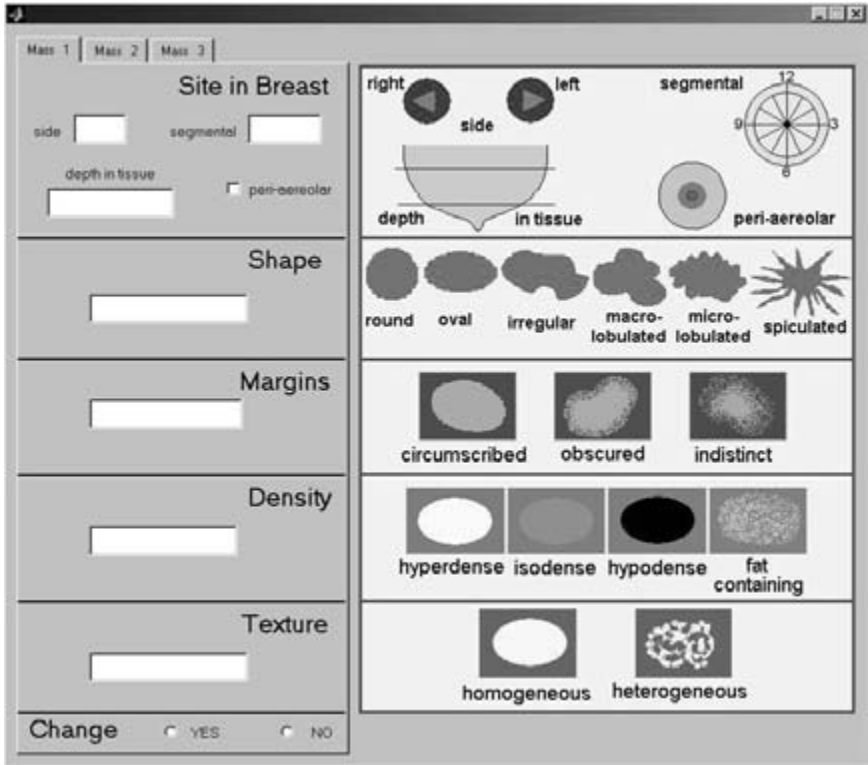


FIGURE 12.1

Graphical user interface for the categorization of breast masses. Reproduced with permission from H. Alto, R.M. Rangayyan, R.B. Paranjape, J.E.L. Desautels, and H. Bryant, "An indexed atlas of digital mammograms for computer-aided diagnosis of breast cancer", *Annales des Télécommunications*, 58(5): 820 – 835, 2003. © GET – Lavoisier. Figure courtesy of C. LeGuillou, École Nationale Supérieure des Télécommunications de Bretagne, Brest, France.




FIGURE 12.2

Examples of breast mass regions and contours with the corresponding values of fractional concavity f_{cc} , spiculation index SI , compactness cf , acutance A , and sum entropy F_8 . (a) Circumscribed benign mass. (b) Macrolobulated benign mass. (c) Microlobulated malignant tumor. (d) Spiculated malignant tumor. Note that the masses and their contours are of widely differing size, but have been scaled to the same size in the illustration. The first letter of the case identifier indicates a malignant diagnosis with ‘m’ and a benign diagnosis with ‘b’ based upon biopsy. The symbols after the first numerical portion of the identifier represent l: left, r: right, c: cranio-caudal view, o: medio-lateral oblique view, x: axillary view. The last two digits represent the year of acquisition of the mammogram. An additional character of the identifier after the year (a – f), if present, indicates the existence of multiple masses visible in the same mammogram. Reproduced with permission from H. Alto, R.M. Rangayyan, and J.E.L. Desautels, “Content-based retrieval and analysis of mammographic masses”, *Journal of Electronic Imaging*, in press, 2005. © SPIE and IS&T.

as the measurement vector or a pattern vector. When the values x_i are real numbers, \mathbf{x} is a point in an n -dimensional Euclidean space: vectors of similar objects may be expected to form clusters as illustrated in Figure 12.3.

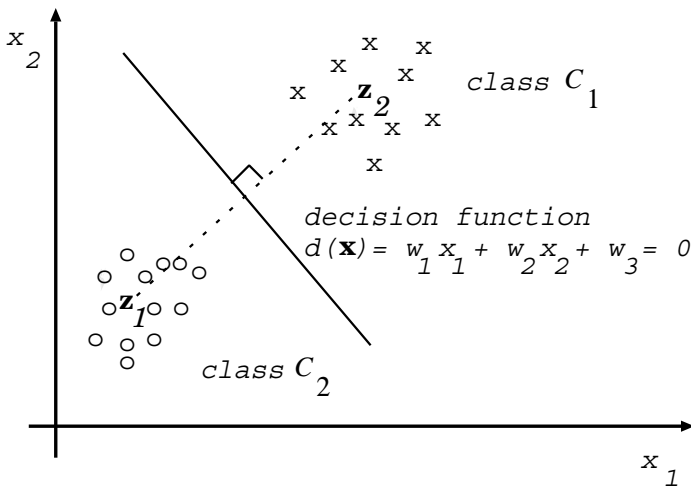


FIGURE 12.3

Two-dimensional feature vectors of two classes, C_1 and C_2 . The prototypes of the two classes are indicated by the vectors \mathbf{z}_1 and \mathbf{z}_2 . The linear decision function $d(\mathbf{x})$ shown (solid line) is the perpendicular bisector of the straight line joining the two prototypes (dashed line). Reproduced with permission from R.M. Rangayyan, *Biomedical Signal Analysis: A Case-Study Approach*, IEEE Press and Wiley, New York, NY, 2002, © IEEE.

For efficient pattern classification, measurements that could lead to disjoint sets or clusters of feature vectors are desired. This point underlines the importance of the appropriate design of the preprocessing and feature extraction procedures. Features or characterizing attributes that are common to all patterns belonging to a particular class are known as *intrasets* or *intra-class features*. Discriminant features that represent the differences between pattern classes are called *intersets* or *inter-class features*.

The pattern classification problem is that of generating optimal decision boundaries or decision procedures to separate the data into pattern classes based on the feature vectors provided. Figure 12.3 illustrates a simple linear decision function or boundary to separate 2D feature vectors into two classes.

12.2 Supervised Pattern Classification

The problem considered in supervised pattern classification may be stated as follows: *You are provided with a number of feature vectors with classes assigned to them. Propose techniques to characterize and parameterize the boundaries that separate the classes.*

A given set of feature vectors of known categorization is often referred to as a *training set*. The availability of a training set facilitates the development of mathematical functions that can characterize the separation between the classes. The functions may then be applied to new feature vectors of unknown classes to classify or recognize them. This approach is known as *supervised pattern classification*. A set of feature vectors of known categorization that is used to evaluate a classifier designed in this manner is referred to as a *test set*. After adequate testing and confirmation of the method with satisfactory results, the classifier may be applied to new feature vectors of unknown classes; the results may then be used to arrive at diagnostic decisions. The following subsections describe a few methods that can assist in the development of discriminant and decision functions.

12.2.1 Discriminant and decision functions

A general linear discriminant or decision function is of the form

$$d(\mathbf{x}) = w_1x_1 + w_2x_2 + \dots + w_nx_n + w_{n+1} = \mathbf{w}^T \mathbf{x}, \quad (12.1)$$

where $\mathbf{x} = [x_1, x_2, \dots, x_n, 1]^T$ is the feature vector augmented by an additional entry equal to unity, and $\mathbf{w} = [w_1, w_2, \dots, w_n, w_{n+1}]^T$ is a correspondingly augmented weight vector. A two-class pattern classification problem may be stated as

$$d(\mathbf{x}) = \mathbf{w}^T \mathbf{x} \begin{cases} > 0 & \text{if } \mathbf{x} \in C_1 \\ \leq 0 & \text{if } \mathbf{x} \in C_2 \end{cases}, \quad (12.2)$$

where C_1 and C_2 represent the two classes. The discriminant function may be interpreted as the boundary separating the classes C_1 and C_2 , as illustrated in Figure 12.3.

In the general case of an M -class pattern classification problem, we will need M weight vectors and M decision functions to perform the following decisions:

$$d_i(\mathbf{x}) = \mathbf{w}_i^T \mathbf{x} \begin{cases} > 0 & \text{if } \mathbf{x} \in C_i \\ \leq 0 & \text{otherwise} \end{cases}, \quad i = 1, 2, \dots, M, \quad (12.3)$$

where $\mathbf{w}_i = (w_{i1}, w_{i2}, \dots, w_{in}, w_{i,n+1})^T$ is the weight vector for the class C_i . Three cases arise in solving this problem [1086]:

Case 1: Each class is separable from the rest by a single decision surface:

$$\text{if } d_i(\mathbf{x}) > 0 \text{ then } \mathbf{x} \in C_i. \quad (12.4)$$

Case 2: Each class is separable from every other individual class by a distinct decision surface, that is, the classes are pairwise separable. There are $M(M-1)/2$ decision surfaces given by $d_{ij}(\mathbf{x}) = \mathbf{w}_{ij}^T \mathbf{x}$.

$$\text{if } d_{ij}(\mathbf{x}) > 0 \forall j \neq i \text{ then } \mathbf{x} \in C_i. \quad (12.5)$$

[Note: $d_{ij}(\mathbf{x}) = -d_{ji}(\mathbf{x})$.]

Case 3: There exist M decision functions $d_k(\mathbf{x}) = \mathbf{w}_k^T \mathbf{x}$, $k = 1, 2, \dots, M$, with the property that

$$\text{if } d_i(\mathbf{x}) > d_j(\mathbf{x}) \forall j \neq i, \text{ then } \mathbf{x} \in C_i. \quad (12.6)$$

This is a special instance of Case 2. We may define

$$d_{ij}(\mathbf{x}) = d_i(\mathbf{x}) - d_j(\mathbf{x}) = (\mathbf{w}_i - \mathbf{w}_j)^T \mathbf{x} = \mathbf{w}_{ij}^T \mathbf{x}. \quad (12.7)$$

If the classes are separable under Case 3, they are separable under Case 2; the converse is, in general, not true.

Patterns that may be separated by linear decision functions as above are said to be *linearly separable*. In other situations, an infinite variety of complex decision boundaries may be formulated by using generalized decision functions based upon nonlinear functions of the feature vectors as

$$d(\mathbf{x}) = w_1 f_1(\mathbf{x}) + w_2 f_2(\mathbf{x}) + \dots + w_K f_K(\mathbf{x}) + w_{K+1} \quad (12.8)$$

$$= \sum_{i=1}^{K+1} w_i f_i(\mathbf{x}). \quad (12.9)$$

Here, $\{f_i(\mathbf{x})\}$, $i = 1, 2, \dots, K$, are real, single-valued functions of \mathbf{x} ; $f_{K+1}(\mathbf{x}) = 1$. Whereas the functions $f_i(\mathbf{x})$ may be nonlinear in the n -dimensional space of \mathbf{x} , the decision function may be formulated as a linear function by defining a transformed feature vector $\mathbf{x}^\dagger = [f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_K(\mathbf{x}), 1]^T$. Then, $d(\mathbf{x}) = \mathbf{w}^T \mathbf{x}^\dagger$, with $\mathbf{w} = [w_1, w_2, \dots, w_K, w_{K+1}]^T$. Once evaluated, $\{f_i(\mathbf{x})\}$ is just a set of numerical values, and \mathbf{x}^\dagger is simply a K -dimensional vector augmented by an entry equal to unity. Several methods exist for the derivation of optimal linear discriminant functions [402, 738, 674].

Example of application: The ROIs of 57 breast masses are shown in Figure 12.4 arranged in the order of decreasing acutance A (see Sections 2.15, 7.9.2, and 12.12). Figure 12.5 shows the contours of the 57 masses arranged in the increasing order of fractional concavity f_{cc} (see Section 6.4). Most of the contours of the benign masses are seen to be smooth, whereas most of the contours of the malignant tumors are rough and spiculated. Furthermore, most of the benign masses have well-defined, sharp edges and are well-circumscribed, whereas the majority of the malignant tumors possess ill-defined and fuzzy borders. It is seen that the shape factor f_{cc} facilitates the ordering of the

contours in terms of shape complexity. However, the contours of a few benign masses and a few malignant tumors do not follow the expected trend. In addition, the acutance measure has lower values for most of the malignant tumors than for a majority of the benign masses.

The three shape factors cf , f_{cc} , and SI (see Chapter 6); the 14 texture measures as defined by Haralick [441, 442] (see Section 7.3); and four measures of edge sharpness as defined by Mudigonda et al. [165] (see Section 7.9.2) were computed for the ROIs and their contours. (*Note:* The factor SI was divided by two in this example to reduce it to the range $[0, 1]$.) Figure 12.6 gives a plot of the 3D feature-vector space (f_{cc}, A, F_8) for the 57 masses. The feature F_8 shows poor separation between the benign and malignant samples, whereas the feature A demonstrates some degree of separation. A scatter plot of the three shape factors (f_{cc}, cf, SI) of the 57 masses is given in Figure 12.7. Each of the three shape factors demonstrates high discriminant capability.

Figure 12.8 shows a 2D plot of the shape-factor vectors $[f_{cc}, SI]$ for a training set formed by selecting the vectors for 18 benign masses and 10 malignant tumors. The prototypes for the benign and malignant classes, obtained by averaging the vectors over all the members of the two classes in the training set, are marked as 'B' and 'M', respectively, on the plot. The solid straight line is the perpendicular bisector of the line joining the two prototypes (dashed line), and represents a linear discriminant function. The equation of the straight line is $SI + 0.6826 f_{cc} - 0.5251 = 0$. The decision function is represented by the following rule:

```

if  $SI + 0.6826 f_{cc} - 0.5251 < 0$  then
    benign mass
else
    malignant tumor
end.
```

It is seen that the rule given above will correctly classify all of the training samples.

Figure 12.9 shows the result of application of the linear discriminant function designed and shown in Figure 12.8 to a test set of 19 benign masses and 10 malignant tumors. The test set does not include any of the cases from the training set. It is seen that the classifier will lead to three false negatives in the test set.

12.2.2 Distance functions

Consider M pattern classes represented by their prototype patterns $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_M$. The prototype of a class is typically computed as the average of all the feature vectors belonging to the class. Figure 12.3 illustrates schematically the prototypes \mathbf{z}_1 and \mathbf{z}_2 of the two classes shown.

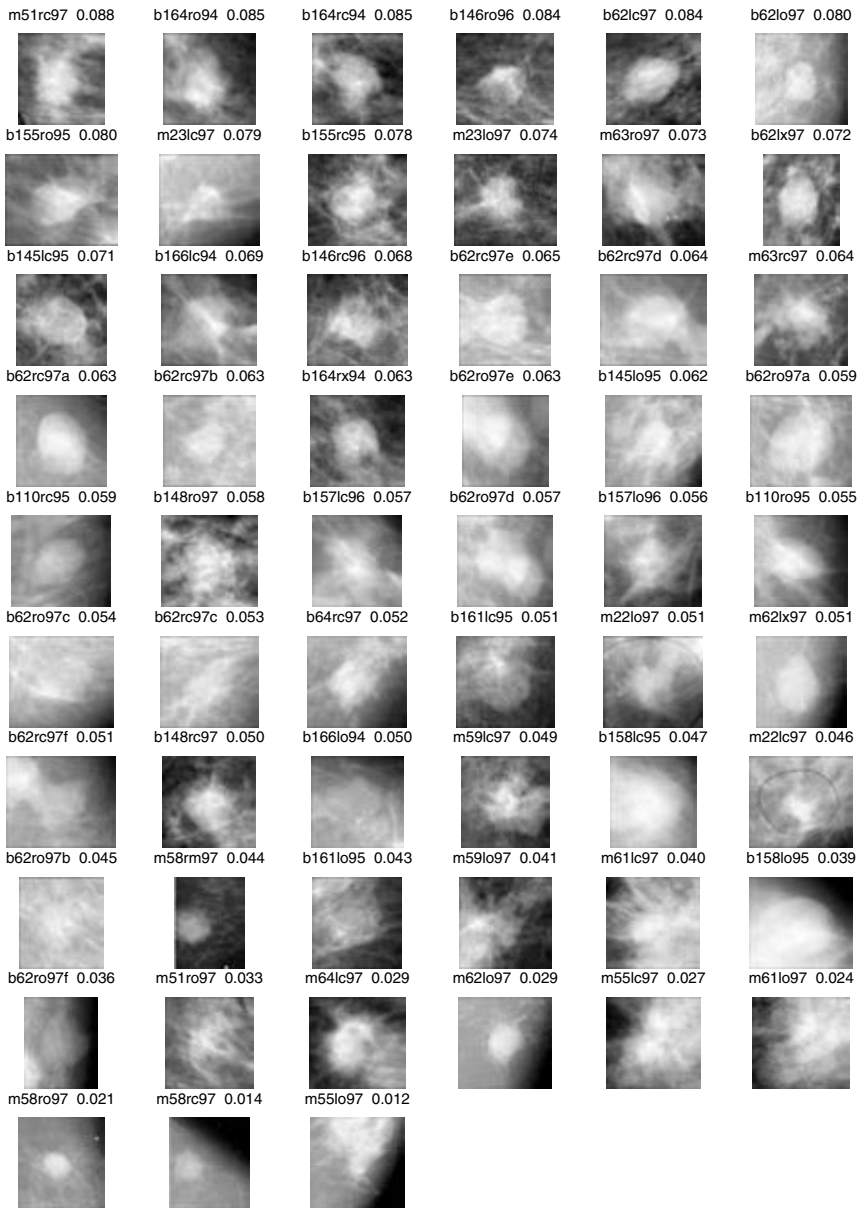


FIGURE 12.4

ROIs of 57 breast masses, including 37 benign masses and 20 malignant tumors. The ROIs are arranged in the order of decreasing acutance A . Note that the masses are of widely differing size, but have been scaled to the same size in the illustration. For details regarding the case identifiers, see Figure 12.2. Reproduced with permission from H. Alto, R.M. Rangayyan, and J.E.L. Desautels, "Content-based retrieval and analysis of mammographic masses", *Journal of Electronic Imaging*, in press, 2005.

© SPIE and IS&T.

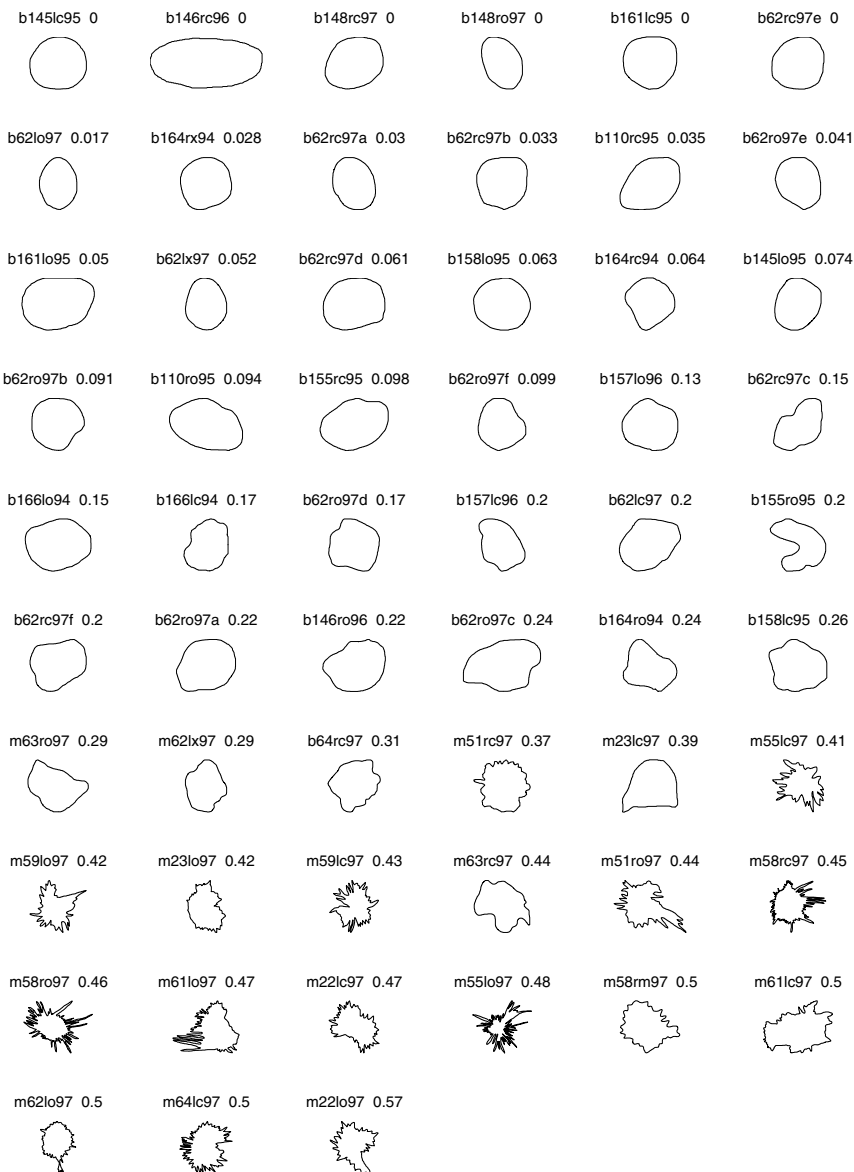


FIGURE 12.5

Contours of 57 breast masses, including 37 benign masses and 20 malignant tumors. The contours are arranged in the order of increasing f_{cc} . Note that the masses and their contours are of widely differing size, but have been scaled to the same size in the illustration. For details regarding the case identifiers, see Figure 12.2. See also Figure 12.30. Reproduced with permission from H. Alto, R.M. Rangayyan, and J.E.L. Desautels, “Content-based retrieval and analysis of mammographic masses”, *Journal of Electronic Imaging*, in press, 2005. © SPIE and IS&T.




FIGURE 12.6

Plot of the 3D feature-vector space (f_{cc} , A , F_8) for the set of 57 masses in Figure 12.4. 'o': benign masses (37). 'x': malignant tumors (20). Reproduced with permission from H. Alto, R.M. Rangayyan, and J.E.L. Desautels, "Content-based retrieval and analysis of mammographic masses", *Journal of Electronic Imaging*, in press, 2005. © SPIE and IS&T.




FIGURE 12.7

Plot of the 3D feature-vector space (f_{cc}, cf, SI) for the set of 57 contours in Figure 12.5. ‘o’: benign masses (37). ‘x’: malignant tumors (20). Figure courtesy of H. Alto.

The Euclidean distance between an arbitrary pattern vector \mathbf{x} and the i^{th} prototype is given as

$$D_i = \|\mathbf{x} - \mathbf{z}_i\| = \sqrt{(\mathbf{x} - \mathbf{z}_i)^T (\mathbf{x} - \mathbf{z}_i)}. \tag{12.10}$$

A simple rule to classify the pattern vector \mathbf{x} would be to choose that class for which the vector has the smallest distance:

$$\text{if } D_i < D_j \ \forall j \neq i, \text{ then } \mathbf{x} \in C_i. \tag{12.11}$$

(See Section 12.12 for the description of an application of the Euclidean distance to the analysis of breast masses and tumors.)

A simple relationship may be established between discriminant functions and distance functions as follows [1086]:

$$\begin{aligned} D_i^2 &= \|\mathbf{x} - \mathbf{z}_i\|^2 = (\mathbf{x} - \mathbf{z}_i)^T (\mathbf{x} - \mathbf{z}_i) \\ &= \mathbf{x}^T \mathbf{x} - 2\mathbf{x}^T \mathbf{z}_i + \mathbf{z}_i^T \mathbf{z}_i = \mathbf{x}^T \mathbf{x} - 2 \left(\mathbf{x}^T \mathbf{z}_i - \frac{1}{2} \mathbf{z}_i^T \mathbf{z}_i \right). \end{aligned} \tag{12.12}$$

Choosing the minimum of D_i^2 is equivalent to choosing the minimum of D_i (because all $D_i > 0$). Furthermore, from the equation above, it follows




FIGURE 12.8

Plot of the 2D feature-vector space (f_{cc} , SI) for the training set of 18 benign masses ('o') and 10 malignant tumors ('x') selected from the dataset in Figure 12.5. The prototypes of the two classes are indicated by the vectors marked 'B' and 'M'. The solid line shown is a linear decision function, obtained as the perpendicular bisector of the straight line joining the two prototypes (dashed line).




FIGURE 12.9

Plot of the 2D feature-vector space (f_{cc}, SI) for the test set of 19 benign masses ('o') and 10 malignant tumors ('x') selected from the dataset in Figure 12.5. The solid line shown is a linear decision function designed as illustrated in Figure 12.8. Three malignant cases are misclassified by the decision function shown.

that choosing the minimum of D_i^2 is equivalent to choosing the maximum of $(\mathbf{x}^T \mathbf{z}_i - \frac{1}{2} \mathbf{z}_i^T \mathbf{z}_i)$. Therefore, we may define the decision function

$$d_i(\mathbf{x}) = \mathbf{x}^T \mathbf{z}_i - \frac{1}{2} \mathbf{z}_i^T \mathbf{z}_i, \quad i = 1, 2, \dots, M. \quad (12.13)$$

A decision rule may then be stated as

$$\text{if } d_i(\mathbf{x}) > d_j(\mathbf{x}) \quad \forall j \neq i, \text{ then } \mathbf{x} \in C_i. \quad (12.14)$$

This is a linear discriminant function, which becomes obvious from the following representation: If z_{ij} , $j = 1, 2, \dots, n$, are the components of \mathbf{z}_i , let $w_{ij} = z_{ij}$, $j = 1, 2, \dots, n$; $w_{i,n+1} = -\frac{1}{2} \mathbf{z}_i^T \mathbf{z}_i$; and $\mathbf{x} = [x_1, x_2, \dots, x_n, 1]^T$. Then, $d_i(\mathbf{x}) = \mathbf{w}_i^T \mathbf{x}$, $i = 1, 2, \dots, M$, where $\mathbf{w}_i = [w_{i1}, w_{i2}, \dots, w_{i,n+1}]^T$. Therefore, distance functions may be formulated as linear discriminant or decision functions.

12.2.3 The nearest-neighbor rule

Suppose that we are provided with a set of N sample patterns $\{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_N\}$ of known classification: each pattern belongs to one of M classes $\{C_1, C_2, \dots, C_M\}$, with $N \gg M$. We are then given a new feature vector \mathbf{x} whose class needs to be determined. Let us compute a distance measure $D(\mathbf{s}_i, \mathbf{x})$ between the vector \mathbf{x} and each sample pattern. Then, the nearest-neighbor rule states that the vector \mathbf{x} is to be assigned to the class of the sample that is the closest to \mathbf{x} :

$$\mathbf{x} \in C_i \text{ if } D(\mathbf{s}_i, \mathbf{x}) = \min\{D(\mathbf{s}_l, \mathbf{x})\}, \quad l = 1, 2, \dots, N. \quad (12.15)$$

A major disadvantage of the above method is that the classification decision is made based upon a single sample vector of known classification. The nearest neighbor may happen to be an outlier that is not representative of its class. It would be more reliable to base the classification upon several samples: we may consider a certain number k of the nearest neighbors of the sample to be classified, and then seek a majority opinion. This leads to the so-called *k-nearest-neighbor* or *k-NN rule*: Determine the k nearest neighbors of \mathbf{x} , and use the majority of equal classifications in this group as the classification of \mathbf{x} . See Section 12.12 for the description of an application of the k-NN method to the analysis of breast masses and tumors.

12.3 Unsupervised Pattern Classification

Let us consider the situation where we are given a set of feature vectors with no categorization or classes attached to them. No prior training information is available. How may we group the vectors into multiple categories?

The design of distance functions and decision boundaries requires a training set of feature vectors of known classes. The functions so designed may then be applied to a new set of feature vectors or samples to perform pattern classification. Such a procedure is known as *supervised* pattern classification due to the initial training step. In some situations a training step may not be possible, and we may be required to classify a given set of feature vectors into either a prespecified or unknown number of categories. Such a problem is labeled as *unsupervised* pattern classification, and may be solved by cluster-seeking methods.

12.3.1 Cluster-seeking methods

Given a set of feature vectors, we may examine them for the formation of inherent groups or clusters. This is a simple task in the case of 2D vectors, where we may plot them, visually identify groups, and label each group with a pattern class. Allowance may have to be made to assign the same class to multiple disjoint groups. Such an approach may be used even when the number of classes is not known at the outset. When the vectors have a dimension higher than three, visual analysis will not be feasible. It then becomes necessary to define criteria to group the given vectors on the basis of similarity, dissimilarity, or distance measures. A few examples of such measures are described below [1086]:

- Euclidean distance

$$D_E^2 = \|\mathbf{x} - \mathbf{z}\|^2 = (\mathbf{x} - \mathbf{z})^T (\mathbf{x} - \mathbf{z}) = \sum_{i=1}^n (x_i - z_i)^2. \quad (12.16)$$

Here, \mathbf{x} and \mathbf{z} are two feature vectors; the latter could be a class prototype, if available. A small value of D_E indicates greater similarity between the two vectors than a large value of D_E .

- Manhattan or city-block distance

$$D_C = \sum_{i=1}^n |x_i - z_i|. \quad (12.17)$$

The Manhattan distance is the shortest path between \mathbf{x} and \mathbf{z} , with each segment being parallel to a coordinate axis [402].

- Mahalanobis distance

$$D_M^2 = (\mathbf{x} - \mathbf{m})^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{m}), \quad (12.18)$$

where \mathbf{x} is a feature vector being compared to a pattern class for which \mathbf{m} is the class mean vector and \mathbf{C} is the covariance matrix. A small value of D_M indicates a higher potential membership of the vector \mathbf{x} in

the class than a large value of D_M . (See Section 12.12 for the description of an application of the Mahalanobis distance to the analysis of breast masses and tumors.)

- Normalized dot product (cosine of the angle between the vectors \mathbf{x} and \mathbf{z})

$$D_d = \frac{\mathbf{x}^T \mathbf{z}}{\|\mathbf{x}\| \|\mathbf{z}\|}. \quad (12.19)$$

A large dot product value indicates a greater degree of similarity between the two vectors than a small value.

The covariance matrix is defined as

$$\mathbf{C} = E[(\mathbf{y} - \mathbf{m})(\mathbf{y} - \mathbf{m})^T], \quad (12.20)$$

where the expectation operation is performed over all feature vectors \mathbf{y} that belong to the class. The covariance matrix provides the covariance of all possible pairs of the features in the feature vector over all samples belonging to the given class being considered. The elements along the main diagonal of the covariance matrix provide the variance of the individual features that make up the feature vector. The covariance matrix represents the scatter of the features that belong to the given class. The mean and covariance need to be updated as more samples are added to a given class in a clustering procedure.

When the Mahalanobis distance needs to be calculated between a sample vector and a number of classes represented by their mean and covariance matrices, a pooled covariance matrix may be used if the numbers of members in the various classes are unequal and low [1088]. If the covariance matrices of two classes are \mathbf{C}_1 and \mathbf{C}_2 , and the numbers of members in the two classes are N_1 and N_2 , the pooled covariance matrix is given by

$$\mathbf{C} = \frac{(N_1 - 1)\mathbf{C}_1 + (N_2 - 1)\mathbf{C}_2}{N_1 + N_2 - 2}. \quad (12.21)$$

Various performance indices may be designed to measure the success of a clustering procedure [1086]. A measure of the tightness of a cluster is the sum of the squared errors performance index:

$$J = \sum_{j=1}^{N_c} \sum_{\mathbf{x} \in S_j} \|\mathbf{x} - \mathbf{m}_j\|^2, \quad (12.22)$$

where N_c is the number of cluster domains, S_j is the set of samples in the j^{th} cluster,

$$\mathbf{m}_j = \frac{1}{N_j} \sum_{\mathbf{x} \in S_j} \mathbf{x} \quad (12.23)$$

is the sample mean vector of S_j , and N_j is the number of samples in S_j .

A few other examples of performance indices are:

- Average of the squared distances between the samples in a cluster domain.
- Intracluster variance.
- Average of the squared distances between the samples in different cluster domains.
- Intercluster distances.
- Scatter matrices.
- Covariance matrices.

A simple cluster-seeking algorithm [1086]: Suppose we have N sample patterns $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$.

1. Let the first cluster center \mathbf{z}_1 be equal to any one of the samples, say $\mathbf{z}_1 = \mathbf{x}_1$.
2. Choose a nonnegative threshold θ .
3. Compute the distance D_{21} between \mathbf{x}_2 and \mathbf{z}_1 . If $D_{21} < \theta$, assign \mathbf{x}_2 to the domain (class) of cluster center \mathbf{z}_1 ; otherwise, start a new cluster with its center as $\mathbf{z}_2 = \mathbf{x}_2$. For the subsequent steps, let us assume that a new cluster with center \mathbf{z}_2 has been established.
4. Compute the distances D_{31} and D_{32} from the next sample \mathbf{x}_3 to \mathbf{z}_1 and \mathbf{z}_2 , respectively. If D_{31} and D_{32} are both greater than θ , start a new cluster with its center as $\mathbf{z}_3 = \mathbf{x}_3$; otherwise, assign \mathbf{x}_3 to the domain of the closer cluster.
5. Continue to apply Steps 3 and 4 by computing and checking the distance from *every* new (unclassified) pattern vector to *every* established cluster center and applying the assignment or cluster-creation rule.
6. Stop when every given pattern vector has been assigned to a cluster.

Observe that the procedure does not require *a priori* knowledge of the number of classes. Recognize also that the procedure does not assign a real-world class to each cluster: it merely groups the given vectors into disjoint clusters. A subsequent step is required to label each cluster with a class related to the actual problem. Multiple clusters may relate to the same real-world class, and may have to be merged.

A major disadvantage of the simple cluster-seeking algorithm is that the results depend upon

- the first cluster center chosen for each domain or class,
- the order in which the sample patterns are considered,

- the value of the threshold θ , and
- the geometrical properties or distributions of the data, that is, the feature-vector space.

The maximin-distance clustering algorithm [1086]: This method is similar to the previous “simple” algorithm, but first identifies the cluster regions that are the farthest apart. The term “maximin” refers to the combined use of maximum and minimum distances between the given vectors and the centers of the clusters already formed.

1. Let \mathbf{x}_1 be the first cluster center \mathbf{z}_1 .
2. Determine the farthest sample from \mathbf{x}_1 , and label it as cluster center \mathbf{z}_2 .
3. Compute the distance from each remaining sample to \mathbf{z}_1 and to \mathbf{z}_2 . For every pair of these computations, save the minimum distance, and select the maximum of the minimum distances. If this “maximin” distance is an appreciable fraction of the distance between the cluster centers \mathbf{z}_1 and \mathbf{z}_2 , label the corresponding sample as a new cluster center \mathbf{z}_3 ; otherwise stop forming new clusters and go to Step 5.
4. If a new cluster center was formed in Step 3, repeat Step 3 using a “typical” or the average distance between the established cluster centers for comparison.
5. Assign each remaining sample to the domain of its nearest cluster center.

The K -means algorithm [1086]: The preceding “simple” and “maximin” algorithms are intuitive procedures. The K -means algorithm is based on iterative minimization of a performance index that is defined as the sum of the squared distances from all points in a cluster domain to the cluster center.

1. Choose K initial cluster centers $\mathbf{z}_1(1), \mathbf{z}_2(1), \dots, \mathbf{z}_K(1)$. K is the number of clusters to be formed. The choice of the cluster centers is arbitrary, and could be the first K of the feature vectors available. The index in parentheses represents the iteration number.
2. At the k^{th} iterative step, distribute the samples $\{\mathbf{x}\}$ among the K cluster domains, using the relation

$$\mathbf{x} \in S_j(k) \text{ if } \|\mathbf{x} - \mathbf{z}_j(k)\| < \|\mathbf{x} - \mathbf{z}_i(k)\| \forall i = 1, 2, \dots, K, i \neq j, \quad (12.24)$$

where $S_j(k)$ denotes the set of samples whose cluster center is $\mathbf{z}_j(k)$.

3. From the results of Step 2, compute the new cluster centers $\mathbf{z}_j(k+1)$, $j = 1, 2, \dots, K$, such that the sum of the squared distances from all points

in $S_j(k)$ to the new cluster center is minimized. In other words, the new cluster center $\mathbf{z}_j(k+1)$ is computed so that the performance index

$$J_j = \sum_{\mathbf{x} \in S_j(k)} \|\mathbf{x} - \mathbf{z}_j(k+1)\|^2, \quad j = 1, 2, \dots, K, \quad (12.25)$$

is minimized. The $\mathbf{z}_j(k+1)$ that minimizes this performance index is simply the sample mean of $S_j(k)$. Therefore, the new cluster center is given by

$$\mathbf{z}_j(k+1) = \frac{1}{N_j(k)} \sum_{\mathbf{x} \in S_j(k)} \mathbf{x}, \quad j = 1, 2, \dots, K, \quad (12.26)$$

where $N_j(k)$ is the number of samples in $S_j(k)$. The name “ K -means” is derived from the manner in which cluster centers are sequentially updated.

4. If $\mathbf{z}_j(k+1) = \mathbf{z}_j(k)$ for $j = 1, 2, \dots, K$, the algorithm has converged: terminate the procedure; otherwise go to Step 2.

The behavior of the K -means algorithm is influenced by:

- the number of cluster centers specified (K),
- the choice of the initial cluster centers,
- the order in which the sample patterns are considered, and
- the geometrical properties or distributions of the data, that is, the feature-vector space.

Example: Figures 12.10 to 12.14 show four cluster plots of the shape factors f_{cc} and SI of the 57 breast mass contours shown in Figure 12.5 (see Section 12.12 for details). Although the categories of the samples would be unknown in a practical situation, the samples are identified in the plots with the + symbol for malignant tumors and the \circ symbol for the benign masses. (The categorization represents the ground-truth or true classification of the samples based upon biopsy.)

The plots in Figures 12.10 to 12.14 show the progression of the K -means algorithm from its initial state to the converged state. $K = 2$ in this example, representing the benign and malignant categories. The only prior knowledge or assumption used is that the samples are to be split into two clusters, that is, there are two classes. Figure 12.10 shows two samples selected to represent the cluster centers, marked with the diamond and asterisk symbols. The straight line indicates the decision boundary, which is the perpendicular bisector of the straight line joining the two cluster centers. The K -means algorithm converged, in this case, at the fifth iteration (that is, there was no change in the

cluster centers after the fifth iteration). The final decision boundary results in the misclassification of four of the malignant samples as being benign. It is interesting to note that even though the two initial cluster centers belong to the benign category, the algorithm has converged to a useful solution. (See Section 12.12.1 for examples of application of other pattern classification techniques to the same dataset.)

12.4 Probabilistic Models and Statistical Decision

Pattern classification methods such as discriminant functions are dependent upon the set of training samples provided. Their success when applied to new cases will depend upon the accuracy of the representation of the various pattern classes by the training samples. How can we design pattern classification techniques that are independent of specific training samples and are optimal in a broad sense?

Probability functions and probabilistic models may be developed to represent the occurrence and statistical attributes of classes of patterns. Such functions may be based upon large collections of data, historical records, or mathematical models of pattern generation. In the absence of information as above, a training step with samples of known categorization will be required to estimate the required model parameters. It is common practice to assume a Gaussian PDF to represent the distribution of the features for each class, and estimate the required mean and variance parameters from the training sets. When PDFs are available to characterize pattern classes and their features, optimal decision functions may be designed, based upon statistical functions and decision theory. The following subsections describe a few methods in this category.

12.4.1 Likelihood functions and statistical decision

Let $P(C_i)$ be the probability of occurrence of class C_i , $i = 1, 2, \dots, M$; this is known as the *a priori*, *prior*, or unconditional probability. The *a posteriori* or *posterior* probability that an observed sample pattern \mathbf{x} came from C_i is expressed as $P(C_i|\mathbf{x})$. If a classifier decides that \mathbf{x} comes from C_j when it actually came from C_i , the classifier is said to incur a *loss* L_{ij} , with $L_{ii} = 0$ or a fixed operational cost and $L_{ij} > L_{ii} \forall j \neq i$.

Because \mathbf{x} may belong to any one of the M classes under consideration, the expected loss, known as the *conditional average risk* or *loss*, in assigning \mathbf{x} to C_j is [1086]

$$R_j(\mathbf{x}) = \sum_{i=1}^M L_{ij} P(C_i|\mathbf{x}). \quad (12.27)$$




FIGURE 12.10

Initial state of the K -means algorithm. The symbols in the cluster plot represent the 2D feature vectors (f_{cc}, SI) for 37 benign (\circ) and 20 malignant ($+$) breast masses. (See Figure 12.5 for the contours of the masses.) The cluster centers (class means) are indicated by the solid diamond and the $*$ symbols. The straight line indicates the decision boundary between the two classes. Figure courtesy of F.J. Ayres.




FIGURE 12.11

Second iteration of the K -means algorithm. Details as in Figure 12.10. Figure courtesy of F.J. Ayres.





FIGURE 12.12

Third iteration of the *K*-means algorithm. Details as in Figure 12.10. Figure courtesy of F.J. Ayres.

**FIGURE 12.13**

Fourth iteration of the *K*-means algorithm. Details as in Figure 12.10. Figure courtesy of F.J. Ayres.




FIGURE 12.14

Final state of the *K*-means algorithm after the fifth iteration. Details as in Figure 12.10. Figure courtesy of F.J. Ayres.

A classifier could compute $R_j(\mathbf{x})$, $j = 1, 2, \dots, M$, for each sample \mathbf{x} , and then assign \mathbf{x} to the class with the smallest conditional loss. Such a classifier will minimize the total expected loss over all decisions, and is called the *Bayes classifier*. From a statistical point of view, the Bayes classifier represents the optimal classifier.

According to Bayes rule, we have [721, 1086]

$$P(C_i|\mathbf{x}) = \frac{P(C_i) p(\mathbf{x}|C_i)}{p(\mathbf{x})}, \quad (12.28)$$

where $p(\mathbf{x}|C_i)$ is called the *likelihood function* of class C_i or the *state-conditional PDF* of \mathbf{x} , and $p(\mathbf{x})$ is the PDF of \mathbf{x} regardless of class membership (unconditional). [Note: $P(y)$ is used to represent the probability of occurrence of an event y ; $p(y)$ is used to represent the PDF of a random variable y . Probabilities and PDFs involving a multidimensional feature vectors are multivariate functions with dimension equal to that of the feature vector.] Bayes rule shows how observing the sample \mathbf{x} changes the *a priori* probability $P(C_i)$ to the *a posteriori* probability $P(C_i|\mathbf{x})$. In other words, Bayes rule provides a mechanism to update the *a priori* probability $P(C_i)$ to the *a posteriori* probability $P(C_i|\mathbf{x})$ due to the observation of the sample \mathbf{x} . Then, we can express the expected loss as [1086]

$$R_j(\mathbf{x}) = \frac{1}{p(\mathbf{x})} \sum_{i=1}^M L_{ij} p(\mathbf{x}|C_i) P(C_i). \quad (12.29)$$

Because $\frac{1}{p(\mathbf{x})}$ is common for all j , we could modify $R_j(\mathbf{x})$ to

$$r_j(\mathbf{x}) = \sum_{i=1}^M L_{ij} p(\mathbf{x}|C_i) P(C_i). \quad (12.30)$$

In a two-class case with $M = 2$, we obtain the following expressions [1086]:

$$r_1(\mathbf{x}) = L_{11} p(\mathbf{x}|C_1) P(C_1) + L_{21} p(\mathbf{x}|C_2) P(C_2). \quad (12.31)$$

$$r_2(\mathbf{x}) = L_{12} p(\mathbf{x}|C_1) P(C_1) + L_{22} p(\mathbf{x}|C_2) P(C_2). \quad (12.32)$$

$$\mathbf{x} \in C_1 \text{ if } r_1(\mathbf{x}) < r_2(\mathbf{x}), \quad (12.33)$$

that is,

$$\begin{aligned} \mathbf{x} \in C_1 \text{ if } [L_{11} p(\mathbf{x}|C_1) P(C_1) + L_{21} p(\mathbf{x}|C_2) P(C_2)] \\ < [L_{12} p(\mathbf{x}|C_1) P(C_1) + L_{22} p(\mathbf{x}|C_2) P(C_2)], \end{aligned} \quad (12.34)$$

or equivalently,

$$\mathbf{x} \in C_1 \text{ if } [(L_{21} - L_{22}) p(\mathbf{x}|C_2) P(C_2)] < [(L_{12} - L_{11}) p(\mathbf{x}|C_1) P(C_1)]. \quad (12.35)$$

This expression may be rewritten as

$$\mathbf{x} \in C_1 \text{ if } \frac{p(\mathbf{x}|C_1)}{p(\mathbf{x}|C_2)} > \frac{P(C_2)(L_{21} - L_{22})}{P(C_1)(L_{12} - L_{11})}. \quad (12.36)$$

The left-hand side of the inequality above, which is a ratio of two likelihood functions, is often referred to as the *likelihood ratio*:

$$l_{12}(\mathbf{x}) = \frac{p(\mathbf{x}|C_1)}{p(\mathbf{x}|C_2)}. \quad (12.37)$$

Then, Bayes decision rule for $M = 2$ is [1086]:

1. Assign \mathbf{x} to class C_1 if $l_{12}(\mathbf{x}) > \theta_{12}$, where θ_{12} is a threshold given by $\theta_{12} = \frac{P(C_2)(L_{21} - L_{22})}{P(C_1)(L_{12} - L_{11})}$.
2. Assign \mathbf{x} to class C_2 if $l_{12}(\mathbf{x}) < \theta_{12}$.
3. Make an arbitrary or heuristic decision if $l_{12}(\mathbf{x}) = \theta_{12}$.

The rule may be generalized to the M -class case as [1086]:

$$\mathbf{x} \in C_i \text{ if } \sum_{k=1}^M L_{ki} p(\mathbf{x}|C_k) P(C_k) < \sum_{q=1}^M L_{qj} p(\mathbf{x}|C_q) P(C_q),$$

$$j = 1, 2, \dots, M, j \neq i. \quad (12.38)$$

In most pattern classification problems, the loss is nil for correct decisions. The loss could be assumed to be equal to a certain quantity for all erroneous decisions. Then, $L_{ij} = 1 - \delta_{ij}$, where

$$\delta_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}, \quad (12.39)$$

and

$$r_j(\mathbf{x}) = \sum_{i=1}^M (1 - \delta_{ij}) p(\mathbf{x}|C_i) P(C_i) \quad (12.40)$$

$$= p(\mathbf{x}) - p(\mathbf{x}|C_j) P(C_j),$$

because

$$\sum_{i=1}^M p(\mathbf{x}|C_i) P(C_i) = p(\mathbf{x}). \quad (12.41)$$

The Bayes classifier will assign a pattern \mathbf{x} to class C_i if

$$p(\mathbf{x}) - p(\mathbf{x}|C_i)P(C_i) < p(\mathbf{x}) - p(\mathbf{x}|C_j)P(C_j), \quad j = 1, 2, \dots, M, j \neq i, \quad (12.42)$$

that is,

$$\mathbf{x} \in C_i \text{ if } p(\mathbf{x}|C_i)P(C_i) > p(\mathbf{x}|C_j)P(C_j), j = 1, 2, \dots, M, j \neq i. \quad (12.43)$$

This is nothing more than using the decision functions

$$d_i(\mathbf{x}) = p(\mathbf{x}|C_i) P(C_i), i = 1, 2, \dots, M, \quad (12.44)$$

where a pattern \mathbf{x} is assigned to class C_i if $d_i(\mathbf{x}) > d_j(\mathbf{x}) \forall j \neq i$ for that pattern. Using Bayes rule, we get

$$d_i(\mathbf{x}) = P(C_i|\mathbf{x}) p(\mathbf{x}), i = 1, 2, \dots, M. \quad (12.45)$$

Because $p(\mathbf{x})$ does not depend upon the class index i , this can be reduced to

$$d_i(\mathbf{x}) = P(C_i|\mathbf{x}), i = 1, 2, \dots, M. \quad (12.46)$$

The different decision functions given above provide alternative yet equivalent approaches, depending upon whether $p(\mathbf{x}|C_i)$ or $P(C_i|\mathbf{x})$ is used (or available). The estimation of $p(\mathbf{x}|C_i)$ would require a training set for each class C_i . It is common to assume a Gaussian distribution and estimate its mean and variance using the training set.

12.4.2 Bayes classifier for normal patterns

The univariate normal or Gaussian PDF for a single random variable x is given by

$$p(x) = \frac{1}{\sqrt{2\pi} \sigma} \exp \left[-\frac{1}{2} \left(\frac{x - m}{\sigma} \right)^2 \right], \quad (12.47)$$

which is completely specified by two parameters: the mean

$$m = E[x] = \int_{-\infty}^{\infty} x p(x) dx, \quad (12.48)$$

and the variance

$$\sigma^2 = E[(x - m)^2] = \int_{-\infty}^{\infty} (x - m)^2 p(x) dx. \quad (12.49)$$

In the case of M pattern classes and pattern vectors \mathbf{x} of dimension n governed by multivariate normal PDFs, we have

$$p(\mathbf{x}|C_i) = \frac{1}{(2\pi)^{n/2} |\mathbf{C}_i|^{1/2}} \exp \left[-\frac{1}{2} (\mathbf{x} - \mathbf{m}_i)^T \mathbf{C}_i^{-1} (\mathbf{x} - \mathbf{m}_i) \right], \quad (12.50)$$

$i = 1, 2, \dots, M$, where each PDF is completely specified by its mean vector \mathbf{m}_i and its $n \times n$ covariance matrix \mathbf{C}_i , with

$$\mathbf{m}_i = E_i[\mathbf{x}], \quad (12.51)$$

and

$$\mathbf{C}_i = E_i[(\mathbf{x} - \mathbf{m}_i)(\mathbf{x} - \mathbf{m}_i)^T]. \quad (12.52)$$

Here, $E_i[\]$ denotes the expectation operator over the patterns belonging to class C_i .

Normal distributions occur frequently in nature, and have the advantage of analytical tractability. A multivariate normal PDF reduces to a product of univariate normal PDFs when the elements of \mathbf{x} are mutually independent (in which case the covariance matrix is a diagonal matrix).

We had earlier formulated the decision functions

$$d_i(\mathbf{x}) = p(\mathbf{x}|C_i) P(C_i), \quad i = 1, 2, \dots, M; \quad (12.53)$$

see Equation 12.44. Given the exponential in the normal PDF, it is convenient to use

$$d_i(\mathbf{x}) = \ln [p(\mathbf{x}|C_i) P(C_i)] = \ln p(\mathbf{x}|C_i) + \ln P(C_i), \quad (12.54)$$

which is equivalent in terms of classification performance because the natural logarithm \ln is a monotonically increasing function. Then [1086],

$$d_i(\mathbf{x}) = \ln P(C_i) - \frac{n}{2} \ln 2\pi - \frac{1}{2} \ln |\mathbf{C}_i| - \frac{1}{2} [(\mathbf{x} - \mathbf{m}_i)^T \mathbf{C}_i^{-1} (\mathbf{x} - \mathbf{m}_i)], \quad (12.55)$$

$i = 1, 2, \dots, M$. The second term does not depend upon i ; therefore, we can simplify $d_i(\mathbf{x})$ to

$$d_i(\mathbf{x}) = \ln P(C_i) - \frac{1}{2} \ln |\mathbf{C}_i| - \frac{1}{2} [(\mathbf{x} - \mathbf{m}_i)^T \mathbf{C}_i^{-1} (\mathbf{x} - \mathbf{m}_i)], \quad i = 1, 2, \dots, M. \quad (12.56)$$

The decision functions above are hyperquadrics; hence, the best that a Bayes classifier for normal patterns can do is to place a general second-order decision surface between each pair of pattern classes. In the case of true normal distributions of patterns, the decision functions as above will be optimal on an average basis: they minimize the expected loss with the simplified loss function $L_{ij} = 1 - \delta_{ij}$ [1086].

If all the covariance matrices are equal, that is, $\mathbf{C}_i = \mathbf{C}$, $i = 1, 2, \dots, M$, we get

$$d_i(\mathbf{x}) = \ln P(C_i) + \mathbf{x}^T \mathbf{C}^{-1} \mathbf{m}_i - \frac{1}{2} \mathbf{m}_i^T \mathbf{C}^{-1} \mathbf{m}_i, \quad i = 1, 2, \dots, M, \quad (12.57)$$

after omitting terms independent of i . The Bayesian classifier is now represented by a set of linear decision functions.

Before one may apply the decision functions as above, it would be appropriate to verify the Gaussian nature of the PDFs of the variables on hand by conducting statistical tests [168, 1087]. Furthermore, it would be necessary to derive or estimate the mean vector and covariance matrix for each class; sample statistics computed from a training set may serve this purpose.

Example: Figure 12.15 shows plots of Gaussian PDF models applied to the shape factor f_{cc} (fractional concavity) of the 57 breast mass contours shown in Figure 12.5 (see Chapter 6 and Section 12.12 for details). The two Gaussians represent the state-conditional PDFs of f_{cc} for the benign and malignant categories. Also shown are the posterior probabilities that the class of a sample is benign or malignant given an observed value of f_{cc} . The posterior probability functions were derived using Bayes rule as in Equation 12.28, with the values of the prior probabilities of the two classes being equal to 0.5. It is seen that the posterior probabilities are both equal to 0.5 at $f_{cc} = 0.32$; the probability of a malignant classification is higher than that of a benign classification for $f_{cc} > 0.32$. Due to the use of equal prior probabilities, the transition point is the same as the point where the two Gaussian models for the state-conditional PDFs cross each other.

Figure 12.16 shows the same Gaussian models for the state-conditional PDFs as in Figure 12.15. However, the posterior probability functions were derived using the prior probability value of 0.9 for the benign category; the prior probability for the malignant category is then 0.1. In this case, the probability of a malignant classification is higher than that of a benign classification for $f_{cc} > 0.36$. The prior assumption that 90% of all masses encountered will be benign has pushed the decision threshold on f_{cc} to a higher value.

Figure 12.17 illustrates the 2D cluster plot and Gaussian PDF models for the shape factors f_{cc} and SI for the same dataset as described above; see Chapter 6 and Section 12.12 for details. The decision boundary indicated by the solid line is the optimal boundary under the assumption of 2D Gaussian PDFs for the two features and the two classes. Two malignant samples are misclassified by the decision boundary shown. (See Section 12.12.1 for examples of application of other pattern classification techniques to the same dataset.)

An interesting point to note from the examples above is that the Gaussian PDF models used are not capable of accommodating the prior knowledge that the shape factors are limited to the range $[0, 1]$. Other PDF models such as the Rayleigh distribution (see Section 3.1.2) should be used if this aspect is important.

12.5 Logistic Regression

Logistic classification is a statistical technique based on a logistic regression model that estimates the probability of occurrence of an event [1091, 1092, 1093]. The technique is designed for problems where patterns are to be classified into one of two classes. When the response variable is binary, theoretical and empirical considerations indicate that the response function is often curvi-




FIGURE 12.15

Plots of Gaussian state-conditional PDF models for the shape factor f_{cc} for benign (dashed line) and malignant (dotted line) breast masses. (See Figure 12.5 for the contours of the masses.) The f_{cc} values of the samples are indicated on the horizontal axis for 37 benign masses with \circ and for 20 malignant tumors as \times . The posterior probability functions for the benign (solid line) and malignant (dash-dot line) classes are also shown. Equal prior probabilities of 0.5 were used for the two classes. See also Figure 12.16. Figure courtesy of F.J. Ayres.




FIGURE 12.16

Same as in Figure 12.15, but with the prior probability of the benign class equal to 0.9. Figure courtesy of F.J. Ayres.




FIGURE 12.17

Plots of 2D Gaussian state-conditional PDF models for the shape factors f_{cc} and SI for benign masses (dash-dot line) and malignant tumors (dashed line). See Figure 12.5 for the contours of the masses. Feature values for 37 benign masses are indicated with \circ , and for 20 malignant tumors as \times . The benign class prototype (mean) is indicated by the solid diamond; that for the malignant class is indicated by the $*$ symbol. The dashed and dash-dot contours indicate two constant-Mahalanobis-distance contours (level sets) each for the two Gaussian PDF models (see Equation 12.18) for the malignant and benign classes, respectively. The solid contour indicates the decision boundary, as given by Equation 12.53 to Equation 12.56, with the decision function being equal for the two classes. The prior probabilities for the two classes were assumed to be equal to 0.5. Figure courtesy of F.J. Ayres.

linear. The typical response function is shaped as a forward or backward tilted “S”, and is known as a sigmoidal function. The function has asymptotes at 0 and 1.

In logistic pattern classification, an event is defined as the membership of a pattern vector in one of the two classes of concern. The method computes a variable that depends upon the given parameters and is constrained to the range $[0, 1]$ so that it may be interpreted as a probability. The probability of the pattern vector belonging to the second class is simply the difference between unity and the estimated value.

For the case of a single feature or parameter, the logistic regression model is given as

$$P(\text{event}) = \frac{\exp(b_0 + b_1 x)}{1 + \exp(b_0 + b_1 x)}, \quad (12.58)$$

or equivalently,

$$P(\text{event}) = \frac{1}{1 + \exp[-(b_0 + b_1 x)]}, \quad (12.59)$$

where b_0 and b_1 are coefficients estimated from the data, and x is the independent (feature) variable. The relationship between the independent variable and the estimated probability is nonlinear, and follows an S-shaped curve that closely resembles the integral of a Gaussian function. In the case of an n -dimensional feature vector \mathbf{x} , the model can be written as

$$P(\text{event}) = \frac{1}{1 + \exp(-z)}, \quad (12.60)$$

where z is the linear combination

$$z = b_0 + b_1 x_1 + b_2 x_2 + \cdots + b_n x_n = \mathbf{b}^T \mathbf{x}, \quad (12.61)$$

that is, z is the dot product of the augmented feature vector \mathbf{x} with a coefficient vector or weight vector \mathbf{b} .

In linear regression, the coefficients of the model are estimated using the method of least squares; the selected regression coefficients are those that result in the smallest sums of squared distances between the observed and the predicted values of the dependent variable. In logistic regression, the parameters of the model are estimated using the maximum likelihood method [1087, 1091]; the coefficients that make the observed results “most likely” are selected. Because the logistic regression model is nonlinear, an iterative algorithm is necessary for the estimation of the coefficients [1092, 1093]. A training set is required to design a classifier based upon logistic regression. See Sections 5.5.2, 5.5.3, and 12.12.1 for illustrations of the application of logistic regression to the classification of breast masses and tumors.

12.6 The Training and Test Steps

In the situation when a limited number of sample vectors with known classification are available, questions arise as to how many of the samples may be used to design or train a classifier, with the understanding that the classifier so designed needs to be tested using an independent set of samples of known classification as well. When a sufficiently large number of samples are available, they may be randomly split into two approximately equal sets, one for use as the training set and the other to be used as the test set. The random-splitting procedure may be repeated a number of times to generate several classifiers. Finally, one of the classifiers so designed may be selected based upon its performance in both the training and test steps.

12.6.1 The leave-one-out method

The leave-one-out method [1087] is suitable for the estimation of the classification accuracy of a pattern classification technique, particularly when the number of available samples is small. In this method, one of the available samples is excluded, the classifier is designed with the remaining samples, and then the classifier is applied to the excluded sample. The validity of the classification so performed is noted. This procedure is repeated with each available sample: if N training samples are available, N classifiers are designed and tested. The training and test sets for any one classifier so designed and tested are independent. However, while the training set for each classifier has $N - 1$ samples, the test set has only one sample. In the final analysis, every sample will have served as a training sample ($N - 1$ times) as well as a test sample (once). An average classification accuracy is then computed using all of the test results.

Let us consider a simple case in which the covariances of the sample sets of two classes are equal. Assume that two sample sets, $S_1 = \{\mathbf{x}_1^{(1)}, \dots, \mathbf{x}_{N_1}^{(1)}\}$ from class C_1 , and $S_2 = \{\mathbf{x}_1^{(2)}, \dots, \mathbf{x}_{N_2}^{(2)}\}$ from class C_2 are given. Here, N_1 and N_2 are the numbers of samples in the sets S_1 and S_2 , respectively. Assume also that the prior probabilities of the two classes are equal to each other. Then, according to the Bayes classifier, and assuming \mathbf{x} to be governed by a multivariate Gaussian PDF, a sample \mathbf{x} is assigned to the class C_1 if

$$(\mathbf{x} - \mathbf{m}_1)^T(\mathbf{x} - \mathbf{m}_1) - (\mathbf{x} - \mathbf{m}_2)^T(\mathbf{x} - \mathbf{m}_2) > \theta, \quad (12.62)$$

where θ is a threshold, and the sample mean $\tilde{\mathbf{m}}_i$ is given by

$$\tilde{\mathbf{m}}_i = \frac{1}{N_i} \sum_{j=1}^{N_i} \mathbf{x}_j^{(i)}. \quad (12.63)$$

In the leave-one-out method, one sample $\mathbf{x}_k^{(i)}$ is excluded from the training set and then used as the test sample. The mean estimate for class C_i without $\mathbf{x}_k^{(i)}$, labeled as $\tilde{\mathbf{m}}_{ik}$, may be computed as

$$\tilde{\mathbf{m}}_{ik} = \frac{1}{N_i - 1} \left[\sum_{j=1}^{N_i} \mathbf{x}_j^{(i)} - \mathbf{x}_k^{(i)} \right], \quad (12.64)$$

which leads to

$$\mathbf{x}_k^{(i)} - \tilde{\mathbf{m}}_{ik} = \frac{N_i}{N_i - 1} (\mathbf{x}_k^{(i)} - \tilde{\mathbf{m}}_i). \quad (12.65)$$

Then, testing a sample $\mathbf{x}_k^{(1)}$ from C_1 can be carried out as

$$\begin{aligned} & (\mathbf{x}_k^{(1)} - \tilde{\mathbf{m}}_{1k})^T (\mathbf{x}_k^{(1)} - \tilde{\mathbf{m}}_{1k}) - (\mathbf{x}_k^{(1)} - \tilde{\mathbf{m}}_2)^T (\mathbf{x}_k^{(1)} - \tilde{\mathbf{m}}_2) \\ &= \left(\frac{N_1}{N_1 - 1} \right)^2 (\mathbf{x}_k^{(1)} - \tilde{\mathbf{m}}_1)^T (\mathbf{x}_k^{(1)} - \tilde{\mathbf{m}}_1) - (\mathbf{x}_k^{(1)} - \tilde{\mathbf{m}}_2)^T (\mathbf{x}_k^{(1)} - \tilde{\mathbf{m}}_2) \\ &> \theta. \end{aligned} \quad (12.66)$$

Observe that when $\mathbf{x}_k^{(1)}$ is tested, only $\tilde{\mathbf{m}}_1$ is changed and $\tilde{\mathbf{m}}_2$ is not changed. Likewise, when a sample $\mathbf{x}_k^{(2)}$ from C_2 is tested, the decision rule is

$$\begin{aligned} & (\mathbf{x}_k^{(2)} - \tilde{\mathbf{m}}_1)^T (\mathbf{x}_k^{(2)} - \tilde{\mathbf{m}}_1) - (\mathbf{x}_k^{(2)} - \tilde{\mathbf{m}}_{2k})^T (\mathbf{x}_k^{(2)} - \tilde{\mathbf{m}}_{2k}) \\ &= (\mathbf{x}_k^{(2)} - \tilde{\mathbf{m}}_1)^T (\mathbf{x}_k^{(2)} - \tilde{\mathbf{m}}_1) - \left(\frac{N_2}{N_2 - 1} \right)^2 (\mathbf{x}_k^{(2)} - \tilde{\mathbf{m}}_2)^T (\mathbf{x}_k^{(2)} - \tilde{\mathbf{m}}_2) \\ &< \theta. \end{aligned} \quad (12.67)$$

The leave-one-out method provides the least-biased (practically unbiased) estimate of the classification accuracy of a given classification method for a given training set, and is useful when the number of samples available with known classification is small.

12.7 Neural Networks

In many practical problems, we may have no knowledge of the prior probabilities of patterns belonging to one class or another. No general classification rules may exist for the patterns on hand. Clinical knowledge may not yield symbolic knowledge bases that could be used to classify patterns that demonstrate exceptional behavior. In such situations, conventional pattern classification methods as described in the preceding sections may not be well-suited for the classification of pattern vectors. Artificial neural networks (ANNs),

with the properties of experience-based learning and fault tolerance, should be effective in solving such classification problems [274, 402, 1089, 1090, 1094, 1095, 1096].

Figure 12.18 illustrates a two-layer perceptron with one hidden layer and one output layer for pattern classification. The network learns the similarities among patterns directly from their instances in the training set that is provided initially. Classification rules are inferred from the training data without prior knowledge of the pattern class distributions in the data. Training of an ANN classifier is typically achieved by the *back-propagation* algorithm [274, 402, 1089, 1090, 1094, 1095, 1096]. The actual output of the ANN y_k is calculated as

$$y_k = f \left(\sum_{j=1}^J w_{jk}^{\#} x_j^{\#} - \theta_k^{\#} \right), \quad k = 1, 2, \dots, K, \quad (12.68)$$

where

$$x_j^{\#} = f \left(\sum_{i=1}^I w_{ij} x_i - \theta_j \right), \quad j = 1, 2, \dots, J, \quad (12.69)$$

and

$$f(\beta) = \frac{1}{1 + \exp(-\beta)}. \quad (12.70)$$

In the equations above, θ_j and $\theta_k^{\#}$ are node offsets; w_{ij} and $w_{jk}^{\#}$ are node weights; x_i are the elements of the pattern vectors (input parameters); and I , J , and K are the numbers of nodes in the input, hidden, and output layers, respectively. The weights and offsets are updated by

$$w_{jk}^{\#}(n+1) = w_{jk}^{\#}(n) + \eta[y_k(1-y_k)(d_k - y_k)]x_j^{\#} + \alpha[w_{jk}^{\#}(n) - w_{jk}^{\#}(n-1)], \quad (12.71)$$

$$\theta_k^{\#}(n+1) = \theta_k^{\#}(n) + \eta[y_k(1-y_k)(d_k - y_k)](-1) + \alpha[\theta_k^{\#}(n) - \theta_k^{\#}(n-1)], \quad (12.72)$$

$$w_{ij}(n+1) = w_{ij}(n) + \eta \left[x_j^{\#}(1 - x_j^{\#}) \sum_{k=1}^K \{y_k(1 - y_k)(d_k - y_k)w_{jk}^{\#}\} \right] x_i + \alpha[w_{ij}(n) - w_{ij}(n-1)], \quad (12.73)$$

and

$$\theta_j(n+1) = \theta_j + \eta \left[x_j^{\#}(1 - x_j^{\#}) \sum_{k=1}^K \{y_k(1 - y_k)(d_k - y_k)w_{jk}^{\#}\} \right] (-1) + \alpha[\theta_j(n) - \theta_j(n-1)], \quad (12.74)$$

where d_k are the desired outputs, α is a momentum term, η is a gain term, and n refers to the iteration number. Equations 12.71 and 12.72 represent the




FIGURE 12.18

A two-layer perceptron. Reproduced with permission from L. Shen, R.M. Rangayyan, and J.E.L. Desautels, "Detection and classification of mammographic calcifications", *International Journal of Pattern Recognition and Artificial Intelligence*, 7(6):1403–1416, 1993. © World Scientific Publishing Company.

back-propagation steps, with $y_k(1 - y_k)x_j^\#$ being the sensitivity of y_k to $w_{jk}^\#$, that is, $\frac{\partial y_k}{\partial w_{jk}^\#}$.

The classifier training algorithm is repeated until the errors between the desired outputs and actual outputs for the training data are smaller than a predetermined threshold value.

Example of application: In the work of Shen et al. [274, 320], the shape features (mf , ff , cf) (see Section 6.6) were employed as inputs (x_i , $i = 1, 2, 3$) to the ANN as above, and calcifications were classified into two groups: benign or malignant. Therefore, the numbers of input (I) and output (K) nodes are 3 and 2, respectively. Feature sets for training of the ANN were computed for 143 calcifications, of which 64 were benign and 79 malignant. The calcifications were obtained from 18 mammograms of biopsy-proven cases chosen from the Radiology Teaching Library of the Foothills Hospital, Calgary,

Alberta, Canada. Boundaries of the calcifications were obtained by region growing after manual selection of seed pixels and tolerance [334]. Figure 6.25 provides a 3D plot of the feature vectors for the 143 calcifications used as the training data.

Three parameters — the number of hidden nodes J , the gain term η , and the momentum term α — need to be determined before training the two-layer perceptron. There is no general rule available for the selection of these parameters. One of the most common methods is trial-and-error by choosing the set of parameters with which the highest training speed (the smallest number of iterations) is achieved. However, the major disadvantage of this method is that the classification effectiveness after training is not considered. To use the training data set more efficiently and to overcome the above-mentioned shortcoming of the trial-and-error method, Shen et al. included a leave-one-out type of algorithm [1087, 1097] in the procedure for determining the three parameters (J , η , and α), as described next.

First, η and α were held at fixed values in order to determine the most suitable number of hidden nodes J . Table 12.1 and Table 12.2 list the results (the average number of iterations and the number of erroneous classifications) of the training and parameter determination method under two circumstances: ($\eta = 1.0, \alpha = 0.7$) and ($\eta = 2.0, \alpha = 0.7$), respectively. Based upon Tables 12.1 and 12.2, $J = 10$ could be selected as it achieved the fewest number of classification errors in a reasonable number of iterations. (The use of a higher number of nodes did not result in a significant reduction in the number of iterations and did not reduce the error of classification.)

After determining the most suitable value for J , the best value for η was determined by fixing $J = 10$ and $\alpha = 0.7$. The corresponding results are listed in Table 12.3. It is seen that $\eta = 2.0$ is the best value of those evaluated.

Finally, the value for α was determined with $J = 10$ and $\eta = 2.0$. Table 12.4 provides the results of various trials. It is clear that the best value of those tried for α is 0.7.

After obtaining the most suitable values for the parameter set, the perceptron was trained again by using all of the training data with $J = 10$, $\eta = 2.0$, and $\alpha = 0.7$. All of the 143 calcifications in the training set were correctly classified in 1,268 iterations. The weight set obtained by this procedure was utilized for the classification of calcifications in the test set.

Sections of size $1,024 \times 768$, 768×512 , 512×768 , and 512×768 pixels of four typical mammograms from complete images of up to $2,560 \times 4,096$ pixels with biopsy-proven calcifications were utilized for the test step. Two of the sections had a total of 58 benign calcifications whereas the other two contained 241 ± 10 malignant calcifications. Based upon visual inspection by a radiologist, the detection rates of the multitolerance region-growing algorithm (see Section 5.4.9) were 81% with no false calcifications and $85 \pm 3\%$ with 29 false calcifications for the benign and malignant sections, respectively [274]. After the detection procedure, the calcifications were classified by the ANN classifier. The correct classification rate for the detected benign calcifications

TABLE 12.1

Results of the Training and ANN Parameter Determination Algorithm with $\eta = 1.0$ and $\alpha = 0.7$.

Number of hidden nodes (J)	Number of iterations (mean)	Number of erroneous classifications out of 143 samples
1	1,810	4
3	1,809	4
5	1,774	4
7	1,729	4
9	1,697	4
10	1,697	3
15	1,676	3
20	1,675	3
30	1,676	3

Reproduced with permission from L. Shen, R.M. Rangayyan, and J.E.L. Desautels, "Detection and classification of mammographic calcifications", *International Journal of Pattern Recognition and Artificial Intelligence*, 7(6):1403–1416, 1993. © World Scientific Publishing Company.

TABLE 12.2

Results of the Training and ANN Parameter Determination Algorithm with $\eta = 2.0$ and $\alpha = 0.7$.

Number of hidden nodes (J)	Number of iterations (mean)	Number of erroneous classifications out of 143 samples
5	1,448	3
10	1,391	3
12	1,380	4
15	1,387	3
20	1,401	4

Reproduced with permission from L. Shen, R.M. Rangayyan, and J.E.L. Desautels, "Detection and classification of mammographic calcifications", *International Journal of Pattern Recognition and Artificial Intelligence*, 7(6):1403–1416, 1993. © World Scientific Publishing Company.

TABLE 12.3

Results of the Training and ANN Parameter Determination Algorithm with $J = 10$ and $\alpha = 0.7$.

Gain η	Number of iterations (mean)	Number of erroneous classifications out of 143 samples
0.3	4,088	4
0.7	2,145	4
1.0	1,697	3
1.5	1,431	4
2.0	1,391	3
2.3	1,430	3
3.0	2,691	5

Reproduced with permission from L. Shen, R.M. Rangayyan, and J.E.L. Desautels, "Detection and classification of mammographic calcifications", *International Journal of Pattern Recognition and Artificial Intelligence*, 7(6):1403–1416, 1993. © World Scientific Publishing Company.

TABLE 12.4

Results of the Training and ANN Parameter Determination Algorithm with $J = 10$ and $\eta = 2.0$.

Momentum α	Number of iterations (mean)	Number of erroneous classifications out of 143 samples
0.5	2,426	4
0.7	1,391	3
0.9	4,152	7

Reproduced with permission from L. Shen, R.M. Rangayyan, and J.E.L. Desautels, "Detection and classification of mammographic calcifications", *International Journal of Pattern Recognition and Artificial Intelligence*, 7(6):1403–1416, 1993. © World Scientific Publishing Company.

was 94%, whereas the correct classification rate for the correctly detected malignant calcifications was 87%.

Classification errors for benign calcifications arose mainly from overlapping calcifications. One possible solution to this problem is two-view analysis [1098]. A likely reason for erroneous classification of malignant calcifications is that not all calcifications within a malignant calcification cluster may possess rough contours; a cluster analysis procedure may assist in overcoming this situation.

12.8 Measures of Diagnostic Accuracy

Pattern recognition or classification decisions that are made in the context of medical diagnosis have implications that go beyond statistical measures of accuracy and validity. We need to provide a clinical or diagnostic interpretation of statistical or rule-based decisions made with pattern vectors.

Consider the simple situation of screening, which represents the use of a test to determine the presence or absence of a specific disease in a certain study population: the decision to be made is binary. Let us represent by A the event that a subject has the particular pathology (or is abnormal), and by N the event that the subject does not have the disease (which may not necessarily mean that the subject is normal). Let the prior probabilities $P(A)$ and $P(N)$ represent the fractions of subjects with the disease and the normal subjects, respectively, in the test population. Let T^+ represent a positive screening test result (indicative of the presence of the disease) and T^- a negative result. The following possibilities arise [1099]:

- A *true positive* (TP) or a “hit” is the situation when the test is positive for a subject with the disease. The true-positive fraction (*TPF*) or *sensitivity* S^+ is given as $P(T^+|A)$ or

$$S^+ = \frac{\text{number of TP decisions}}{\text{number of subjects with the disease}} . \quad (12.75)$$

The sensitivity of a test represents its capability to detect or identify the presence of the disease of concern.

- A *true negative* (TN) represents the case when the test is negative for a subject who does not have the disease. The true-negative fraction (*TNF*) or *specificity* S^- is given as $P(T^-|N)$ or

$$S^- = \frac{\text{number of TN decisions}}{\text{number of subjects without the disease}} . \quad (12.76)$$

The specificity of a test indicates its accuracy in recognizing the absence of the disease of concern.

- A *false negative* (FN) or a “miss” is said to occur when the test is negative for a subject who has the disease of concern; that is, the test has missed the case. The probability of this error, known as the false-negative fraction (*FNF*), is $P(T^-|A)$.
- A *false positive* (FP) or a false alarm is defined as the case where the result of the test is positive when the individual being tested does not have the disease. The probability of this type of error, known as the false-positive fraction (*FPF*), is $P(T^+|N)$.

Table 12.5 summarizes the classification possibilities. Observe that

- $FNF + TPF = 1$,
- $FPF + TNF = 1$,
- $S^- = 1 - FPF = TNF$, and
- $S^+ = 1 - FNF = TPF$.

A summary measure of accuracy may be defined as [1099]

$$\text{accuracy} = S^+ P(A) + S^- P(N), \tag{12.77}$$

where $P(A)$ is the fraction of the study population that actually has the disease (that is, the prevalence of the disease) and $P(N)$ is the fraction of the study population that is actually free of the disease.

TABLE 12.5
Schematic Representation of a
Classification Matrix.

Actual group	Predicted group	
	Normal	Abnormal
Normal	$S^- = TNF$	FPF
Abnormal	FNF	$S^+ = TPF$

Reproduced with permission from R.M. Rangayyan, *Biomedical Signal Analysis: A Case-Study Approach*, IEEE Press and Wiley, New York, NY. 2002, © IEEE.

The efficiency of a test may also be indicated by its predictive values. The *positive predictive value* PPV of a test, defined as

$$PPV = 100 \frac{TP}{TP + FP}, \quad (12.78)$$

represents the percentage of the cases labeled as positive by the test that are actually positive. The *negative predictive value* NPV , defined as

$$NPV = 100 \frac{TN}{TN + FN}, \quad (12.79)$$

represents the percentage of cases labeled as negative by the test that are actually negative.

When a new test or method of diagnosis is being developed and tested, it will be necessary to use another previously established method as a reference to confirm the presence or absence of the disease. Such a reference method is often called the *gold standard*. When computer-based methods need to be tested, it is common practice to use the diagnosis or classification provided by an expert in the field as the gold standard. Results of biopsy, other established laboratory or investigative procedures, or long-term clinical follow-up in the case of normal subjects may also serve this purpose. The term “actual group” in Table 12.5 indicates the result of the gold standard, and the term “predicted group” refers to the result of the test conducted.

Health-care professionals (and the general public) would be interested in knowing the probability that a subject with a positive test result actually has the disease: this is given by the conditional probability $P(A|T^+)$. The question could be answered by using Bayes rule [1087], using which we can obtain

$$P(A|T^+) = \frac{P(A) P(T^+|A)}{P(A)P(T^+|A) + P(N)P(T^+|N)}. \quad (12.80)$$

Observe that $P(T^+|A) = S^+$ and $P(T^+|N) = 1 - S^-$. In order to determine the posterior probability as above, the sensitivity and specificity of the test, as well as the prior probabilities of negative cases and positive cases (the rate of prevalence of the disease) should be known.

A cost matrix may be defined, as in Table 12.6, to reflect the overall cost-effectiveness of a test or method of diagnosis. The cost of conducting the test and arriving at a TN decision is indicated by C_N : this could be seen as the cost of subjecting an otherwise-normal subject to the test for the purposes of screening for a disease. The cost of the test when a TP is found is shown as C_A : this might include the costs of further tests, treatment, follow-up, etc., which are secondary to the test itself, but part of the screening and health-care program. The value C_{FP} indicates the cost of an FP result, that is, a false alarm: this represents the cost of erroneously subjecting an individual without the disease to further tests or therapy. Whereas it may be easy to identify the costs of clinical tests or treatment procedures, it is difficult to quantify

the traumatic and psychological effects of an FP result and the consequent procedures on a normal subject. The cost C_{FN} is the cost of an FN result: the presence of the disease in a patient is not diagnosed, the condition worsens with time, the patient faces more complications of the disease, and the health-care system or the patient has to bear the costs of further tests and delayed therapy.

A loss factor due to misclassification may be defined as

$$L = FPF \times C_{FP} + FNF \times C_{FN}. \tag{12.81}$$

The total cost of the screening program may be computed as

$$C_S = TPF \times C_A + TNF \times C_N + FPF \times C_{FP} + FNF \times C_{FN}. \tag{12.82}$$

Metz [1099] provides more details on the computation of the costs of diagnostic tests.

TABLE 12.6
Schematic Representation of the
Cost Matrix of a Diagnostic Method.

	Predicted group	
Actual group	Normal	Abnormal
Normal	C_N	C_{FP}
Abnormal	C_{FN}	C_A

Reproduced with permission from R.M. Rangayyan, *Biomedical Signal Analysis: A Case-Study Approach*, IEEE Press and Wiley, New York, NY. 2002, © IEEE.

12.8.1 Receiver operating characteristics

Measures of overall correct classification of patterns as percentages provide limited indications of the accuracy of a diagnostic method. The provision of a separate correct classification rate for each category, such as sensitivity and specificity, can facilitate improved analysis. However, these measures do not indicate the dependence of the results upon the decision threshold. Furthermore, the effect of the rate of incidence or prevalence of the particular disease is not considered.

From another perspective, it is desirable to have a screening or diagnostic test that is both highly sensitive and highly specific. In reality, however, such a test is usually not achievable. Most tests are based on clinical measurements that can assume limited ranges of a variable (or a few variables) with an inherent trade-off between sensitivity and specificity. The relationship between sensitivity and specificity is illustrated by the receiver operating characteristics (ROC) curve, which facilitates improved analysis of the classification accuracy of a diagnostic method [1099, 1100, 1101].

Consider the situation illustrated in Figure 12.19. For a given diagnostic test with the decision variable z , we have predetermined state-conditional PDFs of the decision variable z for actually negative or normal cases indicated as $p(z|N)$, and for actually positive or abnormal cases indicated as $p(z|A)$. As indicated in Figure 12.19, the two PDFs will almost always overlap, given that no method can be perfect. The user or operator needs to determine a decision threshold (indicated by the vertical line) so as to strike a compromise between sensitivity and specificity. Lowering the decision threshold will increase TPF at the cost of increased FPF . (Observe that TNF and FNF may be derived easily from FPF and TPF , respectively.)




FIGURE 12.19

State-conditional PDFs of a diagnostic decision variable z for normal and abnormal cases. The vertical line represents the decision threshold. Reproduced with permission from R.M. Rangayyan, *Biomedical Signal Analysis: A Case-Study Approach*, IEEE Press and Wiley, New York, NY. 2002, © IEEE.

An ROC curve is a graph that plots (FPF, TPF) points obtained for a range of decision threshold or cut points of the decision method (see Figure 12.20). The cut point could correspond to the threshold of the probability of prediction. By varying the decision threshold, we get different decision fractions, within the range $[0, 1]$. An ROC curve describes the inherent detection (diagnostic or discriminant) characteristics of a test or method: a receiver (user) may choose to operate at any point along the curve. The ROC curve is independent of the prevalence of the disease or disorder being investigated because it is based upon normalized decision fractions. Because all cases may be simply labeled as negative or all may be labeled as positive, an ROC curve has to pass through the points $(0, 0)$ and $(1, 1)$.




FIGURE 12.20

Examples of ROC curves. Reproduced with permission from R.M. Rangayyan, *Biomedical Signal Analysis: A Case-Study Approach*, IEEE Press and Wiley, New York, NY. 2002, © IEEE.

In a diagnostic situation where a human operator or specialist is required to provide the diagnostic decision, ROC analysis is usually conducted by requiring the specialist to rank each case as one of five possibilities [1099]:

1. definitely or almost definitely negative (normal),

2. probably negative,
3. possibly positive,
4. probably positive,
5. definitely or almost definitely positive (abnormal).

Item 3 above may be replaced by “indeterminate”, if desired. Various values of TPF and FPF are then calculated by varying the decision threshold from level 5 to level 1 according to the decision items listed above. The resulting (FPF, TPF) points are then plotted to form an ROC curve. The maximum likelihood estimation method [1102] is commonly used to fit a binormal ROC curve to data as above.

A summary measure of the effectiveness of a test is given by the area under the ROC curve, traditionally labeled as A_z . It is clear from Figure 12.20 that A_z is limited to the range $[0, 1]$. A test that gives a larger area under the ROC curve indicates a better method than one with a smaller area: in Figure 12.20, the method corresponding to ROC3 is better than the method corresponding to ROC2; both are better than the method represented by ROC1 with $A_z = 0.5$. An ideal method will have an ROC curve that follows the vertical line from $(0, 0)$ to $(0, 1)$, and then the horizontal line from $(0, 1)$ to $(1, 1)$, with $A_z = 1$: the method has $TPF = 1$ with $FPF = 0$, which is ideal. (*Note:* This would require the PDFs represented in Figure 12.19 to be nonoverlapping.) Examples of ROC curves are provided in Sections 12.10, 12.11, and 12.12.

In a form of ROC analysis known as the free-response ROC or FROC, the sensitivity is plotted against the number of false positives per image. See Section 8.10.7 and Figure 8.73 for an example of this type of analysis.

12.8.2 McNemar’s test of symmetry

Suppose we have available two methods to perform a certain diagnostic test. How may we compare the classification performance of one against that of the other?

Measures of overall classification accuracies such as a percentage of correct classification or the area under the ROC curve provide simple measures to compare two or more diagnostic methods. If more details are required as to how the classification of groups of cases varies from one method to another, McNemar’s test of symmetry [1103, 1104] would be an appropriate tool.

McNemar’s test is based on the construction of contingency tables that compare the results of two classification methods. The rows of a contingency table represent the outcomes of one of the methods used as the reference, possibly a gold standard (labeled as Method A in Table 12.7); the columns represent the outcomes of the other method, which is usually a new method (Method B) to be evaluated against the gold standard. The entries in the

table are counts that correspond to particular diagnostic categories, which in Table 12.7 are labeled as normal, indeterminate, and abnormal. A separate contingency table should be prepared for each true category of the patterns; for example, normal and abnormal. (The class “indeterminate” may not be applicable as a true category.) The true category of each case may have to be determined by a third method (for example, biopsy or surgery).

TABLE 12.7
Schematic Representation of a Contingency Table for
McNemar’s Test of Asymmetry.

Method A	Method B			Total
	Normal	Indeterminate	Abnormal	
Normal	a (1)	b (2)	c (3)	$R1$
Indeterminate	d (4)	e (5)	f (6)	$R2$
Abnormal	g (7)	h (8)	i (9)	$R3$
Total	$C1$	$C2$	$C3$	N

Reproduced with permission from R.M. Rangayyan, *Biomedical Signal Analysis: A Case-Study Approach*, IEEE Press and Wiley, New York, NY. 2002, © IEEE.

In Table 12.7, the variables $a, b, c, d, e, f, g, h,$ and i denote the counts in each cell, and the numbers in parentheses denote the cell number. The variables $C1, C2,$ and $C3$ denote the total numbers of counts in the corresponding columns; $R1, R2,$ and $R3$ denote the total numbers of counts in the corresponding rows. The total number of cases in the true category represented by the table is $N = C1 + C2 + C3 = R1 + R2 + R3$.

Each cell in a contingency table represents a paired outcome. For example, in evaluating the diagnostic efficiency of Method B versus Method A, cell number 3 will contain the number of samples that were classified as normal by Method A but as abnormal by Method B. The row totals $R1, R2,$ and $R3,$ and the column totals $C1, C2,$ and $C3$ may be used to determine the sensitivity and specificity of the two methods.

High values along the main diagonal (a, e, i) of a contingency table (see Table 12.7) indicate no significant change in diagnostic performance with Method B as compared to Method A. In a contingency table for truly abnormal cases,

a high value in the upper-right portion (cell number 3) will indicate an improvement in diagnosis (higher sensitivity) with Method B as compared to Method A. In evaluating a contingency table for truly normal cases, Method B will have a higher specificity than Method A if a large value is found in cell 7. McNemar's method may be used to perform detailed statistical analysis of improvement in performance based upon contingency tables if large numbers of cases are available in each category [1103, 1104]. Examples of contingency tables are provided in Section 12.10.

12.9 Reliability of Features, Classifiers, and Decisions

In most practical applications of biomedical image analysis, the researcher is presented with the problem of designing a pattern classification and decision-making system using a small number of training samples (images), with no knowledge of the distributions of the features or parameters computed from the images. The size of the training set, relative to the number of features used in the pattern classification system, affects the accuracy and reliability of the decisions made [1105, 1106, 1107]. One should not increase the number of features to be used without a simultaneous increase in the number of training samples, as the two quantities together affect the bias and variance of the classifier. On the other hand, when the training set has a fixed number of samples, the addition of more features beyond a certain limit will lead to poorer performance of the classifier: this is known as the "curse of dimensionality". The situation leads to "over-training": the classifier is trained to recognize the idiosyncrasies of the training set, and does not generalize or extend well to other data. It is desirable to be able to analyze the bias and variance of a classification rule while isolating the effects of the functional form of the distributions of the features used.

Raudys and Jain [1106] give a rule-of-thumb table for the number of training samples required in relation to the number of features used in order to remain within certain limits of classification errors for five pattern classification methods. When the available features are ordered in terms of their individual classification performance, the optimal number of features to be used with a certain classification method and training set may be determined by obtaining unbiased estimates of the classification accuracy with the number of features increased one at a time in order. A point will be reached when the performance deteriorates, which will indicate the optimal number of features to be used. This method, however, cannot take into account the joint performance of various combinations of features: exhaustive combinations of all features may have to be evaluated to take this aspect into consideration. Software packages such as the Statistical Package for the Social

Sciences (SPSS) [1092, 1093] provide programs to facilitate feature evaluation and selection, as well as the estimation of classification accuracies.

12.9.1 Statistical separability and feature selection

In practical applications of pattern classification, several parameters are often required in order to discriminate between multiple classes. Given the fact that most features provide for limited discrimination between classes due to the overlap in the ranges of their values for the various classes, it is natural to use several features. However, there would be costs associated with the measurement, derivation, and/or computation of each feature. It would be advantageous to be able to assess the contribution made by each feature toward the task of discriminating between the classes of interest, and to be able to select the feature set that provides the best separation between classes and the lowest classification error. The notion of statistical separability of features between classes is useful in addressing these concerns [1108].

Normalized distance between PDFs: Consider a feature x that has the means m_1 and m_2 and standard deviation values σ_1 and σ_2 for the two classes C_1 and C_2 . Assuming that the PDFs $p(x|C_1)$ and $p(x|C_2)$ overlap, the area of overlap is related to the error of classification. If the variances are held constant, the overlap between the PDFs decreases as $|m_1 - m_2|$ increases. If the means are held constant, the overlap increases as σ_1 and σ_2 increase (the dispersion of the features increases). These notions are captured by the normalized distance between the means, defined as [1108]

$$d_n = \frac{|m_1 - m_2|}{\sigma_1 + \sigma_2}. \quad (12.83)$$

The measure d_n provides an indicator of the statistical separability of the PDFs. A limitation of d_n , however, is that $d_n = 0$ if $m_1 = m_2$ regardless of σ_1 and σ_2 . Furthermore, the formulation above is valid only for a single feature x ; a generalization to a feature vector \mathbf{x} would be desirable.

Divergence: Let us rewrite the likelihood ratio in Equation 12.37 as

$$l_{ij}(\mathbf{x}) = \frac{p(\mathbf{x}|C_i)}{p(\mathbf{x}|C_j)}. \quad (12.84)$$

Applying the logarithm, we get

$$l'_{ij}(\mathbf{x}) = \ln[l_{ij}(\mathbf{x})] = \ln[p(\mathbf{x}|C_i)] - \ln[p(\mathbf{x}|C_j)]. \quad (12.85)$$

The divergence D_{ij} between the PDFs $p(\mathbf{x}|C_i)$ and $p(\mathbf{x}|C_j)$ is defined as [1108]

$$D_{ij} = E[l'_{ij}(\mathbf{x})|C_i] + E[l'_{ji}(\mathbf{x})|C_j], \quad (12.86)$$

where

$$E[l'_{ij}(\mathbf{x})|C_i] = \int_{\mathbf{x}} l'_{ij}(\mathbf{x}) p(\mathbf{x}|C_i) d\mathbf{x}. \quad (12.87)$$

Divergence has the following properties [1108]:

- $D_{ij} > 0$;
- $D_{ii} = 0$;
- $D_{ij} = D_{ji}$; and
- if the individual features x_1, x_2, \dots, x_n are statistically independent, $D_{ij}(x_1, x_2, \dots, x_n) = \sum_{k=1}^n D_{ij}(x_k)$.

It follows that adding more features that are statistically independent of one another will increase divergence and statistical separability.

In the case of multivariate Gaussian PDFs, we have [1108]

$$D_{ij} = \frac{1}{2} \text{Tr}[(\mathbf{C}_i - \mathbf{C}_j)(\mathbf{C}_j^{-1} - \mathbf{C}_i^{-1})] + \frac{1}{2} \text{Tr}[(\mathbf{C}_i^{-1} + \mathbf{C}_j^{-1})(\mathbf{m}_i - \mathbf{m}_j)(\mathbf{m}_i - \mathbf{m}_j)^T]. \quad (12.88)$$

The second term in the equation above is similar to the normalized distance d_n as defined in Equation 12.83, and becomes zero for PDFs with identical means; however, due to the first term, $D_{ij} \neq 0$ unless the covariance matrices are identical.

In the case of the existence of multiple classes $C_i, i = 1, 2, \dots, m$, the pairwise divergence values may be averaged to obtain a single measure across all of the m classes as [1108]

$$D_{\text{av}} = \sum_{i=1}^m \sum_{j=1}^m p(C_i) p(C_j) D_{ij}. \quad (12.89)$$

A limitation of both d_n and D_{ij} is that they increase without an upper bound as the separation between the means increases. On the other hand, the error of classification is limited to the range 0 – 100% or $[0, 1]$.

Jeffries–Matusita distance: The Jeffries–Matusita (JM) distance provides an improved measure of the separation between PDFs than the normalized distance and divergence. The JM distance between the PDFs $p(\mathbf{x}|C_i)$ and $p(\mathbf{x}|C_j)$ is defined as [1108]

$$J_{ij} = \left\{ \int_{\mathbf{x}} \left[\sqrt{p(\mathbf{x}|C_i)} - \sqrt{p(\mathbf{x}|C_j)} \right]^2 d\mathbf{x} \right\}^{1/2}. \quad (12.90)$$

In the case of multivariate Gaussian PDFs, we have [1108]

$$J_{ij} = \sqrt{2[1 - \exp(-\alpha)]}, \quad (12.91)$$

where

$$\alpha = \frac{1}{8} (\mathbf{m}_i - \mathbf{m}_j)^T \left(\frac{\mathbf{C}_i + \mathbf{C}_j}{2} \right)^{-1} (\mathbf{m}_i - \mathbf{m}_j) + \frac{1}{2} \ln \left[\frac{|(\mathbf{C}_i + \mathbf{C}_j)|/2}{(|\mathbf{C}_i||\mathbf{C}_j|)^{1/2}} \right], \quad (12.92)$$

where $|\mathbf{C}_i|$ is the determinant of \mathbf{C}_i .

An advantage of the JM distance is that it is limited to the range $[0, \sqrt{2}]$. $J_{ij} = 0$ when the means of the PDFs are equal and the covariance matrices are zero matrices. Pairwise JM distances may be averaged over multiple classes similar to the averaging of divergence as in Equation 12.89. The JM distance determines the upper and lower bounds on the error of classification [1108].

It should be observed that divergence and JM distance are defined for a given feature vector \mathbf{x} . The measures would have to be computed for all combinations of features in order to select the best feature set for a particular pattern-classification problem.

12.10 Application: Image Enhancement for Breast Cancer Screening

Accurate detection of breast cancer depends upon the quality of mammograms; in particular, on the visibility of small, low-contrast objects within the breast image. Unfortunately, the contrast between malignant tissue and normal tissue is often low in mammograms, making the detection of malignant patterns difficult. Contrast between malignant tissue and normal dense tissue may be present on a mammogram, but be below the threshold of human perception. As well, microcalcifications in a sufficiently dense mass may not be readily visible because of low contrast. Hence, the fundamental enhancement needed in mammography is an increase in contrast, especially for dense breasts. Although many enhancement techniques reported are able to enhance specific details, they may also produce disturbing artifacts; see Chapter 4, in particular, Section 4.11.

It is important to distinguish between the evaluation of the detection of the presence of features such as microcalcifications in an image, and the evaluation of the diagnostic conclusion about a subject. Whereas some enhancement techniques may enhance the visibility of features such as calcifications, they may also distort their appearance and shape characteristics, which may lead to misdiagnosis; see Morrow et al. [123] for a discussion on this topic. A similar observation was made by Kimme-Smith et al. [264], who stated that “Studies of digitally enhanced mammograms should examine the actual ability to form diagnostic conclusions from the enhanced images, rather than the ability merely to report the increased numbers of clusters of simulated microcalcifications that it is possible to detect. Radiologic evaluation obviously begins with the detection of an abnormality, but if the image of the abnormality is distorted, an incorrect diagnosis may result.” ROC analysis and McNemar’s method could assist in assessing the effect of enhancement on the diagnosis.

In their ROC study to evaluate the effects of digitization and unsharp-mask filtering on the detection of calcifications, Chan et al. [254] used 12 images with calcifications and 20 normal images. The digitization was performed at a spatial resolution of 0.1 mm per pixel, and the enhanced images were printed on film. Nine radiologists interpreted the images. They found that the detectability of calcifications in the digitized mammograms was improved by unsharp-mask filtering, although both the unprocessed digitized and the processed mammograms provided lower accuracy than the conventional mammograms.

Kimme-Smith et al. [264] compared contact, magnified, and TV-enhanced mammographic images of 31 breasts for diagnosis of calcifications. The interpretation was performed by three experienced radiologists and three radiology residents. The TV enhancement procedure used the Wallis filter, which is similar to unsharp masking. They concluded that TV enhancement could not replace microfocus magnification and could lead to misdiagnosis by inexperienced radiologists. Experienced radiologists showed no significant improvement in performance with the enhanced images.

Nab et al. [1109] performed ROC analysis comparing 270 mammographic films with $2K \times 2K$ 12-bit digitized versions (at 0.1 mm per pixel) displayed on monitors. The task for the two radiologists in the study was to indicate the presence or absence of tumors or calcifications. No significant difference in performance was observed between the use of films and their digitized versions.

Kallergi et al. [267] conducted an ROC study with 100 mammograms and four radiologists, including the original films, digitized images (105 μm pixel size) displayed on monitors, and wavelet-enhanced images displayed on monitors (limited to 8-bit gray scale). The diagnostic task was limited to the detection and classification of calcifications. While they observed a statistically significant reduction in the area under the ROC curve with the digitized images, the difference between reading the original films and the wavelet-enhanced images displayed on monitors was not significant. They also noted that interobserver variation was reduced with the use of the wavelet-enhanced images. They concluded that filmless mammography with their wavelet-based enhancement method is comparable to screen-film mammography for detecting and classifying calcifications.

The ANCE method (described in Sections 4.9.1 and 4.11) was used in a preference study comparing the performance of mammographic enhancement algorithms [125]. The other methods used in the study were adaptive unsharp masking, contrast-limited adaptive histogram equalization, and wavelet-based enhancement. In a majority of the cases with microcalcifications, the ANCE algorithm provided the most preferred results. In the set of images with masses, the unenhanced images were preferred in most of the cases. Rangayyan et al. [124, 266, 271, 322] evaluated the role of the ANCE technique for enhancement of mammograms in a breast cancer screening program using ROC and McNemar's methods. The methods and results of these works are described in detail in the following sections.

12.10.1 Case selection, digitization, and presentation

In order to evaluate the diagnostic utility of the ANCE technique, two ROC studies were conducted using two different datasets: *difficult cases* and *interval-cancer cases*. The difficult-cases dataset is a collection of cases for which the radiologist had been unsure enough to call for a biopsy. The cases were difficult in terms of both the detection of the abnormality present and the diagnosis as normal, benign, or malignant. The investigation described in this section was conducted to test if ANCE could be used to improve the distinction between benign and malignant cases [266].

Interval-cancer cases are cases in a screening program where cancer is detected prior to a scheduled return screening visit; they may be indicative of the inability to detect an already present cancer or an unusually rapid-growing cancer. In these cases, the radiologist had declared that there was no evidence of cancer on the previous mammograms. The purpose of the study, described in this section, was to test if interval cancers could be detected earlier with appropriate digital enhancement and analysis [271]. The goal of interpretation in the study was screening, and not the detection of signs such as calcifications or masses; as such, no record was maintained of the number or sizes of the signs, as done by Kallergi et al. [267] and Nishikawa et al. [1110]. Localization of pathology was not required: the radiologists had to find lesions, if any, and assess them for the likelihood of malignancy, but did not have to mark their locations on the films.

Difficult cases: An experienced radiologist selected 21 difficult cases, related to 14 subjects with benign breast disease and seven subjects with malignant disease, from files over the period 1987 – 1992 at the Foothills Hospital, Calgary, Alberta, Canada. Four films, including the MLO and CC views of each breast, were available for each of 18 cases, but only two films of one breast each were available for three cases, leading to a total of 78 screen-film mammograms. Biopsy results were also available for each subject.

Each film was digitized using an Eikonix 1412 scanner (Eikonix Inc., Bedford, MA) to 4,096 by about 2,048 pixels with 12-bit gray-scale resolution. (The size of the digitized image differed from film to film depending upon the the size of the actual image in the mammogram.) Sampling as above represents a spot size on the film of about $0.062\text{ mm} \times 0.062\text{ mm}$. Films were illuminated by a Plannar 1417 light box (Gordon Instruments, Orchard Park, NY). Although the light box is designed to have a uniform light intensity distribution, it was necessary to correct for nonuniformities in illumination. After correction, pixel gray levels were determined to be accurate to 10 bits, with a dynamic range of approximately $0.02 - 2.52\text{ OD}$ [174].

The digital images were down-sampled by a factor of two for processing and display for interpretation on a Megascan 2111 monitor (Advanced Video Products Inc., Littleton, MA). Although the memory buffer of the Megascan system is of size $4,096 \times 4,096 \times 12$ bits, the display buffer is limited to $2,560 \times 2,048 \times 8$ bits, with panning and zooming facilities. The original

screen-film mammograms used in the study were presented to the radiologists using a standard mammogram film viewer.

Six radiologists from the Foothills Hospital interpreted the original, the unprocessed digitized, and the enhanced mammograms separately. Only one of the radiologists had prior experience with digitized and enhanced mammographic images. The images were presented in random order and the radiologists were given no additional information about the patients. The radiologists ranked each case as (1) definitely or almost definitely benign, (2) probably benign, (3) possibly malignant, (4) probably malignant, or (5) definitely or almost definitely malignant.

Interval-cancer cases: Two hundred and twenty-two screen-film mammograms of 28 interval-cancer patients and six control patients with benign breast disease were selected for this study from files over the period 1991–1995 at the Screen Test Centres of the Alberta Program for the Early Detection of Breast Cancer [61]. Some of the cases of cancer were diagnosed by physical exam or mammography performed after the preceding visit to the screening program but prior to the next scheduled visit. The radiologists who interpreted the mammograms taken prior to the diagnosis of the cancer had declared that there was no evidence of cancer on the films. The small number of benign cases were included to prevent “over-diagnosis”; the radiologists were not informed of the proportion of benign to malignant cases in the dataset. Most of the files included multiple sets of films taken at different times; all sets except one included at least four films each (the MLO and CC views of each breast) in the dataset. (More specifically, the dataset included fifty-two 4-film sets, one 3-film set, one 5-film set, and one 6-film set.) Previous films of all of the interval-cancer cases had initially been reported as being normal. Biopsy results were available for each subject.

The aim of this study was to investigate the possibility of earlier detection of interval breast cancers with the aid of appropriate image processing techniques. Because a few sets of films taken at different times were available for each subject, each set of mammograms of each subject was labeled as a separate case. All films of the subjects with malignant disease within the selected period were labeled as being malignant, even though the cases had not been previously interpreted as such. By this process, 55 cases were obtained, of which 47 were malignant and eight were benign (the numbers of subjects being 28 with malignant disease and six with benign disease).

The films were digitized as described previously in this section, and processed using the ANCE technique with the full digitized resolution available. The digitized version and the ANCE-processed version were printed on film using a KODAK XL 7700 digital continuous-tone printer (Eastman Kodak Company, Rochester, NY) with pixel arrays up to $2,048 \times 1,536$ (8-bit pixels). Gray-level remapping (10 bits/pixel to 8 bits/pixel) and down-sampling by a factor of two were applied before a digitized/enhanced mammogram image was sent for printing with two different LUTs.

Three reference radiologists from the Screen Test Program separately interpreted the original films of the involved side, their digitized versions, and their ANCE-processed versions on a standard mammogram film viewer. Only one of the radiologists had prior experience with digitized and enhanced mammographic images. Interpretation of the digitized images (without ANCE processing) was included in the test to evaluate the effect on diagnostic accuracy of digitization and printing with the resolution and equipment used. The images were presented in random order, and the radiologists were given no additional information about the patients. The radiologists ranked each case as (1) definitely or almost definitely benign, (2) probably benign, (3) indeterminate, (4) probably malignant, or (5) definitely or almost definitely malignant. Note that the diagnostic statement for rank (3) is different in this study from that described previously in this section.

Images were interpreted in random order by the radiologists; images were presented in the same random order to each radiologist individually. Each radiologist interpreted all of the images in a single sitting. Multiple sets of films of a given subject (taken at different times) were treated as different cases and interpreted separately to avoid the development of familiarity and bias. The original, digitized, and enhanced versions of any given case were mixed for random ordering, treated as separate cases, and were interpreted separately to prevent the development of familiarity and bias. All available views of a case were read together as one set. It should be recognized that the initial (original) diagnosis of the cases was performed by different teams of radiologists experienced in the interpretation of screening mammograms, which further limits the scope of bias in the study being described.

12.10.2 ROC and statistical analysis

ROC analysis [1100, 1101] was used to compare the radiologists' performance in detecting abnormalities in the various images. The maximum likelihood estimation method [1102] was used to fit a binormal ROC curve to each radiologist's confidence rating data for each set of mammograms. The slope and intercept parameters of the binormal ROC curve (when plotted on normal probability scales) were calculated for each fitted curve. To estimate the average performance of the group of radiologists on each set of images, composite ROC curves [1101] were calculated by averaging the slope and the intercept parameters of the individual ROC curves. Finally, the area under the binormal ROC curve (as plotted in the unit square) was computed, which represents the overall abnormality detection accuracy for each type of images.

In addition to ROC analysis of the interval-cancer cases, McNemar's test of symmetry [1103, 1104] was performed on a series of 3×3 contingency tables obtained by cross-tabulating (i) diagnostic confidence using the original mammograms (categories 1 or 2, 3, and 4 or 5) against the diagnostic confidence using the digitized mammograms, and (ii) diagnostic confidence using the digitized mammograms against the diagnostic confidence using the en-

hanced mammograms. Separate 3×3 tables, as illustrated in Figure 12.21, were formed for the malignant cases and the benign cases. Cases in which there is no change in the diagnostic confidence will fall on the diagonal (upper left to lower right, labeled as D in Figure 12.21) of the table. For the malignant cases, improvement in the diagnostic accuracy is illustrated by a 3×3 table with the majority of the cases in the three upper right-hand cells (labeled as U in Figure 12.21). Conversely, for the benign cases, improvement in the diagnostic accuracy will be illustrated by a 3×3 table with the majority of the cases in the three lower left-hand cells (labeled as L in Figure 12.21). The hypothesis of significant improvement can be tested statistically using McNemar's test of symmetry [1103, 1104], namely that the probability of an observation being classified into a cell $[i, j]$ is the same as the probability of an observation being classified into the cell $[j, i]$.

		Level with Enhanced Methodology		
		1 or 2	3	4 or 5
Level with Original Methodology	1 or 2	D	U	U
	3	L	D	U
	4 or 5	L	L	D

FIGURE 12.21

Illustration of the contingency table for McNemar's test. Figure courtesy of L. Shen [320].

The validity of McNemar's test depends on the assumption that the cell counts are at least moderately large. In order to avoid the limitations due to this factor, and also to avoid the problem of excessive multiple comparisons, the data across the individual radiologists were combined in two different ways before applying McNemar's test. The first method (referred to as "averaged") averaged the radiologists' diagnostic ratings before forming the 3×3 tables. In the second method (referred to as "combined"), the 3×3 tables for each of the radiologists were formed first and then combined by summing the corresponding cells.

Because this analysis involves multiple p -values, the Bonferroni correction was used to adjust the p -values [1111]. When multiple p -values are produced, the probability of making a Type I error increases. (Rejection of the null hy-

pothesis when it is true is called a Type I error.) The Bonferroni method for adjusting multiple p -values requires that when k hypothesis tests are performed, each p -value be multiplied by k , so that the adjusted p -value is $p^* = kp$. In order to reduce the number of p -values, symmetry was tested for each situation (malignant/benign and averaged/combined) for the two tables original-to-digitized and digitized-to-enhanced, but not the original-to-enhanced (which follows from the other two).

ROC analysis of difficult cases: Because the population involved in this study was such that the original mammograms were sufficiently abnormal to cause the initial attending radiologist to call for biopsy, the aim was to test whether specificity could be improved with the ANCE method.

The composite ROC curves representing breast cancer diagnosis by the six radiologists in this study are compared in Figure 12.22, which illustrates several points: First, the process of digitization (and down-sampling to an effective pixel size of $0.124 \text{ mm} \times 0.124 \text{ mm}$) degraded the quality of the images and thereby made the radiologists' performance worse, especially in the low-FPF range. However, better performance of the radiologists is seen with the digitized images at high FPFs (better sensitivity with worse specificity). Second, the ANCE method improved the radiologists' performance in all ranges of FPF (more significantly in the low-FPF range) as compared with the unprocessed digitized images, although it is still lower than that with the original films in the low range of FPF. The A_z values for the original, digitized, and enhanced mammograms were computed to be 0.67, 0.63, and 0.67, respectively. (Kallergi et al. [267] also observed a drop in the area under the ROC curve when digitized images were interpreted from monitor display as compared with the original films; their wavelet-based enhancement method provided an improvement over the digitized version, although the enhanced images did not provide any statistically significant benefit over the original films.) The A_z values are lower than those normally encountered in most studies (in the range 0.85 – 0.95) due to the fact that the cases selected were difficult enough to call for biopsy. The labeling of all mammograms taken prior to the detection of cancer as abnormal will also have had a bearing on this result. Regardless, the numerical results indicate that the ANCE technique improved the radiologists' overall performance, especially over unprocessed digitized mammograms, and allowed the radiologists to discriminate between the two populations slightly better while interpreting the enhanced mammograms as compared with the original films.

McNemar's tests on difficult cases: Table 12.8 and Table 12.9 contain details of the variation of the radiologists' diagnostic performance for malignant cases and benign cases, respectively, with the difficult-cases dataset. (*Note:* In the tables, B refers to benign, U to undecided or indeterminate, and M to malignant ratings.) For almost every table for individual readers, the numbers were too small to perform the McNemar chi-square test (including the average). In other cases, the numbers would be too small to detect a




FIGURE 12.22

Comparison of composite ROC curves for the detection of abnormalities by interpreting the original, unprocessed digitized, and enhanced images of 21 difficult cases. Reproduced with permission from R.M. Rangayyan, L. Shen, Y. Shen, J.E.L. Desautels, H. Bryant, T.J. Terry, N. Horeczko, and M.S. Rose, "Improvement of sensitivity of breast cancer diagnosis with adaptive neighborhood contrast enhancement of mammograms", *IEEE Transactions on Information Technology in Biomedicine*, 1(3):161-170, 1997. © IEEE.

TABLE 12.8

Variation of Radiologists' Diagnostic Performance with Original, Unprocessed Digitized, and ANCE-processed Digitized Mammograms for the Seven Malignant Cases in the Difficult-cases Dataset.

Rad. Images		Change in diagnostic confidence level					
		B : level 1 or 2;		U : level 3;		M : level 4 or 5	
		B → U (U → B)	U → M (M → U)	B → M (M → B)	B → B	U → U	M → M
#1	O → D	1 (1)	0 (2)	0 (0)	1	0	2
	D → E	1 (0)	1 (0)	0 (0)	1	2	2
	O → E	0 (0)	0 (2)	1 (0)	1	1	2
#2	O → D	0 (0)	0 (1)	0 (0)	1	1	4
	D → E	0 (0)	0 (0)	0 (0)	1	2	4
	O → E	0 (0)	0 (1)	0 (0)	1	1	4
#3	O → D	1 (0)	1 (2)	0 (0)	0	2	1
	D → E	0 (1)	0 (0)	0 (0)	0	4	2
	O → E	1 (1)	1 (2)	0 (0)	0	1	1
#4	O → D	0 (1)	0 (1)	0 (1)	1	0	3
	D → E	1 (0)	1 (0)	0 (0)	2	0	3
	O → E	0 (0)	0 (0)	0 (1)	1	1	4
#5	O → D	0 (0)	1 (0)	0 (0)	1	2	3
	D → E	1 (0)	1 (1)	0 (0)	0	1	3
	O → E	1 (0)	2 (1)	0 (0)	0	1	2
#6	O → D	0 (1)	0 (0)	0 (2)	1	0	3
	D → E	0 (0)	0 (1)	2 (0)	2	0	2
	O → E	0 (0)	1 (1)	0 (1)	1	0	3
Av.	O → D	0 (0)	0 (2)	0 (0)	1	1	3
	D → E	0 (0)	1 (0)	0 (0)	1	2	3
	O → E	0 (0)	1 (2)	0 (0)	1	0	3

Rad. = Radiologist; O = original mammogram; D = unprocessed digitized mammogram; E = ANCE-processed digitized mammogram. The average (Av.) values were obtained by averaging the individual confidence levels [320].

TABLE 12.9

Variation of Radiologists' Diagnostic Performance with Original, Unprocessed Digitized, and ANCE-processed Digitized Mammograms for the 14 Benign Cases in the Difficult-cases Dataset.

Rad. Images		Change in diagnostic confidence level					
		B : level 1 or 2;		U : level 3;		M : level 4 or 5	
		U → B (B → U)	M → U (U → M)	M → B (B → M)	B → B	U → U	M → M
#1	O → D	3 (1)	2 (0)	0 (0)	2	3	3
	D → E	1 (2)	0 (1)	0 (1)	2	4	3
	O → E	2 (1)	2 (1)	0 (1)	1	3	3
#2	O → D	1 (0)	0 (1)	0 (0)	5	4	3
	D → E	0 (0)	0 (0)	0 (0)	6	4	4
	O → E	1 (0)	0 (1)	0 (0)	5	4	3
#3	O → D	3 (1)	1 (2)	0 (1)	4	1	1
	D → E	1 (1)	1 (0)	0 (0)	6	2	3
	O → E	2 (1)	1 (2)	0 (0)	5	2	1
#4	O → D	4 (0)	0 (1)	2 (0)	5	1	1
	D → E	0 (2)	0 (0)	0 (0)	9	1	2
	O → E	3 (1)	0 (1)	2 (0)	4	2	1
#5	O → D	3 (1)	1 (1)	1 (1)	2	2	2
	D → E	2 (0)	1 (1)	1 (1)	5	1	2
	O → E	3 (0)	1 (2)	1 (0)	4	1	2
#6	O → D	5 (0)	0 (1)	2 (0)	5	0	1
	D → E	0 (1)	0 (0)	0 (2)	9	0	2
	O → E	4 (1)	0 (2)	2 (1)	3	0	1
Av.	O → D	5 (1)	1 (1)	1 (0)	3	1	1
	D → E	1 (1)	0 (1)	0 (0)	8	1	2
	O → E	4 (0)	1 (2)	1 (0)	4	1	1

Rad. = Radiologist; O = original mammogram; D = unprocessed digitized mammogram; E = ANCE-processed digitized mammogram. The average (Av.) values were obtained by averaging the individual confidence levels [320].

statistically significant difference. Therefore, the data were combined for the six readers by simply summing the corresponding matrices.

For the benign cases (combined), p -values of 0.004, 0.53, and 0.022 were obtained for original-to-digitized, digitized-to-enhanced, and original-to-enhanced, respectively. For the malignant cases (combined), the p -values were 0.16, 0.36, and 0.69 for original-to-digitized, digitized-to-enhanced, and original-to-enhanced, respectively. The p -values represent no evidence of improvement in the diagnostic accuracy for any of the three tables (original-to-digitized, digitized-to-enhanced, and original-to-enhanced) for the malignant cases in the difficult-cases dataset. However, for the benign cases, there was a statistically significant improvement in the diagnostic accuracy ($p = 0.004$, Bonferroni adjusted value $p^* = 0.024$). There was no evidence of a significant improvement from digitized to enhanced, and although there was a significant improvement from the original to the digitized category (but not significant after Bonferroni adjustment, $p^* = 0.13$), this was attributed to the improvement in moving from the original to the digitized category.

ROC analysis of interval-cancer cases: Figure 12.23 shows the variation of the ROC curves among the three radiologists who interpreted the same set of unprocessed digitized mammograms of the interval-cancer cases. Similar variation was observed with the sets of the original film mammograms and the enhanced mammograms. Details of the variation of the radiologists' diagnostic performance with the original mammograms, unprocessed digitized mammograms, and ANCE-processed mammograms are listed in Table 12.10 and Table 12.11 for the 47 malignant and eight benign cases, respectively.

It is seen from Table 12.10 that, on the average (average of individual diagnostic confidence levels), almost half (21) of the 47 malignant cases, which were originally diagnosed as benign (average diagnostic confidence level of less than 2.5) by the three radiologists with the original films were relabeled as malignant (average diagnostic confidence level of greater than 3.5) with the ANCE-processed versions. Only three malignant cases whose original average diagnostic confidence levels were greater than 3.5 had their average confidence levels reduced to the range of 2.5 to 3.5 when interpreting the enhanced mammograms. However, in general, no significant changes are observed for the benign cases (Table 12.11) with the ANCE procedure.

Composite ROC curves for breast cancer diagnosis with the original, unprocessed digitized, and enhanced images are plotted in Figure 12.24. The following points may be observed in Figure 12.24: First, the radiologists' performance with the enhanced versions is the best among the three, especially for $FPF > 0.3$. This is reasonable, because most of the cancer cases in this dataset were difficult and were initially diagnosed as normal when interpreting the original films. Therefore, the FPF level has to be increased in order to achieve good sensitivity (high TPF). Second, the digitized versions appear to provide better diagnostic results when compared with the original films. This is likely due to the fact that two printouts for each digitized image with two different LUTs (unchanged and lighten2) were provided to the radiolo-




FIGURE 12.23

Variation of conventional ROC curves among three radiologists interpreting the same set of unprocessed digitized mammograms from the interval-cancer cases dataset. Reproduced with permission from R.M. Rangayyan, L. Shen, Y. Shen, J.E.L. Desautels, H. Bryant, T.J. Terry, N. Horeczko, and M.S. Rose, "Improvement of sensitivity of breast cancer diagnosis with adaptive neighborhood contrast enhancement of mammograms", *IEEE Transactions on Information Technology in Biomedicine*, 1(3):161–170, 1997. © IEEE.

TABLE 12.10

Variation of Radiologists' Diagnostic Performance with Original, Unprocessed Digitized, and ANCE-processed Digitized Mammograms for the 47 Malignant Cases in the Interval-cancer Dataset.

		Change in diagnostic confidence level					
		B : level 1 or 2;		U : level 3;	M : level 4 or 5		
Rad.	Images	B → U	U → M	B → M	B →	U →	M →
		(U → B)	(M → U)	(M → B)	B	U	M
#1	O → D	8 (0)	4 (0)	15 (1)	2	1	16
	D → E	0 (0)	9 (0)	2 (0)	0	1	35
	O → E	1 (0)	5 (1)	24 (0)	0	0	16
#2	O → D	8 (0)	7 (0)	5 (0)	4	7	16
	D → E	1 (1)	12 (2)	3 (0)	0	3	25
	O → E	4 (1)	13 (1)	13 (0)	0	0	15
#3	O → D	7 (1)	0 (1)	4 (3)	9	6	16
	D → E	5 (2)	8 (3)	2 (1)	6	4	16
	O → E	6 (0)	4 (3)	8 (3)	6	3	14
Av.	O → D	9 (0)	2 (4)	10 (0)	4	4	14
	D → E	1 (0)	14 (2)	3 (0)	0	3	24
	O → E	2 (0)	5 (3)	21 (0)	0	1	15

Rad. = Radiologist; O = original mammogram; D = unprocessed digitized mammogram; E = ANCE-processed digitized mammogram. The average (Av.) values were obtained by averaging the individual confidence levels. Reproduced with permission from R.M. Rangayyan, L. Shen, Y. Shen, J.E.L. Desautels, H. Bryant, T.J. Terry, N. Horeczko, and M.S. Rose, "Improvement of sensitivity of breast cancer diagnosis with adaptive neighborhood contrast enhancement of mammograms", *IEEE Transactions on Information Technology in Biomedicine*, 1(3):161-170, 1997. © IEEE.

TABLE 12.11

Variation of Radiologists' Diagnostic Performance with Original, Unprocessed Digitized, and ANCE-processed Digitized Mammograms for the Eight Benign Cases in the Interval-cancer Dataset.

		Change in diagnostic confidence level					
		B : level 1 or 2;		U : level 3;	M : level 4 or 5		
Rad.	Images	U → B (B → U)	M → U (U → M)	M → B (B → M)	B → B	U → U	M → M
#1	O → D	0 (0)	1 (4)	0 (0)	1	0	2
	D → E	0 (1)	1 (0)	0 (0)	0	1	5
	O → E	0 (1)	1 (3)	0 (0)	0	1	2
#2	O → D	0 (0)	1 (0)	0 (0)	1	2	4
	D → E	0 (1)	0 (1)	0 (0)	0	2	4
	O → E	0 (1)	1 (1)	0 (0)	0	1	4
#3	O → D	0 (0)	1 (0)	0 (1)	2	1	1
	D → E	0 (0)	0 (1)	0 (1)	1	1	1
	O → E	0 (0)	0 (0)	0 (2)	1	1	2
Av.	O → D	0 (0)	1 (2)	0 (0)	1	1	3
	D → E	0 (1)	1 (1)	0 (0)	0	1	4
	O → E	0 (1)	1 (2)	0 (0)	0	1	3

Rad. = Radiologist; O = original mammogram; D = unprocessed digitized mammogram; E = ANCE-processed digitized mammogram. The average (Av.) values were obtained by averaging the individual confidence levels. Reproduced with permission from R.M. Rangayyan, L. Shen, Y. Shen, J.E.L. Desautels, H. Bryant, T.J. Terry, N. Horeczko, and M.S. Rose, "Improvement of sensitivity of breast cancer diagnosis with adaptive neighborhood contrast enhancement of mammograms", *IEEE Transactions on Information Technology in Biomedicine*, 1(3):161-170, 1997. © IEEE.

gists; the lighten2 LUT (see Figure 12.25) provided by Kodak performs some enhancement. Two print LUTs were used as the radiologists did not favor the use of the hyperbolic tangent (sigmoid) function, which is an approximate model of an X-ray film system, during initial setup tests.




FIGURE 12.24

Comparison of composite ROC curves for the detection of abnormalities by interpreting the original, unprocessed digitized, and enhanced images from the interval-cancer dataset. Reproduced with permission from R.M. Rangayyan, L. Shen, Y. Shen, J.E.L. Desautels, H. Bryant, T.J. Terry, N. Horeczko, and M.S. Rose, "Improvement of sensitivity of breast cancer diagnosis with adaptive neighborhood contrast enhancement of mammograms", *IEEE Transactions on Information Technology in Biomedicine*, 1(3):161-170, 1997. © IEEE.

The A_z values for the original, digitized, and enhanced mammograms were computed to be 0.39, 0.47, and 0.54, respectively. These numbers are much lower than the commonly encountered area values due to the fact that the cases selected are difficult cases, and more importantly, due to the fact that signs of earlier stages of the interval cancers were either not present on the previous films or were not visible. The radiologists interpreting the mammograms taken prior to the diagnosis of the cancer had declared that there was no evidence of cancer on the films; hence, the improvement indicated by the ROC curve is significant in terms of the diagnostic outcome. This also explains why A_z is less than 0.5 for the original and digitized mammograms.




FIGURE 12.25

The two LUTs used for printing mammograms. Reproduced with permission from R.M. Rangayyan, L. Shen, Y. Shen, J.E.L. Desautels, H. Bryant, T.J. Terry, N. Horeczko, and M.S. Rose, “Improvement of sensitivity of breast cancer diagnosis with adaptive neighborhood contrast enhancement of mammograms”, *IEEE Transactions on Information Technology in Biomedicine*, 1(3):161–170, 1997. © IEEE.

The results indicate that the ANCE technique can improve sensitivity and assist radiologists in detecting breast cancer at earlier stages.

McNemar’s tests on interval-cancer cases: For the benign cases (averaged), and for the original-to-enhanced, the numbers were too small to provide a valid chi-square statistic for McNemar’s test. Therefore, for the benign cases, two tables (digitized-to-enhanced and original-to-enhanced) combined over the three radiologists were tested. No significant difference in diagnostic accuracy was found for the benign cases, for either the digitized-to-enhanced table with $p = 0.097$ (Bonferroni adjusted p -value $p^* = 0.58$) or for the original-to-enhanced table with $p = 0.083$ ($p^* = 0.50$).

For each of the four tables for the malignant cases, a significant improvement was observed in the diagnostic accuracy, with the following p -values: original-to-digitized (combined) $p < 0.001$, $p^* < 0.001$; digitized-to-enhanced (combined) $p = 0.0001$, $p^* = 0.0006$; original-to-digitized (averaged) $p = 0.002$, $p^* = 0.012$; digitized-to-enhanced (averaged) $p = 0.0046$, $p^* = 0.028$.

In summary, no significant changes were seen in the diagnostic accuracy for the benign control cases. For the malignant cases, a significant improvement was seen in the diagnostic accuracy in all four tables tested, even after using a Bonferroni adjustment for the multiple p -values ($k = 6$).

12.10.3 Discussion

The results of the interval-cancer study indicate that the ANCE method had a positive impact on the interpretation of mammograms in terms of early detection of breast cancer (improved sensitivity). The ANCE-processed mammograms increased the detectability of signs of malignancy at earlier stages (of the interval-cancer cases) as compared with the original and unprocessed digitized mammograms. In terms of the average diagnostic confidence levels of three experts, 19 of 28 interval-cancer patients were not diagnosed during their earlier mammography tests with the original films only. However, had the ANCE procedure been used, all of these cases would have been diagnosed as malignant at the corresponding earlier times. Only one of six patients initially labeled as having benign disease with the original mammogram films was interpreted as malignant after enhancement. Although the resultant high sensitivity (TPF) comes with increased FPF of over 0.3, such an improvement in the detection of breast cancer at early stages is important.

The results obtained with the set of difficult cases are not as conclusive as the results with the interval-cancer cases. Three reasons for this could be (i) lack of familiarity of five of the six radiologists with digitized and enhanced mammographic images; (ii) interpreting the images on a monitor; and (iii) use of down-sampled images at a lower resolution of $124 \mu\text{m}$. Better results may be achieved if mammograms are digitized and processed with the desired spatial resolution of $50 \mu\text{m}$ and dynamic range of $0 - 3.5 OD$, and printed at the same resolution on film. (No monitor is as yet available to display images of the order of $4,096 \times 4,096$ pixels at 12 bits/pixel.)

The results of statistical analysis using McNemar's tests show (more conclusively than ROC analysis) that the ANCE procedure resulted in a statistically significant improvement in the diagnosis of interval-cancer cases, with no significant effect on the benign control cases. Statistical tests such as McNemar's test complement ROC analysis in certain circumstances, such as those in the study being discussed with small numbers of difficult cases. Both methodologies are useful as they analyze the results from different perspectives: ROC analysis provides a measure of the accuracy of the procedure in terms of sensitivity and specificity, whereas McNemar's test analyzes the statistical significance and consistency of the change (improvement) in performance. ROC analysis could include a chi-square test of statistical significance, if large numbers of cases are available.

In the study with the difficult-cases dataset, both the ROC and statistical analysis using McNemar's tests have shown that the digital versions led to some improvements in distinguishing benign cases from malignant cases (specificity). However, the improvement in the unprocessed digitized mammograms may have come from the availability of a zooming utility.

Although the ANCE algorithm includes procedures to control noise enhancement, increased noise was observed in the processed images. Improvements in noise control could lead to better specificity while increasing the

sensitivity of breast cancer detection. The results could also be improved by interpreting a combination of the original or digitized mammograms with their enhanced versions; increased familiarity with the enhanced mammograms may assist the radiologists in the detection of abnormalities. New laser printers (such as the Kodak 8610) can print images of the order of $4,096 \times 4,096$ pixels with 12-bit gray scale on film; this could lead to improved quality in the reproduction of the enhanced images and consequent improved interpretation by radiologists.

Digital image enhancement has the potential to improve the accuracy of breast cancer diagnosis and lead to earlier detection of breast cancer. Parallel computing strategies may assist in the practical application of the ANCE technique in a screening program [246, 1112, 1113, 1114].

12.11 Application: Classification of Breast Masses and Tumors via Shape Analysis

Based upon the differences in the shape characteristics of benign masses and malignant tumors as observed on mammograms, several methods have been proposed for their classification by using shape factors (see Chapter 6). Ackerman and Gose [615] analyzed breast lesions on xeroradiographs and investigated the use of four measures of malignancy: calcification, spiculation, roughness, and area-to-perimeter ratio. Their spiculation and roughness measures required the location of the center of the lesion as a reference point for computing radial projections. The center of the lesion was simply defined as the average position of the left-to-right and top-to-bottom borders of the rectangle bounding the lesion. Given a suspicious area on a xeroradiograph, their computer-aided classification methods obtained similar operational characteristic curves as that of the radiologist. In another study on xeroradiographs, Ackerman et al. [1084] used 36 radiographic properties of lesions to estimate the probability of malignancy. Using the properties in an automated clustering scheme, they achieved an FN rate of zero at an FP rate of 45%.

Pohlman et al. [1115] used measures of tumor circularity and surface roughness to classify breast tumors. By using a logistic regression model, they reported an area (A_z) of 0.9 under the ROC curve. The shape features were based on the radial distances of a mass boundary from its centroid. The surface roughness was calculated as the percentage of angles with multiple boundary points. In another study, Pohlman et al. [407] segmented lesions from their background using an adaptive region-growing technique and achieved a 97% detection rate with a set of 51 mammograms. They also used six morphological descriptors for benign-versus-malignant classification of the detected lesions, and achieved areas under the ROC curve ranging from 0.76 to 0.93.

Their detection method required manual selection of seed points for region growing and adequate segmentation was obtained over several trials.

Kilday et al. [1116] developed a set of seven shape features based on tumor circularity and radial distance measures (RDM) from the centroid to the points on the boundary. The features included compactness, mean of RDM, standard deviation of RDM, entropy of the RDM histogram, area ratio, zero crossings, and boundary roughness. A three-group classification of breast tumors as fibroadenoma, cyst, and cancer was performed by using the features in a linear discriminant function. They reported a classification accuracy of 51% using the leave-one-out method.

Bruce and Kallergi [1117] studied the effect of the resolution of the images on the detection and classification of mammographic mass shapes as round, lobular, or irregular, using the same shape features as proposed by Kilday et al. [1116] along with wavelet-based scalar energy features. Methods based upon Markov random fields were employed to extract mass regions. Features computed from the regions extracted in images at two different resolutions (220 μm and 180 μm) resulted in similar classification trends. The best overall classification rate of 75.9% was obtained by using wavelet-based features computed from manually segmented mass regions. Later on, Bruce and Adhami [1118] classified manually segmented mass shapes as round, nodular, or stellate using multiresolution shape features derived via the application of the discrete wavelet transform modulus-maxima method. They reported to have achieved at best 80% classification accuracy using the multiresolution features in linear discriminant analysis with boundaries of masses extracted from 60 digitized mammograms. Their morphological description as round or oval shapes for benign masses, and nodular or stellate shapes for malignant tumors may not hold good with all possible mass shapes. It is well known that some benign masses possess stellate shapes and a small proportion of malignant tumors are circumscribed [163, 345, 376]. The database used by Bruce and Adhami [1118] did not represent a good mixture of shapes of all types of masses.

Rangayyan et al. [163] used moments of distances of contour points from the centroid, compactness of the boundary, Fourier descriptors, and chord-length statistics to characterize the roughness of tumor boundaries. Whereas circumscribed-versus-spiculated classification of masses was achieved at accuracies of up to 94.4%, the benign-versus-malignant classification accuracy obtained by using only shape factors based upon contours was limited to about 76%, with a database of 54 masses (28 benign and 26 malignant). Using the same dataset, Menut et al. [354] achieved similar benign-versus-malignant classification accuracy (76%) by performing parabolic modeling of tumor boundaries and using the mean and variance values of the narrowness and width of the individual parabolic segments for classification. The results mentioned above emphasize the difficulties involved in the benign-versus-malignant classification of masses based only on morphological features.

Other methods developed to detect distortions in mammographic images as a result of the presence of masses have included steps to follow the morphological signs or orientations of masses during various stages of detection [339, 635, 1119].

In the work of Rangayyan et al. [166, 345], the shape parameters cf , f_{cc} , and SI (see Chapter 6) were applied to classify a set of contours of 28 benign breast masses and 26 malignant tumors. Figure 6.27 shows the 54 contours arranged in the order of increasing shape complexity as characterized by the feature vector (cf, f_{cc}, SI) ; Figure 6.28 shows a scatter plot of the three features. The features were used in the BMDP 7M stepwise discriminant analysis program [674] to perform pattern classification. The program realizes a jack-knife validation procedure using the leave-one-out algorithm. The classification performance of the features was validated using ROC methodology. ROC plots were obtained by using the BMDP software package and varying the cut points for benign and malignant prior probabilities between 0 and 1 in steps of 0.1. The procedure does not affect the discriminant ratings of the variables and influences only the computation of the constant term in the discriminant function, thus resulting in varying classification accuracies. The ROC curves for some of the feature combinations are shown in Figure 12.26. The area A_z under each ROC curve was computed using the trapezoidal rule.

Table 12.12 provides details on the benign-versus-malignant classification performance of various combinations of the three features (cf, f_{cc}, SI) . All of the features, individually and in different combinations, could effectively discriminate circumscribed benign masses from spiculated malignant tumors. The parameter cf , being a global measure of shape complexity, failed to classify almost all spiculated benign masses, and classified four out of seven circumscribed malignant tumors correctly.

Concavity analysis is sensitive to the presence of spicules in a mass boundary. Fractional concavity (f_{cc}) increases with the number of spicules and their length; however, it does not take into account the degree of spicularity of the individual spicules present in the boundary. Circumscribed malignant tumors typically contain a large number of microlobulations in their boundaries that could appear as alternating concave and convex segments; hence, f_{cc} could distinguish six out of the seven circumscribed malignant tumors correctly, and resulted in a high sensitivity of 88.5%. However, f_{cc} does not represent the characteristics of the spicules intricately in terms of their depth and narrowness; hence, it failed to distinguish spiculated benign masses from spiculated malignant tumors, and resulted in a poor specificity of 60.7% with $A_z = 0.75$.

Because the degree of spiculation of boundary segments is characterized by the spiculation index, a significant portion of the spiculated benign cases (five out of 12) with large convexities and a few narrow spicules were correctly classified as benign by SI . The boundary of one of such benign masses is shown in Figure 12.27 as an example. Both the parameters cf and f_{cc} computed for this boundary are high and misclassified the mass as malignant. However, a careful observation of the boundary reveals that although the mass is




FIGURE 12.26

ROC plots for SI , f_{cc} , (SI, f_{cc}) , and (SI, f_{cc}, cf) . SI: spiculation index. fcc: fractional concavity. C: modified compactness cf . The set of contours used is illustrated in Figure 6.27. Reproduced with permission from R.M. Rangayyan, N.R. Mudigonda, and J.E.L. Desautels, “Boundary modeling and shape analysis methods for classification of mammographic masses”, *Medical and Biological Engineering and Computing*, 38:487–496, 2000. © IFMBE.

TABLE 12.12

Numbers of Masses and Tumors Correctly Classified as Benign or Malignant by the Three Shape Factors cf , f_{cc} , and SI .

Features	Benign		Malignant		% Accuracy			
	Circ.	Spic.	Circ.	Spic.	Ben.	Mal.	Total	A_z
SI	16/16	5/12	3/7	19/19	75.0	84.6	79.6	0.82
f_{cc}	15/16	2/12	6/7	17/19	60.7	88.5	74.1	0.75
cf	15/16	1/12	4/7	19/19	57.1	88.5	72.2	0.76
cf, SI	16/16	5/12	3/7	19/19	75.0	84.6	79.6	0.80
f_{cc}, SI	16/16	4/12	3/7	19/19	71.4	84.6	77.8	0.80
cf, f_{cc}	15/16	1/12	5/7	19/19	57.1	92.3	74.1	0.72
cf, f_{cc}, SI	16/16	4/12	5/7	19/19	71.4	92.3	81.5	0.79

Circ.: circumscribed; Spic.: spiculated. Ben.: benign; Mal.: malignant. See Figure 6.27 for an illustration of the contours. Reproduced with permission from R.M. Rangayyan, N.R. Mudigonda, and J.E.L. Desautels, "Boundary modeling and shape analysis methods for classification of mammographic masses", *Medical and Biological Engineering and Computing*, 38:487–496, 2000. © IFMBE.

spiculated in nature, a major portion of its boundary does not possess sharp and narrow spicules. Hence, the parameter SI , which is sensitive to narrow spicules, correctly classified this mass as a benign mass.




FIGURE 12.27

Shape factors for the boundary of a spiculated benign mass: $cf = 0.8$, $f_{cc} = 0.53$, and $SI = 0.3$. Only SI correctly classified the mass as benign [166].

SI provided an improved specificity of 75% in classifying the benign masses in the database used. This is an encouraging result because none of the other shape parameters in earlier studies [163, 354] could be effective in separating the spiculated benign cases in the MIAS database. Even though SI failed to correctly classify four of the seven circumscribed malignant cases (although with narrow classification margins), it resulted in the best classification accuracy among all combinations of the three features with $A_z = 0.82$. The misclassified cases, although malignant, did not have prominent spicules in their boundaries. It may be observed from Table 12.12 that combining SI with the other features generally yielded improved classification results with high values of A_z ; the ROC curves for some of the feature combinations are shown in Figure 12.26. Because spicules are characterized and emphasized by the narrowness of their angles, SI is particularly sensitive to the stellate or star-like distortions in malignant tumors; the recognition of such distortion has been the focus of several studies on tumor detection [339, 635, 644].

A benign-versus-malignant classification accuracy of 82% was obtained, with $A_z = 0.79$, by combining f_{cc} and SI , which are sensitive to local variations in a boundary, with the global shape feature cf . A total of eight FPs out of the 28 benign masses and two FNs out of the 26 malignant tumors were observed with the combination of all of the three features.

SI , f_{cc} , and cf individually resulted in benign-versus-malignant classification accuracies of 80%, 74%, and 72%, respectively. Using the same dataset of 54 contours, Rangayyan et al. [163] and Menut et al. [354] reported benign-versus-malignant classification accuracies of no more than 76% using various combinations of several other shape factors.

In the linear discriminant analysis model, the criterion to optimize the feature weights is based on maximizing the variance between the classes while minimizing the variance within each class. Also, the size of the class influences the computation of the feature weights. Although the dataset in the study being described here is evenly divided between benign masses (28) and malignant tumors (26), it includes an unusual proportion of spiculated benign masses from the MIAS database (12 out of 28 benign masses in a total of 54 cases). Considering the above, the performance of the features as described above could be regarded to be good. The addition of other features based upon density variations and textural information [163, 165, 275, 676] may result in improved benign-versus-malignant discrimination; see Section 12.12.

12.12 Application: Content-based Retrieval and Analysis of Breast Masses

Alto et al. [528, 529, 1120, 1121] applied combinations of the three shape factors cf , f_{cc} , and SI ; 14 texture measures as defined by Haralick [441, 442] (see Section 7.3); and four measures of edge sharpness as defined by Mudigonda et al. [165] (see Section 7.9.2) for content-based image retrieval (CBIR) and pattern classification studies with a set of 57 ROIs of breast masses and tumors. The cases were selected from Screen Test: Alberta Program for the Early Detection of Breast Cancer [61], and include 20 cases (22 breasts affected) exhibiting a total of 28 masses visible as 57 ROIs on 45 mammograms. Twenty of the ROIs correspond to biopsy-proven malignant tumors, and the remaining 37 to biopsy-proven benign masses. The film mammograms were digitized using the Lumiscan 85 scanner at a resolution of $50 \mu\text{m}$ with 12 bits per pixel. The 57 ROIs are shown in Figure 12.4 arranged in the order of decreasing acutance. Figure 12.5 shows the contours of the 57 masses arranged in increasing order of f_{cc} . Figures 12.5 and 12.4 demonstrate that a few of the benign masses and malignant tumors may appear to be out of place if their classification is based only on the features used in the two illustrations of rank-ordering.

Figure 12.28 illustrates the region corresponding to a macrolobulated benign mass, the contour of the mass with its concave and convex parts identified for the computation of f_{cc} , the ribbon of pixels extracted to compute the texture features, and the normals to the contour for the computation of

edge-sharpness measures. A single combined GCM was created from the pixel values extracted in the four directions with respect to the pixel under consideration. Fourteen texture features were computed according to Haralick's definitions [441, 442] for the mass ribbons. In addition to the features for the original ribbons at a resolution of $50 \mu m$, the texture features were computed using Gaussian-smoothed and down-sampled versions of the ribbons at pixel resolution of $200 \mu m$ and $800 \mu m$.

12.12.1 Pattern classification of masses

Pattern classification experiments were conducted using linear discriminant analysis with the 14 texture features at pixel resolution of 50, 200, and $800 \mu m$. The results indicated $200 \mu m$ to be the most suitable resolution for discrimination between benign and malignant masses. Stepwise logistic regression was performed using the SPSS software package [1092, 1093] to select a subset of features from the three separate sets of shape, edge-sharpness, and texture features. As a result of this evaluation, the shape factor of fractional concavity f_{cc} , the edge-sharpness feature of acutance A as defined in Equation 2.110, and the texture feature of sum entropy F_8 were selected. (Chan et al. [450] found that the three texture features of correlation, difference entropy, and entropy performed better in the classification of breast masses than other combinations of one to eight texture features selected in a specific sequence.) Although the shape features, in particular f_{cc} , gave high sensitivity and specificity, they are highly dependent on the accuracy of the contour, which is not easily drawn even by an experienced radiologist. The results of automatic segmentation methods still need to be confirmed by an expert radiologist, and are often subject to errors and artifacts. It should be observed that large populations of pixels around the given contour are used in the procedures to compute the texture and edge-sharpness measures. The definitions of the ribbon for texture measures and the set of normals to the contour for edge-sharpness measures make the parameters less sensitive to inaccuracies in the contour than the shape factors. These observations lend support to the argument in favor of combining features representing multiple characteristics.

A scatter plot of the three features $[f_{cc}, A, F_8]$ of the 57 masses is given in Figure 12.6. It is seen that while f_{cc} separates the benign and malignant categories well, the texture feature F_8 does not possess good discriminant capability. The measure of acutance A indicates an intermediate degree of discriminant capability. A scatter plot of the three shape factors $[f_{cc}, cf, SI]$ of the 57 masses is given in Figure 12.7. Each of the three shape factors demonstrates high discriminant capability.

The benign-versus-malignant discriminatory performance of the features was validated using several approaches, including linear discriminant analysis, logistic regression, Mahalanobis distance, k-NN, and the ROC methodology. In the pattern classification experiments conducted with the Mahalanobis distance, each mass was treated in turn as the sample to be classified. Mean vec-




FIGURE 12.28

(a) ROI of the benign mass b164ro94 (see Figure 12.4.) (b) ROI overlaid with the contour demonstrating concave parts in black and convex parts in white. (c) Ribbon of pixels for the purpose of computing texture measures, derived by dilating and eroding the contour in (b). (d) Normals to the contour, shown at every tenth point on the contour, used for the computation of edge-sharpness measures. See also Figures 7.24 and 7.25. Reproduced with permission from H. Alto, R.M. Rangayyan, and J.E.L. Desautels, "Content-based retrieval and analysis of mammographic masses", *Journal of Electronic Imaging*, in press, 2005. © SPIE and IS&T.

tors and pooled covariance matrices were computed using the feature vectors of the remaining benign and malignant samples. The Mahalanobis distance was computed from the sample on hand to the mean vectors of the benign and malignant classes; the sample was assigned to the class with the smaller distance.

The sensitivity and specificity values for the ROC plots were obtained with the BMDP 7M stepwise discriminant analysis program [674] by varying the cut points for benign and malignant prior probabilities between 0 and 1 in steps of 0.1. The program realizes a jack-knife validation procedure using the leave-one-out algorithm. Figure 12.29 shows the ROC curves for f_{cc} , A , and F_8 individually and combined. The area A_z under each ROC curve was computed using the trapezoidal rule. The results of all of the experiments mentioned above are listed in Table 12.13. The sensitivity and specificity values obtained at a prior probability value of 0.5 for both the benign and malignant groups are also shown in Table 12.13 for the sake of illustration. The shape factor f_{cc} has demonstrated consistently high classification accuracy regardless of the pattern classification method. The classification performance of the texture features is poor. The measure of acutance has provided slightly better accuracy than the texture measures. The addition of texture and edge-sharpness measures did not significantly alter the performance of f_{cc} . (See Sections 12.3.1 and 12.4.2 for examples of application of other pattern classification methods to the same dataset.)

12.12.2 Content-based retrieval

Systems for CBIR from multimedia databases offer advantages over traditional archiving systems such as single-media, text-based, relational, and hierarchical databases [1122, 1123, 1124, 1125, 1126, 1127, 1128, 1129, 1130, 1131]. The retrieval of relevant multimedia information (such as images, video clips, attributes, and numerical data) from large databases should be accomplished effectively and efficiently in order to assist both the novice and the expert user. Content-based retrieval methods use image content, in the form of feature values, image attributes, and other image descriptors, to identify relevant images or image-related information in response to a query. A query may be an example image, a hand-drawn outline of a shape, a natural language query, or a selection from a set of possible categories provided by the system's user interface. A number of different retrieval methods have been proposed in the literature: some have utilized shape factors whereas others have used other image attributes or textual annotation.




FIGURE 12.29

ROC plots for $[f_{cc}, A, F_8]$ individually and combined. The ROC curves for f_{cc} and $[f_{cc}, A, F_8]$ overlap completely. Reproduced with permission from H. Alto, R.M. Rangayyan, and J.E.L. Desautels, "Content-based retrieval and analysis of mammographic masses", *Journal of Electronic Imaging*, in press, 2005. © SPIE and IS&T.

TABLE 12.13

Accuracy of Classification of Masses as Benign or Malignant Using Combinations of the Shape Factor f_{cc} , Acutance A , and the Texture Measure Sum Entropy F_8 with Pattern Classification Methods and CBIR (Precision) [529].

Features	Logistic regression			Mahalanobis distance			Linear discriminant analysis				k-NN		Precision	
	Sens	Spec	Avg	Sens	Spec	Avg	Sens	Spec	Avg	A_z	$k = 5$	$k = 7$	$k = 5$	$k = 7$
f_{cc}	90	97.3	94.7	90	97.3	94.7	100.0	97.3	98.2	0.99	94.7	94.7	95.1	95.2
A	50	94.6	78.9	75	67.6	70.0	75.0	73.0	73.7	0.74	68.4	73.7	65.3	67.4
F_8	30	86.5	66.7	65	56.8	59.6	75.0	54.1	61.4	0.68	63.2	54.4	58.2	60.9
f_{cc}, A	90	97.3	94.7	90	97.3	94.7	95.0	97.3	96.5	0.98	96.5	94.7	93.0	91.2
f_{cc}, F_8	90	97.3	94.7	90	97.3	94.7	100.0	97.3	98.2	0.99	94.7	94.7	95.1	93.7
A, F_8	55	86.5	75.4	60	70.3	66.7	75.0	73.0	73.7	0.76	75.4	75.4	68.4	68.4
f_{cc}, F_8, A	90	97.3	94.7	95	97.3	96.5	100.0	97.3	98.2	0.99	96.5	96.5	90.9	91.2
14 texture	*	*	*	70	50.0	64.9	65.0	64.9	64.9	0.67	#	#	#	#

*: Logistic regression identified the texture feature F_8 as the only significant feature; results were computed for F_8 only.

#: Experiments not conducted for this feature set. Sens = sensitivity; Spec = specificity; Avg = average accuracy as percentages. See Figures 12.4 and 12.5 for illustrations of the masses and their contours.

A review of CBIR systems by Gudivada and Raghavan [1123] outlines previous approaches to content-based retrieval, and expresses the need to utilize features from a variety of approaches based on attributes, feature extraction, or object recognition, for information representation. Gudivada and Raghavan [1123] and Yoshitaka and Ichikawa [1122] indicate that conventional database systems are not well suited to handle multimedia data containing images, video, and text. Hence, it is necessary to explore more flexible query and retrieval methods.

Representation of breast mass images for CBIR: The first step in the development of a CBIR system is to represent the data or information in a meaningful way in the database so that retrieval is facilitated for a given application [529, 1085, 1121]. The representation of breast masses and tumors in a database requires the design of a reasonable number of descriptors to represent the image features of interest (or diagnostic value) with minimal loss of information. It is well established that most benign masses have contours that are well-circumscribed, smooth, and are round or oval, and have a relatively homogeneous internal texture. On the other hand, malignant tumors typically exhibit ill-differentiated and rough or spiculated contours, with a heterogeneous internal texture [54, 55]. For these reasons, shape factors and texture measures have been proposed for differentiating between benign masses and malignant tumors [163, 165, 275, 345, 354, 428, 451].

Various researchers have chosen to represent the contours of objects in a variety of ways, some of which include: coding the object's contour as an ordered sequence of points or high-curvature points [1124, 1125, 1126]; using chain-code histograms [1125, 1126, 1127, 1128]; and using shape descriptors such as compactness [163, 274, 345, 428, 1118, 1132], concavity/convexity [345, 354, 1132], moments [163, 274, 1128, 1132], Fourier descriptors [274, 428, 1124, 1126], spiculation index [345], and the wavelet transform modulus-maxima [1118]. Loncaric [406] gives an overview of shape analysis techniques from chain codes to fractal geometry. (See Chapter 6 for details on shape analysis.)

Automatically extracted shapes and parameters are considered to be primitive features, whereas logical features are abstract representations of images at various levels of detail and may be synthesized from primitive features. Logical features require more human intervention and domain expertise, and therefore, there is a higher cost associated with preprocessing the data for the database. CBIR approaches differ with respect to the image features that are extracted, the level of abstraction of the features, and the degree of desired domain independence [1123]. Depending upon the application, object-based descriptors such as tumor shape may be preferred to attribute-based descriptors such as color or texture; keywords may also be used where appropriate [1129, 1130]. In the work of Alto et al. [529], the features used are related to radiologically established attributes of breast masses. When the images in a database are indexed with objective measures of diagnostic features, the database may be referred to as an *indexed atlas* [1085, 1133].

The query process: Once the mammographic masses have an appropriate representation in a database, the next step in the development of a CBIR system is to design the query techniques to fit the needs of the end-user. In a CAD application, the end-user could be a radiologist, a radiology intern, or a physician. In standard text-based databases, queries are generally comprised of keywords, natural language queries, or browsing procedures (that is, query by subject). For image or multimedia databases, the same methods may apply only if there are searchable keywords or textual descriptors associated with the images. In a “query by example”, the user specifies a condition by giving examples of the desired image or object, either by cutting and pasting an image or by sketching the example object’s contour [1122]. One of the best-known commercial CBIR systems is Query by Image Content (QBIC) developed at IBM [1129]. QBIC uses visual content such as color percentages, color layout, and texture extracted from images of art collections. A CBIR system developed by Srihari [1130], known as Piction, contains images of newspaper photos annotated with their associated captions. Queries based on text and image features extracted from the photos may then be used to identify human faces found in the newspaper photographs.

A comprehensive list of query classes is given by Gudivada and Raghavan [1123] as: color, texture, sketch, shape, volume, spatial constraints, browsing, objective attributes, subjective attributes, motion, text, and domain concepts. Fewer classes may be used when the database is highly domain-specific, and more are needed when the database is of general scope.

When a query is made in a CBIR system, the retrieved results are typically presented as a set or series of images that are rank-ordered by their degree of similarity (or a distance measure, such as the Euclidean, Manhattan, or Mahalanobis distance, as an indicator of dissimilarity) with respect to the query image. This is different from retrieval in text-based database systems, which generally provide results with an exact match (that is, a single word, set of words, or a phrase). The CBIR work of Alto et al. was focused on retrieving similar masses, of established diagnosis, to assist the radiologist by suggesting a probable diagnosis for the query case on hand. The concept of similarity is especially pertinent in such an application because no two breast masses may be expected to be identical, and a perfect or exact match to a query would be improbable in practice.

Some of the shape-matching procedures suggested by Trimeche et al. [1125] require a comparison of the vertices of polygonal models of the query contour and the database contours. Each vertex is represented by a set of values (such as scale, angle, ratio of consecutive segments, and the ratio to the overall length). The feature vectors could be excessively lengthy if the contour has many vertices, such as a spiculated mass. A matrix containing all possible matches between the vertices of the query shape and each of the candidate contours may then be created. The polygonal model method produced good results with fish contours. The use of shape factors to represent the contours

of masses, as in the work of Alto et al., simplifies the process of comparative analysis.

Visualization of the query results is accomplished by rank-ordering the retrieved results from the minimum to the maximum distance and presenting the top k objects to the user, where k could be 3, 5, 7, \dots , N . One suggested method of visualization of the retrieved results that may enhance the user's perception of the overall information presented was defined by Moghaddam et al. [1134] as a Splat: the retrieved images were displayed in rank order of their visual similarities, with their placement on the page with respect to the query being dictated by their mutual similarities.

Evaluation of retrieval: An important step in developing a CBIR system is to evaluate its efficiency with respect to the retrieval of relevant information. Measures of precision and recall have been proposed to assess the performance of general information retrieval systems, based upon the following definitions [1131]:

- Correct detections

$$A_k = \sum_{n=1}^k V_n, \quad (12.93)$$

where k is the number of retrieved objects, and $V_n \in \{0, 1\}$ with $V_n = 1$ if the retrieved object is relevant to the query and $V_n = 0$ if it is irrelevant. In the present application, a relevant object is a retrieved benign mass for a benign query sample; a retrieved malignant tumor would be considered irrelevant. In the case of a malignant query sample, a retrieved benign mass would be irrelevant and a malignant tumor would be relevant.

- False alarms

$$B_k = \sum_{n=1}^k (1 - V_n). \quad (12.94)$$

- Misses

$$M_k = \left(\sum_{n=1}^N V_n \right) - A_k, \quad (12.95)$$

where N is the total number of objects in the database.

- Correct dismissals

$$D_k = \left(\sum_{n=1}^N (1 - V_n) \right) - B_k. \quad (12.96)$$

- Recall, defined as the ratio of the number of relevant retrieved objects to all relevant objects in the database, and computed as

$$R_k = \frac{A_k}{A_k + M_k}. \quad (12.97)$$

- Precision, defined as the ratio of the number of relevant retrieved objects to all retrieved objects, and computed as

$$P_k = \frac{A_k}{A_k + B_k}. \quad (12.98)$$

- Fallout, defined as the ratio of the number of retrieved irrelevant objects to all irrelevant objects in the database, and computed as

$$F_k = \frac{B_k}{B_k + D_k}. \quad (12.99)$$

The following plots may be used to evaluate the effectiveness of CBIR systems:

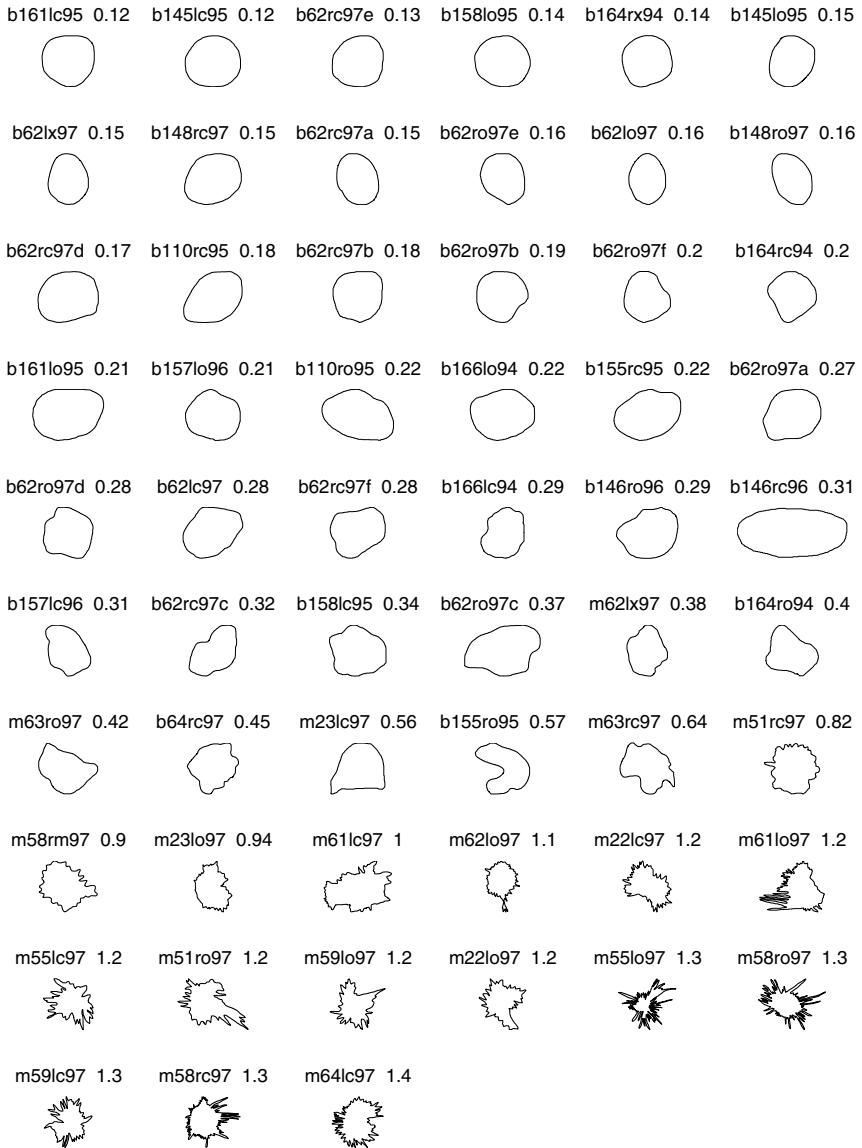
- Retrieval effectiveness: precision versus recall.
- ROC: correct detections versus false alarms.
- Relative operating characteristics: correct detections versus fallout.
- Response ratio: $\frac{A_k}{B_k}$ versus A_k .

In general, an effective CBIR system will demonstrate high precision for all values of recall [1131].

Results of CBIR with breast masses: Alto et al. [529] applied a content-based retrieval algorithm to the 57 masses and their contours shown in Figures 12.4 and 12.5. The retrieval algorithm uses the Euclidean distance between a query sample's feature vector and the feature vector of each of the remaining masses in the database, and rank-orders the masses corresponding to the vectors that are most similar to the query vector (that is, the shortest Euclidean distance). The rank-ordered masses are presented to the user, annotated with the biopsy-proven diagnosis for each retrieved mass. The 57 masses were each used, in turn, as the query sample.

Figure 12.30 shows the contours of the 57 masses rank-ordered by the Euclidean distance from the origin in the three-feature space of $[f_{cc}, cf, SI]$. This is equivalent to sorting the contours by the magnitudes of the feature vectors $[f_{cc}, cf, SI]$. Observe that the use of three shape factors has led to a more comprehensive characterization of shape roughness than only one (f_{cc}) as in Figure 12.5, resulting in a different order of sorting.

Figures 12.31 through 12.34 show the retrieval results for the four masses illustrated in Figure 12.2 using various feature vectors including f_{cc} , A , and F_8 . (In each case, the first mass at the left is the query sample. The retrieved samples are arranged in the increasing order of distance from the query sample from left to right. The contour of the mass in each ROI is provided above the ROI.) The results indicate that the masses retrieved and their sequence

**FIGURE 12.30**

Contours of 57 breast masses, including 37 benign masses and 20 malignant tumors. The contours are arranged in the order of increasing magnitude of the feature vector $[f_{cc}, cf, SI]$, which is given next to each sample. Note that the masses and their contours are of widely differing size, but have been scaled to the same size in the illustration. For details regarding the case identifiers, see Figure 12.2. See also Figure 12.5. Figure courtesy of H. Alto [528].

depend upon the features used to characterize and index the masses. Regardless, all of the results illustrated clearly lead to decisions that agree with the known diagnoses of the query samples, except for the case in Figure 12.33 (b).

Although the texture and edge-sharpness measures resulted in poor classification accuracies on their own, the results of retrieval using both of the features with the shape factor f_{cc} indicate the need to include these measures so as to provide a broader scope of representation of radiographic features than shape complexity alone.

The results of content-based retrieval, as illustrated in Figures 12.31 – 12.34, lend easily to pattern classification with the k-NN method. The k-NN method may be applied by simple visual inspection of the first k cases in the results of retrieval; the classification of the query sample is made based upon the known classification of the majority of the first k objects. Alto et al. applied the k-NN method to the retrieval results with $k = 5, 7, 9,$ and 11 . Correct detections in these cases refer to the retrieval of at least 3, 4, 5, or 6, respectively, correct cases by virtue of their diagnosis corresponding to the known diagnosis of the query (test) sample. The results of k-NN analysis and retrieval precision are presented in Table 12.13 for $k = 5$ and 7 . The high levels of classification accuracy and retrieval precision with the use of shape factors indicate the importance of shape in the analysis of breast masses and tumors. A study by Sahiner et al. [428] has also indicated the importance of shape parameters in the classification of breast masses and tumors.

See Zheng et al. [1135] for the application of CBIR to image analysis in pathology (histology).

12.12.3 Extension to telemedicine

The concepts of an indexed atlas and CBIR may be combined with mobile software agents for web-based medical image retrieval and analysis of medical images, telemedicine, and remote medical consultation applications [1085]. Software agents are autonomous, intelligent, software objects that can process, analyze, and make decisions about data [1136, 1137]. A mobile agent is a self-contained software program that can move within a computer network and perform tasks for the user. A mobile agent can reduce search time and function with limited computational resources and low-bandwidth communication links: this is accomplished by having the agent process or evaluate the data at the source, and then transmit only the pertinent data to the user.

Mobile agents can serve a variety of functions, and may be used to [1136, 1137, 1138, 1139, 1140]:

- search information residing at remote nodes and report back to the source (information agent);
- find under-utilized network resources to perform computationally intensive processing tasks (computation agent); and

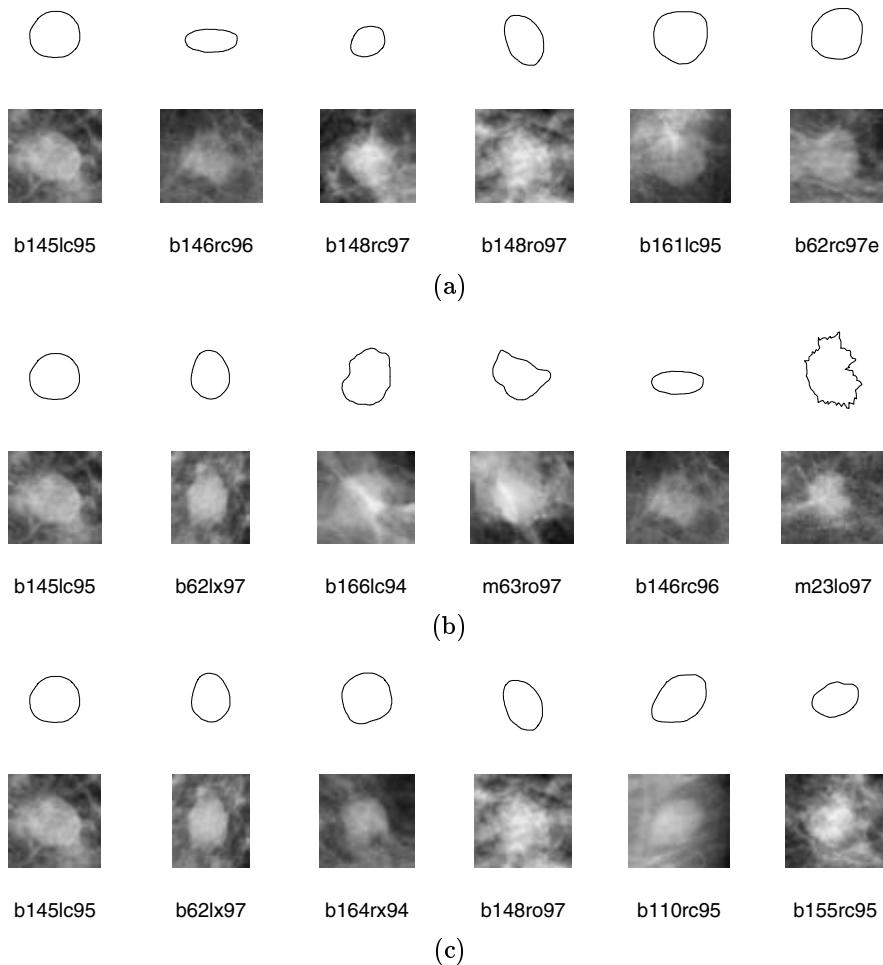


FIGURE 12.31

Content-based retrieval with the circumscribed benign query sample b145lc95 (a) using the shape factor f_{cc} only, (b) using acutance A only, and (c) using the three features $[f_{cc}, A, F_8]$. For details of case identification, see Figure 12.2. Reproduced with permission from H. Alto, R.M. Rangayyan, and J.E.L. Desautels, “Content-based retrieval and analysis of mammographic masses”, *Journal of Electronic Imaging*, in press, 2005. © SPIE and IS&T.

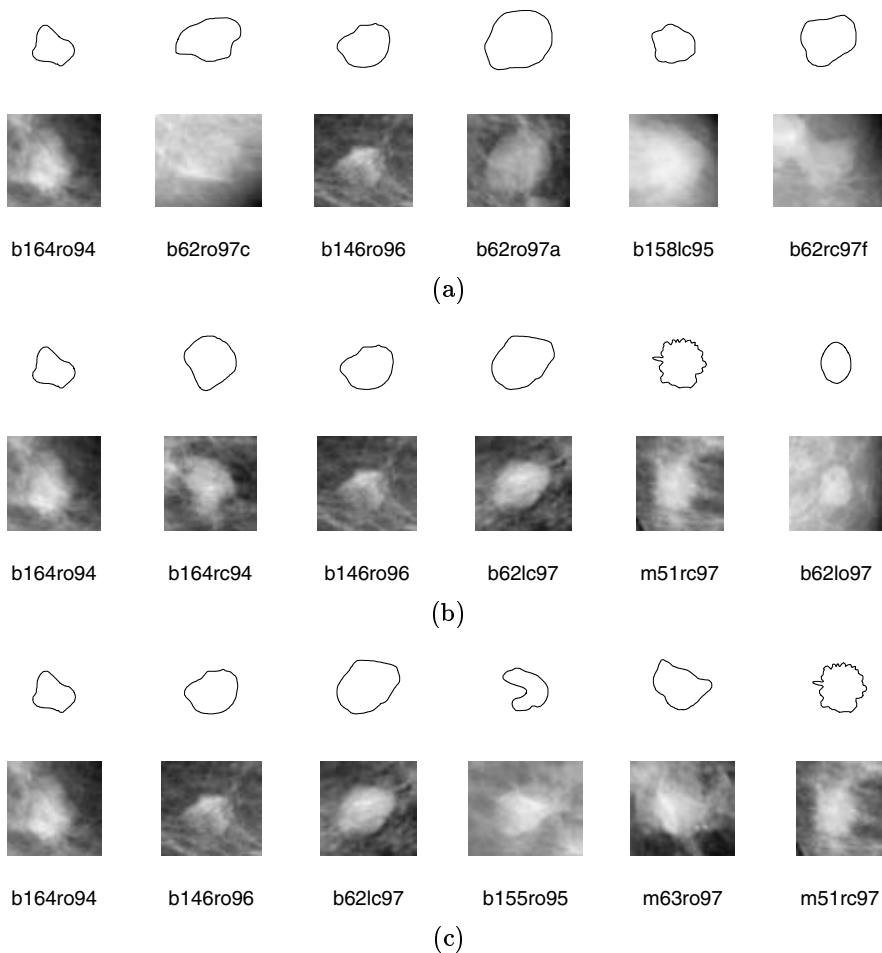


FIGURE 12.32

Content-based retrieval with the macrolobulated benign query sample b164ro94 (a) using the shape factor f_{cc} only, (b) using acutance A only, and (c) using the three features $[f_{cc}, A, F_8]$. For details of case identification, see Figure 12.2. Reproduced with permission from H. Alto, R.M. Rangayyan, and J.E.L. Desautels, “Content-based retrieval and analysis of mammographic masses”, *Journal of Electronic Imaging*, in press, 2005. © SPIE and IS&T.

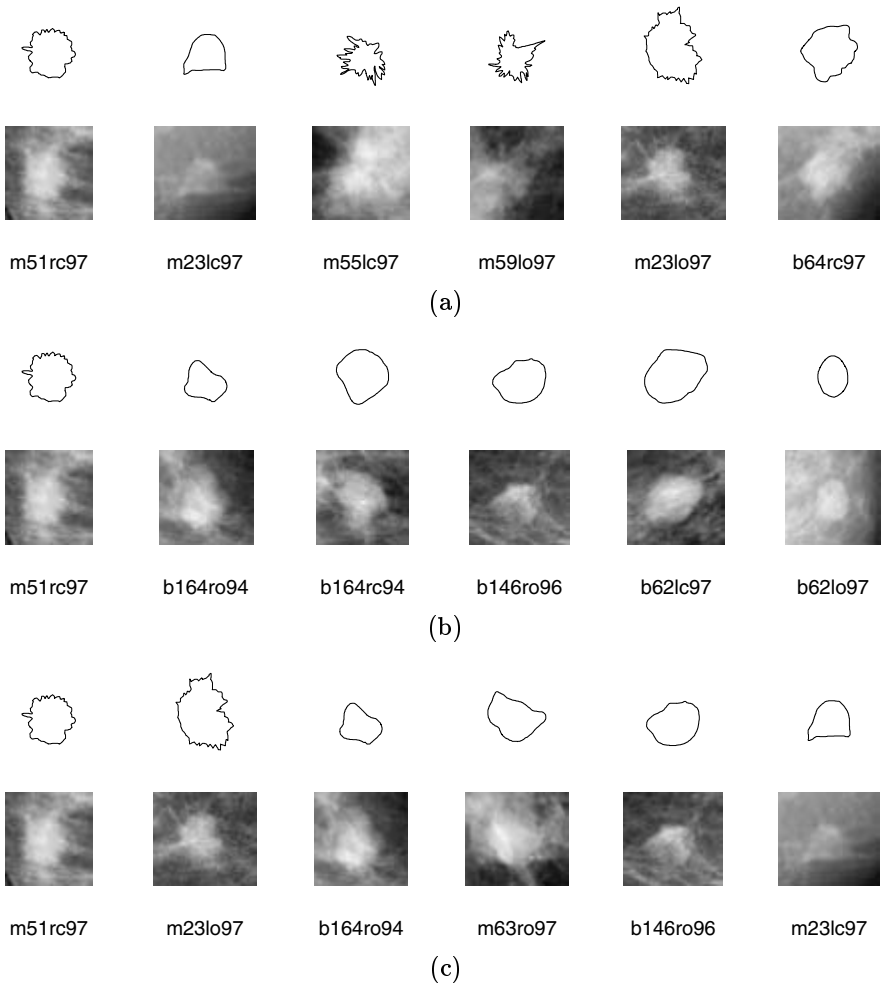


FIGURE 12.33

Content-based retrieval with the microlobulated malignant query sample m51rc97 (a) using the shape factor f_{cc} only, (b) using acutance A only, and (c) using the three features $[f_{cc}, A, F_8]$. For details of case identification, see Figure 12.2. Reproduced with permission from H. Alto, R.M. Rangayyan, and J.E.L. Desautels, “Content-based retrieval and analysis of mammographic masses”, *Journal of Electronic Imaging*, in press, 2005. © SPIE and IS&T.

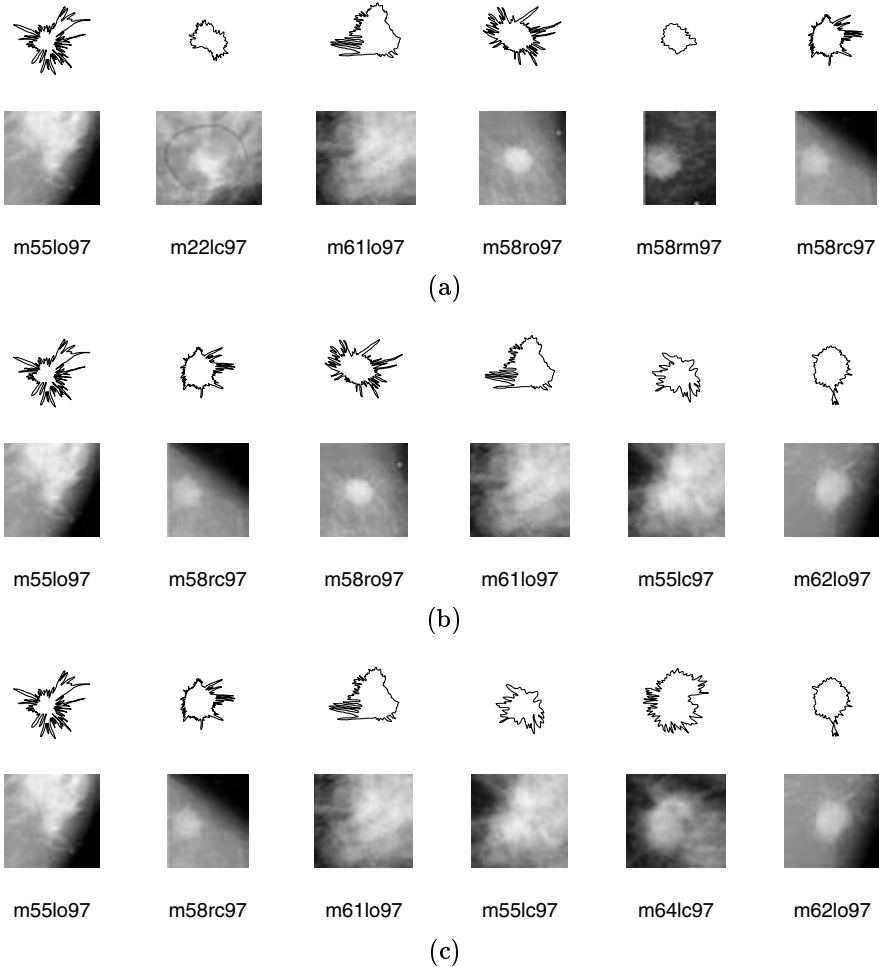


FIGURE 12.34

Content-based retrieval with the spiculated malignant query sample m55lo97 (a) using the shape factor f_{cc} only, (b) using acutance A only, and (c) using the three features $[f_{cc}, A, F_8]$. For details of case identification, see Figure 12.2. Reproduced with permission from H. Alto, R.M. Rangayyan, and J.E.L. Desautels, “Content-based retrieval and analysis of mammographic masses”, *Journal of Electronic Imaging*, in press, 2005. © SPIE and IS&T.

- send messages back and forth between clients residing at various network nodes (communication agent).

Figure 12.35 shows a schematic representation of the combined use of an indexed atlas, CBIR, and mobile agents in the context of mammography and CAD of breast cancer.

The major strength of mobile agents is that they permit program execution near or at a distributed data source by moving to each site in order to perform a computational task. In addition, mobile-agent systems typically have the characteristics of low network traffic, load balancing, fault tolerance, and asynchronous interaction. Agents can function independently of one another, as well as cooperate to solve problems. The use of mobile agents with a CBIR system brings specific benefits and difficulties. For example, mobile agents can move to sites with better or more data, and faster computers. They can replicate themselves and use the inherent power of parallelism to improve productivity. The basic strengths of mobile-agent systems include the inherent parallelism of multiple agents conducting simultaneous searches, parallel searching with intelligent preprocessing, and agent-to-agent communication. Specific difficulties with the mobile-agent paradigm include issues of security, complexity, and control. Security is important when dealing with patient information. Data encryption and restricting the agents to operations within secure networks could address security concerns.

12.13 Remarks

We have studied how biomedical images may be processed and analyzed to extract quantitative features that may be used to classify the images as well as lead toward diagnostic decisions. The practical development and application of such techniques is usually hampered by a number of limitations related to the extent of discriminant information present in the images selected for analysis, as well as the limitations of the features designed and computed. Artifacts inherent in the images or caused by the image acquisition systems impose further limitations.

The subject of pattern classification is a vast area by itself [402, 1086, 401, 1087]. The topics presented in this chapter provide a brief introduction to the subject.

A pattern classification system that is designed with limited data and information about the chosen images and features will provide results that should be interpreted with due care. Above all, it should be borne in mind that the final diagnostic decision requires far more information than that provided by images and image analysis: this aspect is best left to the physician or health-care specialist in the spirit of computer-*aided* diagnosis.

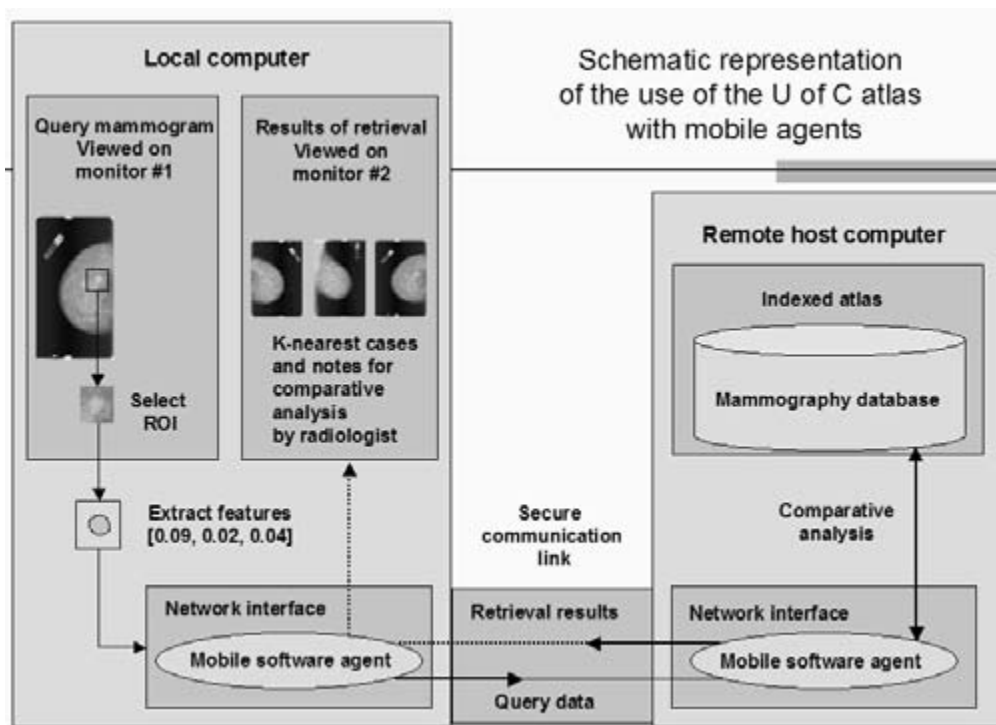


FIGURE 12.35

Use of an indexed atlas, CBIR, and mobile agents in the context of mammography and CAD of breast cancer. *Note:* “U of C” = University of Calgary, Calgary, Alberta, Canada. Reproduced with permission from H. Alto, R.M. Rangayyan, R.B. Paranjape, J.E.L. Desautels, and H. Bryant, “An indexed atlas of digital mammograms for computer-aided diagnosis of breast cancer”, *Annales des Télécommunications*, 58(5): 820 – 835, 2003. © GET – Lavoisier.

12.14 Study Questions and Problems

Selected data files related to some of the problems and exercises are available at the site

www.enel.ucalgary.ca/People/Ranga/enel697

1. The prototype vectors of two classes of images are specified as Class 1 : $[1, 0.5]^T$, and Class 2 : $[3, 3]^T$. A new sample vector is given as $[2, 1]^T$. Give the equations for two measures of similarity or dissimilarity, compute the measures for the sample vector, and classify the sample into Class 1 or Class 2 using each measure.
2. In a three-class pattern classification problem, the three decision boundaries are $d_1(\mathbf{x}) = -x_1 + x_2$, $d_2(\mathbf{x}) = x_1 + x_2 - 5$, and $d_3(\mathbf{x}) = -x_2 + 1$. Draw the decision boundaries on a sheet of graph paper. Classify the sample pattern vector $\mathbf{x} = [6, 5]^T$ using the decision functions.
3. Two pattern class prototype vectors are given to you as $\mathbf{z}_1 = [3, 4]^T$ and $\mathbf{z}_2 = [10, 2]^T$. Classify the sample pattern vector $\mathbf{x} = [4, 5]^T$ using (a) the normalized dot product, and (b) the Euclidean distance.
4. A researcher makes two measurements per sample on a set of 10 normal and 10 abnormal samples.

The set of feature vectors for the normal samples is

$$\{[2, 6]^T, [22, 20]^T, [10, 14]^T, [10, 10]^T, [24, 24]^T, [8, 10]^T, [8, 8]^T, [6, 10]^T, [8, 12]^T, [6, 12]^T\}.$$

The set of feature vectors for the abnormal samples is

$$\{[4, 10]^T, [24, 16]^T, [16, 18]^T, [18, 20]^T, [14, 20]^T, [20, 22]^T, [18, 16]^T, [20, 20]^T, [18, 18]^T, [20, 18]^T\}.$$

Plot the scatter diagram of the samples in both classes in the feature-vector space (on a sheet of graph paper). Design a linear decision function to classify the samples with the lowest possible error of misclassification. Write the decision function as a mathematical rule and draw the same on the scatter diagram.

How many (if any) samples are misclassified by your decision function? Mark the misclassified samples on the plot.

Two new observation sample vectors are provided to you as $\mathbf{x}_1 = [12, 15]^T$ and $\mathbf{x}_2 = [14, 15]^T$. Classify the samples using your decision rule.

Now, classify the samples \mathbf{x}_1 and \mathbf{x}_2 using the k -nearest-neighbor method, with $k = 7$. Measure distances graphically on your graph paper plot and mark the neighbors used in this decision process for each sample.

Comment upon the results — whether the two methods resulted in the same classification result or not — and provide reasons.

12.15 Laboratory Exercises and Projects

1. The file `tumor_shape1.dat` gives the values of several shape factors for 28 benign masses and 26 malignant tumors; the contours are illustrated in Figure 6.27. See Chapter 6 for details regarding the methods; see the file `tumor_shape1.txt` for details regarding the data file. Select the three shape factors (SI, f_{cc}, cf) and form feature vectors for each case. Using each feature vector as the test case and the remaining vectors in the dataset as the training set, classify each case as benign or malignant using (a) the Euclidean distance, (b) the Manhattan distance, and (b) the Mahalanobis distance. Compare the results with the classification provided in the data file and determine the TPF, TNF, FPF, and FNF.

Comment upon the performance of the three distance measures.

Repeat experiment with the data in the file `tumor_shape2.dat` for 37 benign masses and 20 malignant tumors; the contours are illustrated in Figure 12.5. See the file `tumor_shape2.txt` for details regarding the data file. Discuss the differences in the results you obtain with the two datasets.

2. Using the data in the preceding problem, classify each case as benign or malignant using the k -NN method, with $k = 3, 5,$ and 7 . Use the Euclidean distance. Comment upon the results.
3. The files `mfc_ben.dat` and `mfc_mal.dat` give the values of the three shape factors (mf, ff, cf) for 64 benign calcifications and 79 malignant calcifications, respectively. (See Chapter 6 for details.) Design a pattern classification system using the Mahalanobis distance and evaluate its performance in terms of the TPF, TNF, FPF, and FNF.

References

- [1] Lathi BP. *Signal Processing and Linear Systems*. Berkeley-Cambridge, Carmichael, CA, 1998.
- [2] Oppenheim AV, Willsky AS, and Nawab SH. *Signals and Systems*. Prentice Hall, Englewood Cliffs, NJ, 2nd edition, 1997.
- [3] Barrett HH and Swindell W. *Radiological Imaging – Volumes 1 and 2*. Academic, New York, NY, 1981.
- [4] Cho ZH, Jones JP, and Singh M. *Foundations of Medical Imaging*. Wiley, New York, NY, 1993.
- [5] Macovski A. *Medical Imaging Systems*. Prentice Hall, Englewood Cliffs, NJ, 1983.
- [6] Huda W and Slone R. *Review of Radiologic Physics*. Williams and Wilkins, Baltimore, MD, 1995.
- [7] Oppenheim AV and Schaffer RW. *Discrete-time Signal Processing*. Prentice Hall, Englewood Cliffs, NJ, 1989.
- [8] Gonzalez RC and Woods RE. *Digital Image Processing*. Prentice Hall, Upper Saddle River, NJ, 2nd edition, 2002.
- [9] Hall EL. *Computer Image Processing and Recognition*. Academic, New York, NY, 1979.
- [10] Pratt WK. *Digital Picture Processing*. Wiley, New York, NY, 2nd edition, 1991.
- [11] Rosenfeld A and Kak AC. *Digital Picture Processing*. Academic, New York, NY, 2nd edition, 1982.
- [12] Jain AK. *Fundamentals of Digital Image Processing*. Prentice Hall, Englewood Cliffs, NJ, 2nd edition, 1989.
- [13] Sonka M, Hlavac V, and Boyle R. *Image Processing, Analysis and Machine Vision*. Chapman & Hall Computing, London, UK, 1993.
- [14] Robb RA, editor. *Three-Dimensional Biomedical Imaging, Volumes I and II*. CRC Press, Boca Raton, FL, 1985.
- [15] Cooper KE, Cranston WI, and Snell ES. Temperature regulation during fever in man. *Clinical Science*, 27(3):345–356, 1964.
- [16] Cooper KE. Body temperature and its regulation. In *Encyclopedia of Human Biology*, volume 2, pages 73–83. Academic, New York, NY, 1997.
- [17] Sickles EA. Breast thermography. In Feig SA and McLelland R, editors, *Breast Carcinoma: Current Diagnosis and Treatment*, pages 227–231. Masson, New York, NY, 1983.

- [18] Zhou X and Gordon R. Detection of early breast cancer: An overview and future prospects. *Critical Reviews in Biomedical Engineering*, 17(3):203–255, 1989.
- [19] Keyserlingk JR, Ahlgren P, Yu E, Belliveau N, and Yassa M. Functional infrared imaging of the breast. *IEEE Engineering in Medicine and Biology Magazine*, 19(3):30–41, 2000.
- [20] Ohashi Y and Uchida I. Applying dynamic thermography in the diagnosis of breast cancer. *IEEE Engineering in Medicine and Biology Magazine*, 19(3):42–51, 2000.
- [21] Head JF, Wang F, Lipari CA, and Elliott RL. The important role of infrared imaging in breast cancer. *IEEE Engineering in Medicine and Biology Magazine*, 19(3):52–57, 2000.
- [22] Keyserlingk JR, Yassa M, Ahlgren P, and Belliveau N. Preliminary evaluation of preoperative chemohormonotherapy-induced reduction of the functional infrared imaging score in patients with locally advanced breast cancer. In *CDROM Proceedings of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Istanbul, Turkey, October 2001.
- [23] Qi H and Head JF. Asymmetry analysis using automated segmentation and classification for breast cancer detection in thermograms. In *CDROM Proceedings of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Istanbul, Turkey, October 2001.
- [24] Merla A, Ledda A, Di Donato L, Di Luzio S, and Romani GL. Use of infrared functional imaging to detect impaired thermoregulatory control in men with asymptomatic varicocele. *Fertility and Sterility*, 78(1):199–200, 2002.
- [25] Merla A, Ledda A, Di Donato L, and Romani GL. Diagnosis of sub-clinical varicocele by means of infrared functional imaging. In *CDROM Proceedings of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Istanbul, Turkey, October 2001.
- [26] Vlaisavljević V. A comparative study of the diagnostic value of telethermography and contact thermography in the diagnosis of varicocele. In Zorгниotti AW, editor, *Temperature and Environmental Effects on the Testis*, pages 261–265. Plenum, New York NY, 1991.
- [27] Carlsen EN. Transmission spectroscopy: An improvement in light scanning. *RNM Images*, 13(2):22–25, 1983.
- [28] Sickles EA. Breast CT scanning, heavy-ion mammography, NMR imaging, and diaphanography. In Feig SA and McLelland R, editors, *Breast Carcinoma: Current Diagnosis and Treatment*, pages 233–250. Masson, New York, NY, 1983.
- [29] Kopans DB. Nonmammographic breast imaging techniques: Current status and future developments. In Sickles EA, editor, *The Radiologic Clinics of North America*, volume 25, number 5, pages 961–971. Saunders, Philadelphia, PA, September 1987.
- [30] Bozzola JJ and Russell LD. *Electron Microscopy: Principles and Techniques for Biologists*. Jones and Bartlett, Sudbury, MA, 2nd edition, 1999.

- [31] Rangayyan RM. *Biomedical Signal Analysis – A Case-Study Approach*. IEEE and Wiley, New York, NY, 2002.
- [32] Frank C, Woo SLY, Amiel D, Harwood F, Gomez M, and Akeson W. Medial collateral ligament healing – A multidisciplinary assessment in rabbits. *The American Journal of Sports Medicine*, 11(6):379–389, 1983.
- [33] Frank C, McDonald D, Bray D, Bray R, Rangayyan R, Chimich D, and Shrive N. Collagen fibril diameters in the healing adult rabbit medial collateral ligament. *Connective Tissue Research*, 27:251–263, 1992.
- [34] Gubler MC, Heidet L, and Antignac C. Alport’s syndrome, thin basement membrane nephropathy, nail-patella syndrome, and Type III collagen glomerulopathy. In Jennette JC, Olson JL, Schwartz MM, and Silva FG, editors, *Heptinstall’s Pathology of the Kidney*, pages 1207–1230. Lippincott-Raven, Philadelphia, PA, 5th edition, 1998.
- [35] Frank C, MacFarlane B, Edwards P, Rangayyan R, Liu ZQ, Walsh S, and Bray R. A quantitative analysis of matrix alignment in ligament scars: A comparison of movement versus immobilization in an immature rabbit model. *Journal of Orthopaedic Research*, 9(2):219–227, 1991.
- [36] Chaudhuri S, Nguyen H, Rangayyan RM, Walsh S, and Frank CB. A Fourier domain directional filtering method for analysis of collagen alignment in ligaments. *IEEE Transactions on Biomedical Engineering*, 34(7):509–518, 1987.
- [37] Liu ZQ, Rangayyan RM, and Frank CB. Statistical analysis of collagen alignment in ligaments by scale-space analysis. *IEEE Transactions on Biomedical Engineering*, 38(6):580–588, 1991.
- [38] Robb RA. X-ray computed tomography: An engineering synthesis of multiscientific principles. *CRC Critical Reviews in Biomedical Engineering*, 7:264–333, 1982.
- [39] Nudelman S and Roehrig H. Photoelectronic-digital imaging for diagnostic radiology. In Robb RA, editor, *Three-Dimensional Biomedical Imaging, Volume I*, pages 5–60. CRC Press, Boca Raton, FL, 1985.
- [40] Yaffe MJ. Development of full field digital mammography. In Karssemeijer N, Thijssen M, Hendriks J, and van Erning L, editors, *Proceedings of the 4th International Workshop on Digital Mammography*, pages 3–10, Nijmegen, The Netherlands, June 1998.
- [41] Cheung L, Bird R, Chitkara A, Rego A, Rodriguez C, and Yuen J. Initial operating and clinical results of a full field mammography system. In Karssemeijer N, Thijssen M, Hendriks J, and van Erning L, editors, *Proceedings of the 4th International Workshop on Digital Mammography*, pages 11–18, Nijmegen, The Netherlands, June 1998.
- [42] Maidment ADA, Fahrig R, and Yaffe MJ. Dynamic range requirements in digital mammography. *Medical Physics*, 20(6):1621–1633, 1993.
- [43] Herman GT. *Image Reconstruction From Projections: The Fundamentals of Computed Tomography*. Academic, New York, NY, 1980.
- [44] Macovski A. Physical problems of computerized tomography. *Proceedings of the IEEE*, 71(3):373–378, 1983.

- [45] Andersson I. Mammography in clinical practice. *Medical Radiography and Photography*, 62(2):1–41, 1986.
- [46] Larsson SA. Gamma camera emission tomography. *Acta Radiologica Supplementum* 363, 1980.
- [47] Canadian Cancer Society. *Facts on Breast Cancer*, April 1989.
- [48] National Cancer Institute of Canada, Toronto. *Annual Report of the National Cancer Institute of Canada*, 1987.
- [49] Spratt JS and Spratt JA. Growth rates. In Donegan WL and Spratt JS, editors, *Cancer of the Breast*, chapter 10, pages 270–302. Saunders, Philadelphia, PA, 3rd edition, 1988.
- [50] Feig SA, Schwartz F, Nerlinger R, and Edeiken J. Prognostic factors of breast neoplasms detected on screening by mammography and physical examination. *Radiology*, 133:577–582, 1979.
- [51] McLelland R. Screening for breast cancer: Opportunities, status and challenges. In Brünner S and Langfeldt B, editors, *Recent Results in Cancer Research*, volume 119, pages 29–38. Springer-Verlag, Berlin, Germany, 1990.
- [52] Basset LW and Gold RH, editors. *Breast cancer detection: Mammography and other methods in breast imaging*. Grune & Stratton, Orlando, FL, 2nd edition, 1987.
- [53] Cardenosa G. *Breast Imaging Companion*. Lippincott-Raven, Philadelphia, PA, 1997.
- [54] Homer MJ. *Mammographic Interpretation: A Practical Approach*. McGraw-Hill, Boston, MA, 2nd edition, 1997.
- [55] Heywang-Köbrunner SH, Schreer I, and Dershaw DD. *Diagnostic Breast Imaging*. Thieme, Stuttgart, Germany, 1997.
- [56] Haus AG. Recent trends in screen-film mammography: Technical factors and radiation dose. In Brünner S and Langfeldt B, editors, *Recent Results in Cancer Research*, volume 105, pages 37–51. Springer-Verlag, Berlin, Germany, 1987.
- [57] Warren SL. A roentgenologic study of the breast. *American Journal of Roentgenology and Radiation Therapy*, 24:113–124, 1930.
- [58] Egan RL. Experience with mammography in a tumor institute. Evaluation of 1000 studies. *Radiology*, 75:894–900, 1960.
- [59] Sickles EA and Weber WN. High-contrast mammography with a moving grid: Assessment of clinical utility. *American Journal of Radiology*, 146:1137–1139, 1986.
- [60] Sickles EA. The role of magnification technique in modern mammography. In Brünner S and Langfeldt B, editors, *Recent Results in Cancer Research*, volume 105, pages 19–24. Springer-Verlag, Berlin, Germany, 1987.
- [61] Alberta Cancer Board, Alberta, Canada. *Screen Test: Alberta Program for the Early Detection of Breast Cancer – 1999/2001 Biennial Report*, 2001.
- [62] Edholm P. The tomogram – Its formation and content. *Acta Radiologica*, Supplement No. 193:1–109, 1960.

- [63] Rangayyan RM and Kantzas A. Image reconstruction. In Webster JG, editor, *Wiley Encyclopedia of Electrical and Electronics Engineering, Supplement 1*, pages 249–268. Wiley, New York, NY, 2000.
- [64] Rangayyan RM. Computed tomography techniques and algorithms: A tutorial. *Innovation et Technologie en Biologie et Medecine*, 7(6):745–762, 1986.
- [65] Boyd DP, Gould RG, Quinn JR, and Sparks R. Proposed dynamic cardiac 3-D densitometer for early detection and evaluation of heart disease. *IEEE Transactions on Nuclear Science*, NS-26(2):2724–2727, 1979.
- [66] Boyd DP and Lipton MJ. Cardiac computed tomography. *Proceedings of the IEEE*, 71(3):298–307, 1983.
- [67] Radon J. Über die bestimmung von funktionen durch ihre integralwerte längs gewisser mannigfaltigkeiten. *Berichte der Sächsischen Akademie der Wissenschaft*, 69:262–277, 1917.
- [68] Radon J. On the determination of functions from their integral values along certain manifolds (English translation by Parks PC). *IEEE Transactions on Medical Imaging*, 5(4):170–176, 1986.
- [69] Cormack AM. Representation of a function by its line integrals, with some radiological applications. *Journal of Applied Physics*, 34:2722–2727, 1963.
- [70] Cormack AM. Representation of a function by its line integrals, with some radiological applications II. *Journal of Applied Physics*, 35:2908–2913, 1964.
- [71] Bracewell RN and Riddle AC. Inversion of fan-beam scans in radio astronomy. *The Astrophysical Journal*, 150:427–434, 1967.
- [72] Crowther RA, Amos LA, Finch JT, De Rosier DJ, and Klug A. Three dimensional reconstructions of spherical viruses by Fourier synthesis from electron micrographs. *Nature*, 226:421–425, 1970.
- [73] De Rosier DJ and Klug A. Reconstruction of three dimensional images from electron micrographs. *Nature*, 217:130–134, 1968.
- [74] Ramachandran GN and Lakshminarayanan AV. Three-dimensional reconstruction from radiographs and electron micrographs: Application of convolutions instead of Fourier transforms. *Proceedings of the National Academy of Science, USA*, 68:2236–2240, 1971.
- [75] Gordon R, Bender R, and Herman GT. Algebraic Reconstruction Techniques (ART) for three-dimensional electron microscopy. *Journal of Theoretical Biology*, 29:471–481, 1970.
- [76] Oldendorf WH. Isolated flying spot detection of radio-density discontinuities – Displaying the internal structural pattern of a complex object. *IRE Transactions on Bio-Medical Electronics*, BME-8:68–72, 1961.
- [77] Hounsfield GN. Computerized transverse axial scanning (tomography) Part I: Description of system. *British Journal of Radiology*, 46:1016–1022, 1973.
- [78] Ambrose J. Computerized transverse axial scanning (tomography) Part II: Clinical application. *British Journal of Radiology*, 46:1023–1047, 1973.
- [79] Robb RA, Hoffman EA, Sinak LJ, Harris LD, and Ritman EL. High-speed three-dimensional x-ray computed tomography: The Dynamic Spatial Reconstructor. *Proceedings of the IEEE*, 71(3):308–319, 1983.

- [80] Herman GT. *Image Reconstruction From Projections: Implementation and Applications*. Springer-Verlag, Berlin, Germany, 1979.
- [81] Knoll GF. Single-photon emission computed tomography. *Proceedings of the IEEE*, 71(3):320–329, 1983.
- [82] Kak AC and Slaney M. *Principles of Computerized Tomographic Imaging*. IEEE, New York, NY, 1988.
- [83] Greenleaf JF. Computerized tomography with ultrasound. *Proceedings of the IEEE*, 71(3):330–337, 1983.
- [84] Hinshaw WS and Lent AH. An introduction to NMR imaging: From the Bloch equation to the imaging equation. *Proceedings of the IEEE*, 71(3):338–350, 1983.
- [85] Müller G, Chance B, Alfano R, Arridge S, Beuthan J, Gratton E, Kaschke M, Masters B, Svanberg S, and van der Zee P, editors. *Medical Optical Tomography: Functional Imaging and Monitoring*. SPIE, Bellingham, WA, 1993.
- [86] Boulfelfel D. *Restoration of Nuclear Medicine Images*. PhD thesis, Department of Electrical and Computer Engineering, University of Calgary, Calgary, Alberta, Canada, November 1992.
- [87] Boulfelfel D, Rangayyan RM, Hahn LJ, and Kloiber R. Use of the geometric mean of opposing planar projections in pre-reconstruction restoration of SPECT images. *Physics in Medicine and Biology*, 37(10):1915–1929, 1992.
- [88] Ter-Pogossian MM. Positron emission tomography (PET). In Robb RA, editor, *Three-Dimensional Biomedical Imaging, Volume II*, pages 41–56. CRC Press, Boca Raton, FL, 1985.
- [89] Seitz RJ and Roland PE. Vibratory stimulation increases and decreases the regional cerebral blood flow and oxidative metabolism: a positron emission tomography (PET) study. *Acta Neurologica Scandinavica*, 86:60–67, 1992.
- [90] Ogawa M, Magata Y, Ouchi Y, Fukuyama H, Yamauchi H, Kimura J, Yonekura Y, and Konishi J. Scopolamine abolishes cerebral blood flow response to somatosensory stimulation in anesthetized cats: PET study. *Brain Research*, 650:249–252, 1994.
- [91] Cloutier G, Chen D, and Durand LG. Performance of time-frequency representation techniques to measure blood flow turbulence with pulsed-wave Doppler ultrasound. *Ultrasound in Medicine and Biology*, 27(4):535–550, 2001.
- [92] Fenster A and Downey DB. Three-dimensional ultrasound imaging of the prostate. In *Proceedings SPIE 3659: Medical Imaging – Physics of Medical Imaging*, pages 2–11, San Diego, CA, February 1999.
- [93] Robinson BS and Greenleaf JF. Computerized ultrasound tomography. In Robb RA, editor, *Three-Dimensional Biomedical Imaging, Volume II*, pages 57–78. CRC Press, Boca Raton, FL, 1985.
- [94] Lauterbur PC and Lai CM. Zeugmatography by reconstruction from projections. *IEEE Transactions on Nuclear Science*, 27(3):1227–1231, 1980.

- [95] Liang ZP and Lauterbur PC. *Principles of Magnetic Resonance Imaging: A Signal Processing Perspective*. IEEE, New York, NY, 2000.
- [96] Hill BC and Hinshaw WS. Fundamentals of NMR imaging. In Robb RA, editor, *Three-Dimensional Biomedical Imaging, Volume II*, pages 79–124. CRC Press, Boca Raton, FL, 1985.
- [97] Tompkins WJ. *Biomedical Digital Signal Processing*. Prentice Hall, Upper Saddle River, NJ, 1995.
- [98] Cromwell L, Weibell FJ, and Pfeiffer EA. *Biomedical Instrumentation and Measurements*. Prentice Hall, Englewood Cliffs, NJ, 2nd edition, 1980.
- [99] Aston R. *Principles of Biomedical Instrumentation and Measurement*. Merrill, Columbus, OH, 1990.
- [100] Webster JG, editor. *Medical Instrumentation: Application and Design*. Wiley, New York, NY, 3rd edition, 1998.
- [101] Bronzino JD. *Biomedical Engineering and Instrumentation*. PWS Engineering, Boston, MA, 1986.
- [102] Bronzino JD, editor. *The Biomedical Engineering Handbook*. CRC Press, Boca Raton, FL, 1995.
- [103] Bartlett J. *Familiar Quotations*. Little, Brown and Co., Boston, MA, 15th edition, 1980.
- [104] Chakraborty D, Pfeiffer DE, and Brikman I. Perceptual noise measurement of displays. In *Proceedings SPIE 1443: Medical Imaging V – Image Physics*, pages 183–190, San Jose, CA, February 1991.
- [105] Eckert MP and Chakraborty D. Quantitative analysis of phantom images in mammography. In *Proceedings SPIE 2167: Medical Imaging – Image Processing*, pages 887–899, Newport Beach, CA, February 1994.
- [106] Chakraborty DP. Physical measures of image quality in mammography. In *Proceedings SPIE 2708: Medical Imaging 1996 – Physics of Medical Imaging*, pages 179–193, Newport Beach, CA, February 1996.
- [107] Chakraborty DP and Eckert MP. Quantitative versus subjective evaluation of mammography accreditation phantom images. *Medical Physics*, 22(2):133–143, 1995.
- [108] Chakraborty DP. Computer analysis of mammography phantom images (CAMPI): An application to the measurement of microcalcification image quality of directly acquired digital images. *Medical Physics*, 24(8):1269–1277, 1997.
- [109] Bijkerk KR, Thijssen MAO, and Arnoldussen TJM. Modification of the CDMAM contrast-detail phantom for image quality evaluation of full-field digital mammography systems. In Yaffe MJ, editor, *Proceedings of the 5th International Workshop on Digital Mammography*, pages 633–640, Toronto, Canada, June 2000.
- [110] Furuie S, Herman GT, Narayan TK, Kinahan PE, Karp JS, Lewitt RM, and Matej S. A methodology for testing statistically significant differences between fully 3D PET reconstruction algorithms. *Physics in Medicine and Biology*, 39:341–354, 1994.

- [111] Barrett HH. Objective assessment of image quality: Effects of quantum noise and object variability. *Journal of the Optical Society of America - A*, 7(7):1266–1278, July 1990.
- [112] Kayargadde V and Martens JB. Perceptual characterization of images degraded by blur and noise: model. *Journal of the Optical Society of America - A*, 13(6):1178–1188, June 1996.
- [113] Kayargadde V and Martens JB. Perceptual characterization of images degraded by blur and noise: experiments. *Journal of the Optical Society of America - A*, 13(6):1166–1177, June 1996.
- [114] Perrin FH. Methods of appraising photographic systems Part I – Historical review. *Journal of the Society of Motion Picture and Television Engineers*, 69(3):151–156, 1960.
- [115] Higgins GC and Jones LA. The nature and evaluation of the sharpness of photographic images. *Journal of the Society of Motion Picture and Television Engineers*, 58(4):277–290, 1952.
- [116] Rangayyan RM and Elkadiki S. A region-based algorithm for the computation of image edge profile acutance. *Journal of Electronic Imaging*, 4(1):62–70, 1995.
- [117] Olabarriaga SD and Rangayyan RM. Subjective and objective evaluation of image sharpness – Behavior of the region-based image edge profile acutance measure. In *Proceedings SPIE 2712: Medical Imaging 1996 – Image Perception*, pages 154–162, Newport Beach, CA, February 1996.
- [118] Lloyd SP. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.
- [119] Max J. Quantizing for minimum distortion. *IEEE Transactions on Information Theory*, 6:7–12, 1960.
- [120] von Gierke HE. Transmission of vibratory energy through human tissue. In Glasser O, editor, *Medical Physics, Volume 3*, pages 661–669. Year Book Medical, Chicago, IL, 1960.
- [121] Phelps ME, Hoffman EJ, and Ter-Pogossian MM. Attenuation coefficients of various body tissues, fluids and lesions at photon energies of 18 to 136 keV. *Radiology*, 117:573–583, December 1975.
- [122] Levine MD. *Vision in Man and Machine*. McGraw-Hill, New York, NY, 1985.
- [123] Morrow WM, Paranjape RB, Rangayyan RM, and Desautels JEL. Region-based contrast enhancement of mammograms. *IEEE Transactions on Medical Imaging*, 11(3):392–406, 1992.
- [124] Rangayyan RM, Shen L, Shen Y, Desautels JEL, Bryant H, Terry TJ, Horeczko N, and Rose MS. Improvement of sensitivity of breast cancer diagnosis with adaptive neighborhood contrast enhancement of mammograms. *IEEE Transactions on Information Technology in Biomedicine*, 1(3):161–170, 1997.
- [125] Sivaramakrishna R, Obuchowski NA, Chilcote WA, Cardenosa G, and Powell KA. Comparing the performance of mammographic enhancement algorithms – a preference study. *American Journal of Roentgenology*, 175:45–51, 2000.

- [126] Shannon CE. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 623–656, 1948.
- [127] Shannon CE. Communication in the presence of noise. *Proceedings of the IRE*, 37:10–21, 1949.
- [128] Papoulis A. *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, New York, NY, 1965.
- [129] Gray RM. *Entropy and Information Theory*. Springer-Verlag, New York, NY, 1990.
- [130] Jain AK. Image data compression: A review. *Proceedings of the IEEE*, 69(3):349–389, 1981.
- [131] Goodman JW. *Introduction to Fourier Optics*. McGraw-Hill, New York, NY, 1968.
- [132] Hon TC, Rangayyan RM, Hahn LJ, and Kloiber R. Restoration of gamma camera-based nuclear medicine images. *IEEE Transactions on Medical Imaging*, 8(4):354–363, 1989.
- [133] Higashida Y, Baba Y, Hatemura M, Yoshida A, Takada T, and Takahashi M. Physical and clinical evaluation of a $2,048 \times 2,048$ -matrix image intensifier TV digital imaging system in bone radiography. *Academic Radiology*, 3(10):842–848, 1996.
- [134] Pateyron M, Peyrin F, Laval-Jeantet AM, Spanne P, Cloetens P, and Peix G. 3D microtomography of cancellous bone samples using synchrotron radiation. In *Proceedings SPIE 2708: Medical Imaging 1996 – Physics of Medical Imaging*, pages 417–426, Newport Beach, CA, February 1996.
- [135] Schade OH. Image gradation, graininess and sharpness in television and motion-picture systems, Part IV, A & B: Image analysis in photographic and television systems (Definition and sharpness). *Journal of the Society of Motion Picture and Television Engineers*, 64(11):593–617, 1955.
- [136] Schade OH. Optical and photoelectric analog of the eye. *Journal of the Optical Society of America*, 46(9):721–739, 1956.
- [137] Schade OH. An evaluation of photographic image quality and resolving power. *Journal of the Society of Motion Picture and Television Engineers*, 73(2):81–119, 1964.
- [138] Burke JJ and Snyder HL. Quality metrics of digitally derived imagery and their relation to interpreter performance. In *Proceedings SPIE 310: Image Quality*, pages 16–23, 1981.
- [139] Tapiovaara MJ and Wagner RF. SNR and noise measurements for medical imaging: I. A practical approach based on statistical decision theory. *Physics in Medicine and Biology*, 38:71–92, 1993.
- [140] Tapiovaara MJ. SNR and noise measurements for medical imaging: II. Application to fluoroscopic x-ray equipment. *Physics in Medicine and Biology*, 38:1761–1788, 1993.
- [141] Hall CF. Subjective evaluation of a perceptual quality metric. In *Proceedings SPIE 310: Image Quality*, pages 200–204, 1981.

- [142] Crane EM. An objective method for rating picture sharpness: SMT acutance. *Journal of the Society of Motion Picture and Television Engineers*, 73(8):643–647, 1964.
- [143] Crane EM. Acutance and granulance. In *Proceedings SPIE 310: Image Quality*, pages 125–132, 1981.
- [144] Yip KL. Imaging characteristics of CRT multifformat printers. In *Proceedings SPIE 1653: Image Capture, Formatting, and Display*, pages 477–487, 1992.
- [145] Kriss MA. Image analysis of discrete and continuous systems – film and CCD sensors. In *Proceedings SPIE 1398: CAN-AM Eastern*, pages 4–14, 1990.
- [146] Gendron RG. An improved objective method for rating picture sharpness: CMT acutance. *Journal of the Society of Motion Picture and Television Engineers*, 82(12):1009–1012, 1973.
- [147] Westheimer G. Spatial frequency and light spread descriptions of visual acuity and hyperacuity. *Journal of the Optical Society of America*, 67(2):207–212, 1977.
- [148] Westheimer G. The spatial sense of the eye. *Investigative Ophthalmology and Visual Science*, 18(9):893–912, 1979.
- [149] Wolfe RN and Eisen FC. Psychometric evaluation of the sharpness of photographic reproductions. *Journal of the Optical Society of America*, 43(10):914–923, 1953.
- [150] Perrin FH. Methods of appraising photographic systems, Part II – Manipulation and significance of the sine-wave response function. *Journal of the Society of Motion Picture and Television Engineers*, 69(4):239–250, 1960.
- [151] Barten PGJ. Evaluation of subjective image quality with the square-root integral method. *Journal of the Optical Society of America A*, 7(10):2024–2031, 1990.
- [152] Higgins GC. Methods for analyzing the photographic system, including the effects of nonlinearity and spatial frequency response. *Photographic Science and Engineering*, 15(2):106–118, 1971.
- [153] Granger EM and Cupery KN. An optical merit function (SQF) which correlates with subjective image judgments. *Photographic Science and Engineering*, 16(3):221–230, 1972.
- [154] Higgins GC. Image quality criteria. *Journal of Applied Photographic Engineering*, 3(2):53–60, 1977.
- [155] Task HL, Pinkus AR, and Hornseth JP. A comparison of several television display image quality measures. *Proceedings of the Society of Information Display*, 19(3):113–119, 1978.
- [156] Barten PGJ. *Contrast Sensitivity of the Human Eye and its Effects on Image Quality*. SPIE, Bellingham, WA, 1999.
- [157] Carlson CR and Cohen RW. A simple psychophysical model for predicting the visibility of displayed information. *Proceedings of the Society of Information Display*, 21(3):229–246, 1980.
- [158] Saghri JA, Cheatham PS, and Habibi A. Image quality measure based on a human visual system model. *Optical Engineering*, 28(7):813–818, 1989.

- [159] Nill NB and Bouzas BH. Objective image quality measure derived from digital image power spectra. *Optical Engineering*, 31(4):813–825, 1992.
- [160] Lukas FXJ and Budrikis ZL. Picture quality prediction based on a visual model. *IEEE Transactions on Communications*, 30(7):1679–1692, 1982.
- [161] Budrikis ZL. Visual fidelity criterion and modeling. *Proceedings of the IEEE*, 60(7):771–779, 1972.
- [162] Westerink JHDM and Roufs JAJ. A local basis for perceptually relevant resolution measures. In *Society of Information Display Digest*, pages 360–363, 1988.
- [163] Rangayyan RM, El-Faramawy NM, Desautels JEL, and Alim OA. Measures of acutance and shape for classification of breast tumors. *IEEE Transactions on Medical Imaging*, 16(6):799–810, 1997.
- [164] Rangayyan RM and Das A. Image enhancement based on edge profile acutance. *Journal of the Indian Institute of Science*, 78:17–29, 1998.
- [165] Mudigonda NR, Rangayyan RM, and Desautels JEL. Gradient and texture analysis for the classification of mammographic masses. *IEEE Transactions on Medical Imaging*, 19(10):1032–1043, 2000.
- [166] Mudigonda NR. *Image Analysis Methods for the Detection and Classification of Mammographic Masses*. PhD thesis, Department of Electrical and Computer Engineering, University of Calgary, Calgary, Alberta, Canada, December 2001.
- [167] Papoulis A. *Signal Analysis*. McGraw-Hill, New York, NY, 1977.
- [168] Bendat JS and Piersol AG. *Random Data: Analysis and Measurement Procedures*. Wiley, New York, NY, 2nd edition, 1986.
- [169] Auñón JI and Chandrasekar V. *Introduction to Probability and Random Processes*. McGraw-Hill, New York, NY, 1997.
- [170] Ramsey FL and Schafer DW. *The Statistical Sleuth — A Course in Methods of Data Analysis*. Wadsworth Publishing Company, Belmont, CA, 1997.
- [171] Riffenburgh RH. *Statistics in Medicine*. Academic, San Diego, CA, 1993.
- [172] Bailar III JC and Mosteller F, editors. *Medical Uses of Statistics*. NEJM Books, Boston, MA, 2nd edition, 1992.
- [173] Peebles Jr. PZ. *Probability, Random Variables, and Random Signal Principles*. McGraw-Hill, New York, NY, 3rd edition, 1993.
- [174] Kuduvali GR and Rangayyan RM. Performance analysis of reversible image compression techniques for high-resolution digital teleradiology. *IEEE Transactions on Medical Imaging*, 11(3):430–445, 1992.
- [175] Mascarenhas NDA. An overview of speckle noise filtering in SAR images. In Guyenne TD, editor, *Proceedings of the First Latin American Seminar on Radar Remote Sensing: Image Processing Techniques*, volume 1, pages 71–79. European Space Agency, Buenos Aires, Argentina, December 2-4 1996.
- [176] Rabiner LR and Schafer RW. *Digital Processing of Speech Signals*. Prentice Hall, Englewood Cliffs, NJ, 1978.

- [177] Trussell HJ and Hunt BR. Sectioned methods for image restoration. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(2):157–164, 1978.
- [178] Rabie TF, Rangayyan RM, and Paranjape RB. Adaptive-neighborhood image deblurring. *Journal of Electronic Imaging*, 3(4):368–378, 1994.
- [179] Jiang SS and Sawchuk AA. Noise updating repeated Wiener filter and other adaptive noise smoothing filters using local image statistics. *Applied Optics*, 25:2326–2337, July 1986.
- [180] Barner KE and Arce GR. Optimal detection methods for the restoration of images degraded by signal-dependent noise. In Pearlman WA, editor, *Proceedings of SPIE on Visual Communications and Image Processing IV*, volume 1199, pages 115–124. SPIE, 1989.
- [181] Froehlich GK, Walkup JF, and Krile TF. Estimation in signal-dependent film-grain noise. *Applied Optics*, 20:3619–3626, October 1981.
- [182] Naderi F and Sawchuk AA. Estimation of images degraded by film-grain noise. *Applied Optics*, 17:1228–1237, April 1978.
- [183] Downie JD and Walkup JF. Optimal correlation filters with signal-dependent noise. *Journal of the Optical Society of America*, 11:1599–1609, May 1994.
- [184] Lee JS. Speckle analysis and smoothing of synthetic aperture radar images. *Computer Graphics and Image Processing*, 17:24–32, 1981.
- [185] Schultze MA. An edge-enhancing nonlinear filter for reducing multiplicative noise. In Dougherty ER and Astola JT, editors, *Proceedings of SPIE on Nonlinear Image Processing VIII*, volume 3026, pages 46–56. SPIE, February 1997.
- [186] Kuan DT, Sawchuk AA, Strand TC, and Chavel P. Adaptive restoration of images with speckle. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 35:373–383, 1987.
- [187] Lim JS and Nawab H. Techniques for speckle noise removal. *Optical Engineering*, 21:472–480, 1981.
- [188] Kuan DT, Sawchuk AA, Strand TC, and Chavel P. Adaptive noise smoothing filter for images with signal-dependent noise. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-7:165–177, 1985.
- [189] Arsenault HH, Gendron C, and Denis M. Transformation of film-grain noise into signal-independent Gaussian noise. *Journal of the Optical Society of America*, 71:91–94, January 1981.
- [190] Arsenault HH and Levesque M. Combined homomorphic and local-statistics processing for restoration of images degraded by signal-dependent noise. *Applied Optics*, 23:845–850, March 1984.
- [191] Kasturi R, Walkup JF, and Krile TF. Image restoration by transformation of signal-dependent noise to signal-independent noise. *Applied Optics*, 22:3537–3542, November 1983.
- [192] Dougherty ER and Astola J. *An Introduction to Nonlinear Image Processing*. SPIE, Bellingham, WA, 1994.

- [193] Pitas I and Venetsanopoulos AN. Order statistics in digital image processing. *Proceedings of the IEEE*, 80:1893–1923, 1992.
- [194] Theilheimer F. A matrix version of the fast Fourier transform. *IEEE Transactions on Audio and Electroacoustics*, AU-17(2):158–161, 1969.
- [195] Hunt BR. A matrix theory proof of the discrete convolution theorem. *IEEE Transactions on Audio and Electroacoustics*, AU-19(4):285–288, 1971.
- [196] Hunt BR. The application of constrained least squares estimation to image restoration by digital computer. *IEEE Transactions on Computers*, C-22(9):805–812, 1973.
- [197] Helstrom CW. Image restoration by the method of least squares. *Journal of the Optical Society of America*, 57(3):297–303, 1967.
- [198] Wiener NE. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series, with Engineering Applications*. MIT Press, Cambridge, MA, 1949.
- [199] Lim JS. *Two-dimensional Signal and Image Processing*. Prentice Hall, Englewood Cliffs, NJ, 1990.
- [200] Lee JS. Digital image enhancement and noise filtering by use of local statistics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-2:165–168, March 1980.
- [201] Rangayyan RM, Ciuc M, and Faghih F. Adaptive neighborhood filtering of images corrupted by signal-dependent noise. *Applied Optics*, 37(20):4477–4487, 1998.
- [202] Aghdasi F, Ward RK, and Palcic B. Restoration of mammographic images in the presence of signal-dependent noise. In *Proceedings of SPIE vol. 1905 on Biomedical Image Processing and Biomedical Visualization*, pages 740–751, San Jose, CA, 1993.
- [203] The Math Works Inc., Natick, MA. *Image Processing Toolbox for use with MATLAB: User's Guide*, 2nd edition, 1997.
- [204] Lee JS. Refined filtering of image noise using local statistics. *Computer Graphics and Image Processing*, 15:380–389, 1981.
- [205] Chan P and Lim JS. One-dimensional processing for adaptive image restoration. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 33(1):117–129, 1985.
- [206] Hadhoud MM and Thomas DW. The two-dimensional adaptive LMS (TDLMS) algorithm. *IEEE Transactions on Circuits and Systems*, 35(5):485–494, 1988.
- [207] Song WJ and Pearlman WA. Restoration of noisy images with adaptive windowing and non-linear filtering. In *Proceedings of SPIE on Visual Communications and Image Processing*, volume 707, pages 198–206. SPIE, 1986.
- [208] Song WJ and Pearlman WA. A minimum-error, minimum-correlation filter for images. In *Proceedings of SPIE on Applications of Digital Image Processing IX*, volume 697, pages 225–232. SPIE, 1986.
- [209] Song WJ and Pearlman WA. Edge-preserving noise filtering based on adaptive windowing. *IEEE Transactions on Circuits and Systems*, 35(8):1046–1055, 1988.

- [210] Mahesh B, Song WJ, and Pearlman WA. Adaptive estimators for filtering noisy images. *Optical Engineering*, 29(5):488–494, 1990.
- [211] Paranjape RB, Rangayyan RM, and Morrow WM. Adaptive neighborhood mean and median filtering. *Journal of Electronic Imaging*, 3:360–367, October 1994.
- [212] Paranjape RB, Rabie TF, and Rangayyan RM. Image restoration by adaptive neighborhood noise subtraction. *Applied Optics*, 33:1861–1869, May 1994.
- [213] Rangayyan RM and Das A. Filtering multiplicative noise in images using adaptive region-based statistics. *Journal of Electronic Imaging*, 7:222–230, 1998.
- [214] Ciuc M, Rangayyan RM, Zaharia T, and Buzuloiu V. Filtering noise in color images using adaptive-neighborhood statistics. *Journal of Electronic Imaging*, 9(4):484–494, 2000.
- [215] Morrow WM. Region-based image processing with application to mammography. Master’s thesis, Department of Electrical and Computer Engineering, University of Calgary, Calgary, Alberta, Canada, December 1990.
- [216] Hunter CJ, Matyas JR, and Duncan NA. The three-dimensional architecture of the notochordal nucleus pulposus: novel observations on cell structures in the canine intervertebral disc. *Journal of Anatomy*, 202(3):279–291, 2003.
- [217] University of California, Los Angeles, CA. *The Confocal Microscope*: http://www.gonda.ucla.edu/bri_core/confocal.htm, accessed April 2002.
- [218] University of British Columbia, Vancouver, BC, Canada. *Introduction to Laser Scanning Confocal Microscopy*: <http://www.cs.ubc.ca/spider/ladic/intro.html>, accessed April 2002.
- [219] Al-Kofahi KA, Lasek S, Szarowski DH, Pace CJ, Nagy G, Turner JN, and Roysam B. Rapid automated three-dimensional tracing of neurons from confocal image stacks. *IEEE Transactions on Information Technology in Biomedicine*, 6(2):171–187, 2002.
- [220] Haralick RM, Sternberg SR, and Zhuang X. Image analysis using mathematical morphology. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(4):532–550, 1987.
- [221] Giardina CR and Dougherty ER. *Morphological Methods in Image and Signal Processing*. Prentice Hall, Englewood Cliffs, NJ, 1988.
- [222] Dougherty ER. *An Introduction to Morphological Image Processing*. SPIE, Bellingham, WA, 1992.
- [223] Meijering EHW. *Image Enhancement in Digital X-ray Angiography*. PhD thesis, Image Sciences Institute, University Medical Center Utrecht, Utrecht, The Netherlands, 2000.
- [224] Meijering EHW, Zuiderveld KJ, and Viergever MA. Image registration for digital subtraction angiography. *International Journal of Computer Vision*, 31(2/3):227–246, 1999.
- [225] Meijering EHW, Niejessen WJ, Bakker J, van der Molen AJ, de Kort GAP, Lo RTH, Mali WPTM, and Viergever MA. Reduction of patient motion

- artifacts in digital subtraction angiography: Evaluation of a fast and fully automatic technique. *Radiology*, 219:288–293, 2001.
- [226] MacMahon H. Improvement in detection of pulmonary nodules: Digital image processing and computer-aided diagnosis. *RadioGraphics*, 20:1169–1177, 2000.
- [227] MacMahon H. Energy subtraction: The University of Chicago Hospitals' observer test. *Insights & Images, Fujifilm Medical Systems, Stamford, CT, USA, Late Summer:1–2*, 1999.
- [228] Mazur AK, Mazur EJ, and Gordon R. Digital differential radiography: a new diagnostic procedure for locating neoplasms, such as breast cancers, in soft, deformable tissues. In *Proceedings of SPIE vol. 1905 on Biomedical Image Processing and Biomedical Visualization*, pages 443–455, San Jose, CA, 1993.
- [229] Lindley CA. *Practical Image Processing in C*. Wiley, New York, NY, 1991.
- [230] Paranjape RB, Morrow WM, and Rangayyan RM. Adaptive-neighborhood histogram equalization for image enhancement. *CVGIP: Graphical Models and Image Processing*, 54(3):259–267, May 1992.
- [231] Ketchum DJ. Real-time image enhancement techniques. In *Proceedings SPIE/OSA 74*, pages 120–125, 1976.
- [232] Pizer SM, Amburn EP, Austin JD, Cromartie R, Geselowitz A, Geer T, ter Haar Remeny B, Zimmerman JB, and Zuiderveld K. Adaptive histogram equalization and its variations. *Computer Vision, Graphics, and Image Processing*, 39:355–368, 1987.
- [233] Leszczynski KW and Shalev S. A robust algorithm for contrast enhancement by local histogram modification. *Image and Vision Computing*, 7(3):205–209, 1989.
- [234] Rehm K and Dallas WJ. Artifact suppression in digital chest radiographs enhanced with adaptive histogram equalization. In *Proceedings SPIE 1092*, pages 220–300, 1989.
- [235] Buzuloiu V, Ciuc M, Rangayyan RM, and Vertan C. Adaptive-neighborhood histogram equalization of color images. *Journal of Electronic Imaging*, 10(4):445–459, 2001.
- [236] Bogert BP, Healy MJR, and Tukey JW. The quefrency analysis of time series for echoes: Cepstrum, pseudo-autocovariance, cross-cepstrum, and saphe cracking. In Rosenblatt M, editor, *Proceedings of the Symposium on Time Series Analysis*, pages 209–243. Wiley, New York, NY, 1963.
- [237] Oppenheim AV, Schaffer RW, and Stockham Jr. TG. Nonlinear filtering of multiplied and convolved signals. *Proceedings of the IEEE*, 56(8):1264–1291, 1968.
- [238] Oppenheim AV and Schaffer RW. Homomorphic analysis of speech. *IEEE Transactions on Audio and Electroacoustics*, AU-16(2):221–226, 1968.
- [239] Childers DG, Skinner DP, and Kemerait RC. The cepstrum: A guide to processing. *Proceedings of the IEEE*, 65(10):1428–1443, 1977.

- [240] Stockham Jr. TG. Image processing in the context of a visual model. *Proceedings of the IEEE*, 60(7):828–842, 1972.
- [241] Yoon JH, Ro YM, Kim SI, and Park DS. Contrast enhancement of mammography image using homomorphic filter in wavelet domain. In Yaffe MJ, editor, *Proceedings of the 5th International Workshop on Digital Mammography*, pages 617–623, Toronto, Canada, June 2000.
- [242] Gordon R and Rangayyan RM. Feature enhancement of film mammograms using fixed and adaptive neighborhoods. *Applied Optics*, 23(4):560–564, February 1984.
- [243] Rangayyan RM and Nguyen HN. Pixel-independent image processing techniques for enhancement of features in mammograms. In *Proceedings of the 8th IEEE Engineering in Medicine and Biology Conference*, pages 1113–1117, 1986.
- [244] Rangayyan RM and Nguyen HN. Pixel-independent image processing techniques for noise removal and feature enhancement. In *IEEE Pacific Rim Conference on Communications, Computers, and Signal Processing*, pages 81–84, Vancouver, June 1987. IEEE.
- [245] Pavlidis T. *Algorithms for Graphics and Image Processing*. Computer Science Press, Rockville, MD, 1982.
- [246] Rangayyan RM, Alto H, and Gavrilov D. Parallel implementation of the adaptive neighborhood contrast enhancement technique using histogram-based image partitioning. *Journal of Electronic Imaging*, 10:804–813, 2001.
- [247] Dhawan AP, Buelloni G, and Gordon R. Enhancement of mammographic features by optimal adaptive neighborhood image processing. *IEEE Transactions on Medical Imaging*, 5(1):8–15, 1986.
- [248] Dronkers DJ and Zwaag HV. Photographic contrast enhancement in mammography. *Radiologia Clinica et Biologica*, 43:521–528, 1974.
- [249] McSweeney MB, Sprawls P, and Egan RL. Enhanced-image mammography. In *Recent Results in Cancer Research*, volume 90, pages 79–89. Springer-Verlag, Berlin, Germany, 1984.
- [250] Askins BS, Brill AB, Rao GUV, and Novak GR. Autoradiographic enhancement of mammograms. *Diagnostic Radiology*, 130:103–107, 1979.
- [251] Bankman IN, editor. *Handbook of Medical Imaging: Processing and Analysis*. Academic Press, London, UK, 2000.
- [252] Ram G. Optimization of ionizing radiation usage in medical imaging by means of image enhancement techniques. *Medical Physics*, 9(5):733–737, 1982.
- [253] Rogowska J, Preston K, and Sashin D. Evaluation of digital unsharp masking and local contrast stretching as applied to chest radiographs. *IEEE Transactions on Biomedical Engineering*, 35(10):817–827, 1988.
- [254] Chan HP, Vyborny CJ, MacMahon H, Metz CE, Doi K, and Sickles EA. ROC studies of the effects of pixel size and unsharp-mask filtering on the detection of subtle microcalcifications. *Investigative Radiology*, 22:581–589, 1987.

- [255] Dhawan AP and Le Royer E. Mammographic feature enhancement by computerized image processing. *Computer Methods and Programs in Biomedicine*, 27:23–35, 1988.
- [256] Ji TL, Sundareshan MK, and Roehrig H. Adaptive image contrast enhancement based on human visual properties. *IEEE Transactions on Medical Imaging*, 13(4):573–586, 1994.
- [257] Laine AF, Schuler S, Fan J, and Huda W. Mammographic feature enhancement by multiscale analysis. *IEEE Transactions on Medical Imaging*, 13(4):725–740, December 1994.
- [258] Vuylsteke P and Schoeters E. Multiscale image contrast amplification (MUSIC). In *Proceedings of SPIE on Medical Imaging 1994: Image Processing*, volume 2167, pages 551–560, 1994.
- [259] Belikova T, Lashin V, and Zaltsman I. Computer assistance in the digitized mammogram processing to improve diagnosis of breast lesions. In *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 69–78, York, England, 10-12 July 1994.
- [260] Qu G, Huda W, Laine A, Steinbach B, and Honeyman J. Use of accreditation phantoms and clinical images to evaluate mammography image processing algorithms. In *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 345–354, York, England, 10-12 July 1994.
- [261] Tahoces PG, Correa J, Souto M, Gonzalez C, Gomez L, and Vidal JJ. Enhancement of chest and breast radiographs by automatic spatial filtering. *IEEE Transactions on Medical Imaging*, 10(3):330–335, 1991.
- [262] Qian W, Clarke LP, Kallergi M, and Clark RA. Tree-structured nonlinear filters in digital mammography. *IEEE Transactions on Medical Imaging*, 13(4):25–36, 1994.
- [263] Chen J, Flynn MJ, and Rebner M. Regional contrast enhancement and data compression for digital mammographic images. In *Proceedings of SPIE on Biomedical Image Processing and Biomedical Visualization*, volume SPIE-1905, pages 752–758, San Jose, CA, February 1993.
- [264] Kimme-Smith C, Gold RH, Bassett LW, Gormley L, and Morioka C. Diagnosis of breast calcifications: Comparison of contact, magnified, and television-enhanced images. *American Journal of Roentgenology*, 153:963–967, 1989.
- [265] Simpson K and Bowyer KW. A comparison of spatial noise filtering techniques for digital mammography. In *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 325–334, York, England, 10-12 July 1994.
- [266] Rangayyan RM, Shen L, Paranjape RB, Desautels JEL, MacGregor JH, Morrish HF, Burrowes P, Share S, and MacDonald FR. An ROC evaluation of adaptive neighborhood contrast enhancement for digitized mammography. In *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 307–314, York, England, 10-12 July 1994.
- [267] Kallergi M, Clarke LP, Qian W, Gavrielides M, Venugopal P, Berman CG, Holman-Ferris SD, Miller MS, and Clark RA. Interpretation of calcifications

- in screen/film, digitized, and wavelet-enhanced monitor-displayed mammograms: A receiver operating characteristic study. *Academic Radiology*, 3:285–293, 1996.
- [268] Laine A, Fan J, and Schuler S. A framework for contrast enhancement by dyadic wavelet analysis. In *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 91–100, York, England, 10–12 July 1994.
- [269] Laine A, Fan J, and Yan WH. Wavelets for contrast enhancement of digital mammography. *IEEE Engineering in Medicine and Biology Magazine*, 14(5):536–550, September/October 1995.
- [270] Qian W, Clarke LP, and Zheng BY. Computer assisted diagnosis for digital mammography. *IEEE Engineering in Medicine and Biology Magazine*, 14(5):561–569, September/October 1995.
- [271] Shen L, Shen Y, Rangayyan RM, Desautels JEL, Bryant H, Terry TJ, and Horeczko N. Earlier detection of interval breast cancers with adaptive neighborhood contrast enhancement of mammograms. In *Proceedings of SPIE on Medical Imaging 1996: Image Processing*, volume 2710, pages 940–949, Newport Beach, CA, February 1996.
- [272] Palcic B, MacAulay C, Shlien A, Treurniet W, Tezcan H, and Anderson G. Comparison of three different methods for automated classification of cervical cells. *Analytical Cellular Pathology*, 4:429–441, 1992.
- [273] Harauz G, Chiu DKY, MacAulay C, and Palcic B. Probabilistic inference in computer-aided screening for cervical cancer: an event covering approach to information extraction and decision rule formulation. *Analytical Cellular Pathology*, 6:37–50, 1994.
- [274] Shen L, Rangayyan RM, and Desautels JEL. Detection and classification of mammographic calcifications. *International Journal of Pattern Recognition and Artificial Intelligence*, 7(6):1403–1416, 1993.
- [275] Mudigonda NR, Rangayyan RM, and Desautels JEL. Detection of breast masses in mammograms by density slicing and texture flow-field analysis. *IEEE Transactions on Medical Imaging*, 20(12):1215–1227, 2001.
- [276] Guliato D, Rangayyan RM, Carnielli WA, Zuffo JA, and Desautels JEL. Segmentation of breast tumors in mammograms using fuzzy sets. *Journal of Electronic Imaging*, 12(3):369–378, 2003.
- [277] Guliato D, Rangayyan RM, Carnielli WA, Zuffo JA, and Desautels JEL. Fuzzy fusion operators to combine results of complementary medical image segmentation techniques. *Journal of Electronic Imaging*, 12(3):379–389, 2003.
- [278] Ferrari RJ, Rangayyan RM, Desautels JEL, Borges RA, and Frère AF. Automatic identification of the pectoral muscle in mammograms. *IEEE Transactions on Medical Imaging*, 23:232–245, 2004.
- [279] Ferrari RJ, Rangayyan RM, Desautels JEL, Borges RA, and Frère AF. Identification of the breast boundary in mammograms using active contour models. *Medical and Biological Engineering and Computing*, 42:201–208, 2004.

- [280] Ferrari RJ, Rangayyan RM, Borges RA, and Frère AF. Segmentation of the fibro-glandular disc in mammograms using Gaussian mixture modelling. *Medical and Biological Engineering and Computing*, 42:378–387, 2004.
- [281] Marr D and Hildreth E. Theory of edge detection. *Journal of the Royal Society of London B*, 207:187–217, 1980.
- [282] Marr D. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. WH Freeman, San Francisco, CA, 1982.
- [283] Limb JO. A model of threshold vision incorporating inhomogeneity of the visual fields. *Vision Research*, 17(4):571–584, 1977.
- [284] Witkin AP. Scale-space filtering. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1019–1022, Karlsruhe, Germany, August 1983.
- [285] Yuille AL and Poggio TA. Scale theorems for zero-crossing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(1):15–25, 1986.
- [286] Babaud J, Witkin AP, Baudin M, and Duda RO. Uniqueness of the Gaussian kernel for scale-space filtering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(1):26–33, 1986.
- [287] Mokhtarian F and Mackworth A. Scale-based description and recognition of planar curves and two-dimensional shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(1):34–43, 1986.
- [288] Bischof WF and Caelli TM. Parsing scale-space and spatial stability analysis. *Computer Vision, Graphics, and Image Processing*, 42:192–205, 1988.
- [289] Caelli TM, Bischof WF, and Liu ZQ. Filter-based approaches to pattern recognition. *Pattern Recognition*, 6(6):639–650, 1988.
- [290] Hummel RA. Representations based on zero-crossings in scale-space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 204–209, Miami Beach, FL, 1986.
- [291] Binford T. Image understanding: intelligent systems. In *Proceedings of the DARPA Image Understanding Workshop*, pages 18–31, Los Angeles, CA, 1987. Defence Advanced Research Projects Agency.
- [292] Rotem D and Zeevi YY. Image reconstruction from zero-crossings. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 34(5):1269–1277, 1986.
- [293] Katz I. *Coaxial stereo and scale-based matching*. Department of Computer Science, University of British Columbia, Vancouver, BC, Canada, 1985.
- [294] Clark JJ. Authenticating edges produced by zero crossing algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(1):43–57, 1989.
- [295] Liu ZQ, Rangayyan RM, and Frank CB. Directional analysis of images in scale space. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(11):1185–1192, 1991.

- [296] Huertas A and Medioni G. Detection of intensity changes with subpixel accuracy using Laplacian-Gaussian masks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(5):651–664, 1986.
- [297] Berzins V. Accuracy of Laplacian detectors. *Computer Vision, Graphics, and Image Processing*, 27:195–210, 1984.
- [298] Grimson WEL and Hildreth EC. Comments on ‘Digital step edges from zero crossings of second directional derivatives’. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 7(1):121–126, 1985.
- [299] Chen JS and Medioni G. Detection, localization, and estimation of edges. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(2):191–198, 1989.
- [300] Marr D and Poggio T. *Some comments on a recent theory of stereopsis*. MIT Artificial Intelligence Laboratory Memo, MIT, Boston, MA, 1980.
- [301] Richter J and Ullman S. Non-linearities in cortical simple cells and the possible detection of zero-crossings. *Biological Cybernetics*, 53:195–202, 1986.
- [302] Canny J. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(6):670–698, 1986.
- [303] Davis LS. A survey of edge detection techniques. *Computer Graphics and Image Processing*, 4:248–270, 1975.
- [304] Fu KS and Mui JK. A survey on image segmentation. *Pattern Recognition*, 13:3–16, 1981.
- [305] Haralick RM and Shapiro LG. Image segmentation techniques. *Computer Vision, Graphics, and Image Processing*, 29:100–132, 1985.
- [306] Sahoo PK, Soltani S, and Wong AKC. A survey of thresholding techniques. *Computer Vision, Graphics, and Image Processing*, 41:233–260, 1988.
- [307] Pal NR and Pal SK. A review on image segmentation techniques. *Pattern Recognition*, 26:1277–1294, 1993.
- [308] German D, Graffigne C, and Dong P. Boundary detection by constrained optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):609–628, 1990.
- [309] Zucker SW. Region growing: Childhood and adolescence. *Computer Graphics and Image Processing*, 5:382–399, 1976.
- [310] Cheevasuvit F, Maitre H, and Vidal-Madjar D. A robust method for picture segmentation based on split-and-merge procedure. *Computer Vision, Graphics, and Image Processing*, 34:268–281, 1986.
- [311] Adams R and Bischof L. Seeded region growing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(6):641–647, 1994.
- [312] Chang YL and Li XB. Adaptive image region-growing. *IEEE Transactions on Image Processing*, 3(6):868–872, 1994.
- [313] Besl PJ and Jain RC. Segmentation through variable-order surface fitting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10:167–192, 1988.

- [314] Pavlidis T and Liow YT. Integrating region growing and edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(3):225–233, March 1990.
- [315] Meyer F and Beucher S. Morphological segmentation. *Journal of Visual Communication and Image Representation*, 1:21–46, 1990.
- [316] Haddon JF and Boyce JF. Image segmentation by unifying region and boundary information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(10):929–948, 1990.
- [317] Moigne JL and Tilton JC. Refining image segmentation by integration of edge and region data. *IEEE Transactions on Geoscience and Remote Sensing*, 33(3):605–615, 1995.
- [318] LaValle SM and Hutchinson SA. A Bayesian segmentation methodology for parametric image models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(2):211–217, 1995.
- [319] Won CS and Derin H. Unsupervised segmentation of noisy and textured images using Markov random fields. *CVGIP: Graphical Models and Image Processing*, 54(4):308–328, 1992.
- [320] Shen L. *Region-based Adaptive Image Processing Techniques for Mammography*. PhD thesis, Department of Electrical and Computer Engineering, University of Calgary, Calgary, Alberta, Canada, August 1998.
- [321] Shen L and Rangayyan RM. A segmentation-based lossless image coding method for high-resolution medical image compression. *IEEE Transactions on Medical Imaging*, 16(3):301–307, 1997.
- [322] Rangayyan RM, Shen L, Shen Y, Rose MS, Desautels JEL, Bryant HE, Terry TJ, and Horeczko N. Region-based contrast enhancement. In Strickland RN, editor, *Image-Processing Techniques for Tumor Detection*, pages 213–242. Marcel Dekker, New York, NY, 2002.
- [323] Sezan MI, Yip KL, and Daly SJ. Uniform perceptual quantization: Applications to digital radiography. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-17:622–634, 1987.
- [324] Sickles EA. Mammographic features of malignancy found during screening. In Brüner S and Langfeldt B, editors, *Recent Results in Cancer Research*, volume 119, pages 88–93. Springer-Verlag, Berlin, Germany, 1990.
- [325] Spiesberger W. Mammogram inspection by computer. *IEEE Transactions on Biomedical Engineering*, BME-26(4):213–219, 1979.
- [326] Chan HP, Doi K, Galhotra S, Vyborny CJ, MacMahon H, and Jokich PM. Image feature analysis and computer-aided diagnosis in digital radiography: I. Automated detection of microcalcifications in mammography. *Medical Physics*, 14(4):538–548, 1987.
- [327] Chan HP, Doi K, Vyborny CJ, Lam KL, and Schmidt RA. Computer-aided detection of microcalcifications in mammograms: Methodology and preliminary clinical study. *Investigative Radiology*, 23(9):664–671, 1988.
- [328] Chan HP, Doi K, Vyborny CJ, Schmidt RA, Metz CE, Lam KL, Ogura T, Wu YZ, and MacMahon H. Improvement in radiologists' detection of

- clustered microcalcifications on mammograms: The potential of computer-aided diagnosis. *Investigative Radiology*, 25(10):1102–1110, 1990.
- [329] Davies DH and Dance DR. Automatic computer detection of clustered calcifications in digital mammograms. *Physics in Medicine Biology*, 35(8):1111–1118, 1990.
- [330] Fam BW, Olson SL, Winter PF, and Scholz FJ. Algorithm for the detection of fine clustered calcifications on film mammograms. *Radiology*, 169:333–337, 1988.
- [331] Karssemeijer N. Adaptive noise equalization and recognition of microcalcification clusters in mammograms. *International Journal of Pattern Recognition and Artificial Intelligence*, 7:1357–1376, December 1993.
- [332] Brzakovic D and Neskovic M. Mammogram screening using multiresolution-based image segmentation. *International Journal of Pattern Recognition and Artificial Intelligence*, 7:1437–1460, December 1993.
- [333] Netsch T. A scale-space approach for the detection of clustered microcalcifications in digital mammograms. In *Proceedings of the 3rd International Workshop on Digital Mammography*, pages 301–306, Chicago, IL, 9-12 June 1996.
- [334] Shen L, Rangayyan RM, and Desautels JEL. Application of shape analysis to mammographic calcifications. *IEEE Transactions on Medical Imaging*, 13(2):263–274, 1994.
- [335] Bankman IN, Nizialek T, Simon I, Gatewood OB, Weinberg IN, and Brody WR. Segmentation algorithms for detecting microcalcifications in mammograms. *IEEE Transactions on Information Technology in Biomedicine*, 1(2):141–149, 1997.
- [336] Serrano C, Trujillo JD, Acha B, and Rangayyan RM. Use of 2D linear prediction error to detect microcalcifications in mammograms. In *CDROM Proceedings of the II Latin American Congress on Biomedical Engineering*, Havana, Cuba, 23-25 May 2001.
- [337] Kuduvalli GR and Rangayyan RM. An algorithm for direct computation of 2-D linear prediction coefficients. *IEEE Transactions on Signal Processing*, 41(2):996–1000, 1993.
- [338] Kuduvalli GR. *Image Data Compression for High-resolution Digital Telera-diology*. PhD thesis, Department of Electrical and Computer Engineering, University of Calgary, Calgary, Alberta, Canada, May 1992.
- [339] Karssemeijer N. Detection of stellate distortions in mammograms using scale space operators. In Bizais Y, Barillot C, and Paola PD, editors, *Information Processing in Medical Imaging*, pages 335–346. Kluwer Academic, Dordrecht, The Netherlands, 1995.
- [340] Laine A, Huda W, Chen D, and Harris J. Segmentation of masses using continuous scale representations. In Doi K, Giger ML, Nishikawa RM, and Schmidt RA, editors, *3rd International Workshop on Digital Mammography*, pages 447–450, Chicago, IL, 9-12 June 1996.
- [341] Miller L and Ramsey N. The detection of malignant masses by non-linear multiscale analysis. In Doi K, Giger ML, Nishikawa RM, and Schmidt RA,

- editors, *3rd International Workshop on Digital Mammography*, pages 335–340, Chicago, IL, 9-12 June 1996.
- [342] Zhang M, Giger ML, Vyborny CJ, and Doi K. Mammographic texture analysis for the detection of spiculated lesions. In Doi K, Giger ML, Nishikawa RM, and Schmidt RA, editors, *3rd International Workshop on Digital Mammography*, pages 347–350, Chicago, IL, 9-12 June 1996.
- [343] Matsubara T, Fujita H, Endo T, Horita K, Ikeda M, Kido C, and Ishigaki T. Development of mass detection algorithm based on adaptive thresholding technique in digital mammograms. In Doi K, Giger ML, Nishikawa RM, and Schmidt RA, editors, *3rd International Workshop on Digital Mammography*, pages 391–396, Chicago, IL, 9-12 June 1996.
- [344] Kupinski MA and Giger ML. Automated seeded lesion segmentation on digital mammograms. *IEEE Transactions on Medical Imaging*, 17(4):510–517, 1998.
- [345] Rangayyan RM, Mudigonda NR, and Desautels JEL. Boundary modeling and shape analysis methods for classification of mammographic masses. *Medical and Biological Engineering and Computing*, 38:487–496, 2000.
- [346] Cannon RL, Dave JV, and Bezdek JC. Fuzzy C-Means clustering algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(2):248–255, 1986.
- [347] Clark MC, Hall DB, Goldgof DB, Clarke LP, Velthuisen RP, and Silbiger MS. MRI segmentation using fuzzy clustering techniques. *IEEE Engineering in Medicine and Biology*, pages 730–742, November/December 1994.
- [348] Chen CH and Lee GG. On digital mammogram segmentation and microcalcification detection using multiresolution wavelet analysis. *Graphical Models and Image Processing*, 59(5):349–364, 1997.
- [349] Sameti M and Ward RK. A fuzzy segmentation algorithm for mammogram partitioning. In K. Doi, M.L. Giger, R.M. Nishikawa, and R.A. Schmidt, editors, *3rd International Workshop on Digital Mammography*, pages 471–474, Chicago, IL, 9-12 June 1996.
- [350] Zadeh LA. Fuzzy sets. *Information and Control*, 8:338–353, 1965.
- [351] Klir GJ and Yuon B. *Fuzzy Sets and Fuzzy Logic*. Prentice Hall, Englewood Cliffs, NJ, 1995.
- [352] Guliato D. *Combinação de Algoritmos de Segmentação por Operadores de Agregação (in Portuguese)*. PhD thesis, Department of Electrical Engineering, University of São Paulo, São Paulo, São Paulo, Brazil, August 1998.
- [353] Guliato D, Rangayyan RM, Adorno F, and Ribeiro MMG. Analysis and classification of breast masses by fuzzy-set-based image processing. In Peitgen HO, editor, *6th International Workshop on Digital Mammography*, pages 196–197, Bremen, Germany, 22-25 June 2002.
- [354] Menut O, Rangayyan RM, and Desautels JEL. Parabolic modeling and classification of breast tumours. *International Journal of Shape Modeling*, 3(3 & 4):155–166, 1998.

- [355] Udupa JK and Samarasekera S. Fuzzy connectedness and object definition: Theory, algorithms, and applications in image segmentation. *Graphical Models and Image Processing*, 58(3):246–261, May 1996.
- [356] Saha PK, Udupa JK, Conant EF, Chakraborty DP, and Sullivan D. Breast tissue density quantification via digitized mammograms. *IEEE Transactions on Medical Imaging*, 20(8):792–803, 2001.
- [357] Saha PK, Udupa JK, and Odhner D. Scale-based fuzzy connected image segmentation: Theory, algorithms, and validation. *Computer Vision and Image Understanding*, 77:145–174, 2000.
- [358] Hough PVC. Method and means for recognizing complex patterns. US Patent 3,069,654, December 18, 1962.
- [359] Duda RO and Hart PE. Use of the Hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15:11–15, 1972.
- [360] Rangayyan RM and Krishnan S. Feature identification in the time-frequency plane by using the Hough-Radon transform. *Pattern Recognition*, 34:1147–1158, 2001.
- [361] Pavlidis T and Horowitz SL. Segmentation of plane curves. *IEEE Transactions on Computers*, C-23:860–870, August 1974.
- [362] Kass M, Witkin A, and Terzopoulos D. Snakes: active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988.
- [363] Falcão AX, Udupa JK, Samarasekera, Sharma S, Hirsch BE, and Lotufo RA. User-steered image segmentation paradigms: Live wire and live lane. *Models and Image Processing*, 60:233–260, 1998.
- [364] Falcão AX, Udupa JK, and Miyazawa FK. An ultra-fast user-steered image segmentation paradigm: Live wire on the fly. *IEEE Transactions on Medical Imaging*, 19(1):55–62, 2000.
- [365] Deglint HJ, Rangayyan RM, and Boag GS. Three-dimensional segmentation of the tumor mass in computed tomographic images of neuroblastoma. In Fitzpatrick JM and Sonka M, editors, *Proceedings of SPIE Medical Imaging 2004: Image Processing*, volume 5370, pages 475–483, San Diego, CA, February 2004.
- [366] Deglint HJ. Image processing algorithms for three-dimensional segmentation of the tumor mass in computed tomographic images of neuroblastoma. Master’s thesis, Department of Electrical and Computer Engineering, University of Calgary, Calgary, Alberta, Canada, August 2004.
- [367] Lou SL, Lin HD, Lin KP, and Hoogstrate D. Automatic breast region extraction from digital mammograms for PACS and telemammography applications. *Computerized Medical Imaging and Graphics*, 24:205–220, 2000.
- [368] Bick U, Giger ML, Schmidt RA, Nishikawa RM, and Doi K. Density correction of peripheral breast tissue on digital mammograms. *RadioGraphics*, 16(6):1403–1411, November 1996.
- [369] Byng JW, Critten JP, and Yaffe MJ. Thickness-equalization processing for mammographic images. *Radiology*, 203(2):564–568, 1997.

- [370] Chandrasekhar R and Attikiouzel Y. A simple method for automatically locating the nipple on mammograms. *IEEE Transactions on Medical Imaging*, 16(5):483–494, 1997.
- [371] Lau TK and Bischof WF. Automated detection of breast tumors using the asymmetry approach. *Computers and Biomedical Research*, 24:273–295, 1991.
- [372] Miller P and Astley S. Automated detection of mammographic asymmetry using anatomical features. *International Journal of Pattern Recognition and Artificial Intelligence*, 7(6):1461–1476, 1993.
- [373] Méndez AJ, Tahoces PG, Lado MJ, Souto M, Correa JL, and Vidal JJ. Automatic detection of breast border and nipple in digital mammograms. *Computer Methods and Programs in Biomedicine*, 49:253–262, 1996.
- [374] Bick U, Giger ML, Schmidt RA, Nishikawa RM, Wolverton DE, and Doi K. Automated segmentation of digitized mammograms. *Academic Radiology*, 2(1):1–9, 1995.
- [375] Ferrari RJ, Rangayyan RM, Desautels JEL, and Frère AF. Segmentation of mammograms: Identification of the skin-air boundary, pectoral muscle, and fibro-glandular disc. In Yaffe M J, editor, *Proceedings of the 5th International Workshop on Digital Mammography*, pages 573–579, Toronto, ON, Canada, June 2000.
- [376] Suckling J, Parker J, Dance DR, Astley S, Hutt I, Boggis CRM, Ricketts I, Stamatakis E, Cerneaz N, Kok SL, Taylor P, Betal D, and Savage J. The Mammographic Image Analysis Society digital mammogram database. In Gale AG, Astley SM, Dance DR, and Cairns AY, editors, *Proceedings of the 2nd International Workshop on Digital Mammography*, volume 1069 of *Excerpta Medica International Congress Series*, pages 375–378, York, England, July 1994.
- [377] Mackiewicz B. Intracranial boundary detection and radio frequency correction in magnetic resonance images. Master's thesis, School of Computing Science - Simon Fraser University, Burnaby, BC, Canada, August 1995.
- [378] Lobregt S and Viergever MA. A discrete dynamic contour model. *IEEE Transactions on Medical Imaging*, 14(1):12–24, 1995.
- [379] Williams DJ and Shah M. A fast algorithm for active contours and curvature estimation. *Computer Vision, Graphics, and Image Processing: Image Understanding*, 55(1):14–26, 1992.
- [380] Mattis P and Kimball S. GIMP — GNU Image Manipulation Program — version 1.1.17. <http://www.gimp.org> — GNU General Public License — GPL, accessed May 2002.
- [381] Ferrari RJ, Rangayyan RM, Desautels JEL, and Frère AF. Analysis of asymmetry in mammograms via directional filtering with Gabor wavelets. *IEEE Transactions on Medical Imaging*, 20(9):953–964, 2001.
- [382] Karssemeijer N. Automated classification of parenchymal patterns in mammograms. *Physics in Medicine and Biology*, 43(2):365–378, 1998.
- [383] Aylward SR, Hemminger BH, and Pisano ED. Mixture modeling for digital mammogram display and analysis. In Karssemeijer N, Thijssen M, Hendriks

- J, and van Erning L, editors, *Proceedings of the 4th International Workshop on Digital Mammography*, pages 305–312, Nijmegen, The Netherlands, June 1998.
- [384] Manjunath BS and Ma WY. Texture features for browsing and retrieval of image data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(8):837–842, 1996.
- [385] Lee TS. Image representation using 2–D Gabor wavelets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(10):959–971, 1996.
- [386] Mallat S. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7):674–693, 1989.
- [387] Jähne B. *Digital Image Processing*. Springer, San Diego, CA, 4th edition, 1997.
- [388] De Valois RL, Albrecht DG, and Thorell LG. Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research*, 22:545–559, 1982.
- [389] Daugman JG. Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *Journal of the Optical Society of America*, 2(7):1160–1169, 1985.
- [390] Jones P and Palmer LA. An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, 58(6):1233–1258, 1987.
- [391] Tucker AK. *Textbook of mammography*. Churchill Livingstone, New York, NY, 1993.
- [392] Ma WY and Manjunath BS. EdgeFlow: A technique for boundary detection and image segmentation. *IEEE Transactions on Image Processing*, 9(8):1375–1388, 2000.
- [393] Meyer F and Beucher S. Morphological segmentation. *Journal of Visual Communication and Image Representation*, 1(1):21–46, 1990.
- [394] Asar H, Nandhakumar N, and Aggarwal JK. Pyramid-based image segmentation using multisensory data. *Pattern Recognition*, 23(6):583–593, 1990.
- [395] Vincent L and Soille P. Watershed in digital spaces: An efficient algorithm based on immersion simulations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(6):583–598, 1991.
- [396] Xiaohan Y and Yla-Jaaski J. Direct segmentation in 3D and its application to medical images. In *Proceedings of SPIE: Image Processing*, volume 1898, pages 187–192, 1993.
- [397] Hadjarian A, Bala J, Gutta S, Trachiots S, and Pachowicz P. The fusion of supervised and unsupervised techniques for segmentation of abnormal regions. In Karssemeijer N, Thijssen M, Hendriks J, and van Erning L, editors, *4th International Workshop on Digital Mammography*, pages 299–302, Nijmegen, The Netherlands, June 1998. Kluwer Academic Publishers.
- [398] Yager RR. Connectives and quantifiers in fuzzy sets. *Fuzzy Sets and Systems*, 40:39–75, 1991.

- [399] Bloch I. Information combination operators for data fusion: a comparative review with classification. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 26(1):52–67, 1996.
- [400] Sahiner B, Petrick N, Chan HP, Hadjiiski LM, Paramagul C, Helvie MA, and Gurcan MN. Computer-aided characterization of mammographic masses: Accuracy of mass segmentation and its effects on characterization. *IEEE Transactions on Medical Imaging*, 20(12):1275–1284, 2001.
- [401] Gonzalez RC and Thomason MG. *Syntactic Pattern Recognition: An Introduction*. Addison-Wesley, Reading, MA, 1978.
- [402] Duda RO, Hart PE, and Stork DG. *Pattern Classification*. Wiley, New York, NY, 2nd edition, 2001.
- [403] American College of Radiology, Reston, VA. *Illustrated Breast Imaging Reporting and Data System (BI-RADSTM)*, 3rd edition, 1998.
- [404] van Otterloo PJ. *A Contour-oriented Approach to Shape Analysis*. Prentice Hall, New York, NY, 1991.
- [405] Bookstein FL. *The Measurement of Biological Shape and Shape Change*. Springer-Verlag, New York, NY, 1978.
- [406] Loncaric S. A survey of shape analysis techniques. *Pattern Recognition*, 31(8):983–1001, 1998.
- [407] Pohlman S, Powell KA, Obuchowski NA, Chilcote WA, and Grundfest-Broniatowski S. Quantitative classification of breast tumors in digitized mammograms. *Medical Physics*, 23(8):1337–1345, 1996.
- [408] Freeman H. On the encoding of arbitrary geometric configurations. *IRE Transactions on Electronic Computers*, EC-10:260–268, 1961.
- [409] Pavlidis T and Ali F. Computer recognition of handwritten numerals by polygonal approximations. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-5:610–614, November 1975.
- [410] Ventura JA and Chen JM. Segmentation of two-dimensional curve contours. *Pattern Recognition*, 25(10):1129–1140, 1992.
- [411] Suen CY and Zhang TY. A fast parallel algorithm for thinning digital patterns. *Communications of the ACM: Image Processing and Computer Vision*, 27(5):236–239, 1984.
- [412] Blum H. A transformation for extracting new descriptors of shape. In Wathen-Dunn W, editor, *Models for the Perception of Speech and Visual Form*. MIT Press, Cambridge, MA, 1967.
- [413] Lu HE and Wang PSP. A comment on ‘A fast parallel algorithm for thinning digital patterns’. *Communications of the ACM: Image Processing and Computer Vision*, 29(3):239–242, 1986.
- [414] Eng K, Rangayyan RM, Bray RC, Frank CB, Anscomb L, and Veale P. Quantitative analysis of the fine vascular anatomy of articular ligaments. *IEEE Transactions on Biomedical Engineering*, 39(3):296–306, 1992.
- [415] Bray RC, Rangayyan RM, and Frank CB. Normal and healing ligament vascularity: a quantitative histological assessment in the adult rabbit medial collateral ligament. *Journal of Anatomy*, 188:87–95, 1996.

- [416] Shen L. Shape analysis of mammographic calcifications. Master's thesis, Department of Electrical and Computer Engineering, University of Calgary, Calgary, Alberta, Canada, July 1992.
- [417] Gupta L and Srinath MD. Contour sequence moments for the classification of closed planar shapes. *Pattern Recognition*, 20(3):267–272, 1987.
- [418] Dudani SA, Breeding KJ, and McGhee RB. Aircraft identification by moment invariants. *IEEE Transactions on Computer*, C-26(1):39–45, 1983.
- [419] Reeves AP, Prokop RJ, Andrews SE, and Kuhl F. Three dimensional shape analysis using moments and Fourier descriptors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(6):937–943, 1988.
- [420] Hu MK. Visual pattern recognition by moment invariants. *IRE Transactions on Information Theory*, IT-8(2):179–187, 1962.
- [421] You Z and Jain AK. Performance evaluation of shape matching via chord length distribution. *Computer Vision, Graphics, and Image Processing*, 28:185–198, 1984.
- [422] Granlund GH. Fourier preprocessing for hand print character recognition. *IEEE Transactions on Computers*, C-21:195–201, February 1972.
- [423] Persoon E and Fu KS. Shape discrimination using Fourier descriptors. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-7(3):170–179, 1977.
- [424] Zahn CT and Roskies RZ. Fourier descriptors for plane closed curves. *IEEE Transactions on Computers*, C-21:269–281, March 1972.
- [425] Lestrel PE, editor. *Fourier Descriptors and their Applications in Biology*. Cambridge University Press, Cambridge, UK, 1997.
- [426] Kuhl FP and Giardina CR. Elliptic Fourier features of a closed contour. *Computer Graphics and Image Processing*, 18:236–258, 1982.
- [427] Lin CC and Chellappa R. Classification of partial 2-D shapes using Fourier descriptors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-9:686–690, 1987.
- [428] Sahiner BS, Chan HP, Petrick N, Helvie MA, and Hadjiiski LM. Improvement of mammographic mass characterization using spiculation measures and morphological features. *Medical Physics*, 28(7):1455–1465, 2001.
- [429] Lee TK, McLean DI, and Atkins MS. Irregularity index: A new border irregularity measure for cutaneous melanocytic lesions. *Medical Image Analysis*, 7:47–64, 2003.
- [430] Feig SA, Galkin BM, and Muir HD. Evaluation of breast microcalcifications by means of optically magnified tissue specimen radiographs. In Brüner S and Langfeldt B, editors, *Recent Results in Cancer Research*, volume 105, pages 111–123. Springer-Verlag, Berlin, Germany, 1987.
- [431] Sickles EA. Breast calcifications: Mammographic evaluation. *Radiology*, 160:289–293, 1986.
- [432] Rao AR. *A Taxonomy for Texture Description and Identification*. Springer-Verlag, New York, NY, 1990.
- [433] Sutton RN and Hall EL. Texture measures for automatic classification of pulmonary disease. *IEEE Transactions on Computers*, 21(7):667–676, 1972.

- [434] Paget RD and Longstaff D. Terrain mapping of radar satellite. *Journal of Electronic Imaging*, 6(2):6–7, 1996.
- [435] Ojala T, Pietikäinen M, and Nisula J. Determining composition of grain mixtures by texture classification based on feature distribution. *International Journal of Pattern Recognition and Artificial Intelligence*, 10(1):73–81, 1996.
- [436] Swarnakar V, Acharya RS, Sibata C, and Shin K. Fractal-based characterization of structural changes in biomedical images. In *Proceedings of SPIE*, volume 2709, pages 444–455, 1996.
- [437] Uppaluri R, Mitsa T, Hoffman EA, McLennan G, and Sonka M. Texture analysis of pulmonary parenchyma in normal and emphysematous lung. In *Proceedings of SPIE*, volume 2709, pages 456–467, 1996.
- [438] Julesz B and Bergen JR. Textons, the fundamental elements in preattentive vision and perception of textures. *The Bell System Technical Journal*, 62(6):1619–1645, 1983.
- [439] Wechsler H. Texture analysis – A survey. *Signal Processing*, 2:271–282, 1980.
- [440] Haralick RM and Shapiro LG. *Computer and Robot Vision*. Addison-Wesley, Reading, MA, 1992.
- [441] Haralick RM. Statistical and structural approaches to texture. *Proceedings of the IEEE*, 67(5):786–804, 1979.
- [442] Haralick RM, Shanmugam K, and Dinstein I. Textural features for image classification. *IEEE Transactions on Systems, Man, Cybernetics*, SMC-3(6):610–622, 1973.
- [443] van Wijk JJ. Spot noise. *Computer Graphics*, 25(4):309–318, 1991.
- [444] Martins ACG and Rangayyan RM. Texture element extraction via cepstral filtering in the Radon domain. *IETE Journal of Research (India)*, 48(3,4):143–150, 2002.
- [445] Lerski RA, Straughan K, Schad LR, Boyce D, Blüml S, and Zuna I. MR image texture analysis – An approach to tissue characterization. *Magnetic Resonance Imaging*, 11:873–887, 1993.
- [446] Petrosian A, Chan H, Helvie MA, Goodsitt MM, and Adler DD. Computer-aided diagnosis in mammography: Classification of mass and normal tissue by texture analysis. *Physics in Medicine and Biology*, 39:2273–2288, 1994.
- [447] Martins ACG, Rangayyan RM, and Ruschioni RA. Audification and sonification of texture in images. *Journal of Electronic Imaging*, 10(3):690–705, 2001.
- [448] Byng JW, Boyd NF, Fishell E, Jong RA, and Yaffe MJ. Automated analysis of mammographic densities. *Physics in Medicine and Biology*, 41:909–923, 1996.
- [449] Parkkinen J, Selkänaho M, and Oja E. Detecting texture periodicity from the cooccurrence matrix. *Pattern Recognition Letters*, 11:43–50, January 1990.
- [450] Chan HP, Wei D, Helvie MA, Sahiner B, Adler DD, Goodsitt MM, and Petrick N. Computer-aided classification of mammographic masses and nor-

- mal tissue: linear discriminant analysis in texture feature space. *Physics in Medicine and Biology*, 40(5):857–876, 1995.
- [451] Sahiner BS, Chan HP, Petrick N, Helvie MA, and Goodsitt MM. Computerized characterization of masses on mammograms: The rubber band straightening transform and texture analysis. *Medical Physics*, 25(4):516–526, 1998.
- [452] Laws KI. Rapid texture identification. In *Proceedings of SPIE Vol. 238: Image Processing for Missile Guidance*, pages 376–380, 1980.
- [453] Pietikäinen M, Rosenfeld A, and Davis LS. Experiments with texture classification using averages of local pattern matches. *IEEE Transactions on Systems, Man, and Cybernetics*, 13:421–426, 1983.
- [454] Hsiao JY and Sawchuk AA. Supervised textured image segmentation using feature smoothing and probabilistic relaxation techniques. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(12):1279–1292, 1989.
- [455] Miller P and Astley S. Classification of breast tissue by texture analysis. *Image and Vision Computing*, 10(5):277–282, 1992.
- [456] Goldberger AL, Rigney DR, and West BJ. Chaos and fractals in human physiology. *Scientific American*, 262:42–49, February 1990.
- [457] West BJ. Fractal forms in physiology. *International Journal of Modern Physics B*, 4(10):1629–1669, 1990.
- [458] Liu SH. Formation and anomalous properties of fractals. *IEEE Engineering in Medicine and Biology Magazine*, 11(2):28–39, June 1992.
- [459] Deering W and West BJ. Fractal physiology. *IEEE Engineering in Medicine and Biology Magazine*, 11(2):40–46, June 1992.
- [460] Barnsley M. *Fractals Everywhere*. Academic, San Diego, CA, 1988.
- [461] Peitgen HO and Saupe D, editors. *The Science of Fractal Images*. Springer-Verlag, New York, NY, 1988.
- [462] Mandelbrot BB. *Fractals*. WH Freeman and Company, San Francisco, CA, 1977.
- [463] Kantz H, Kurts J, and Mayer-Kress G, editors. *Nonlinear Analysis of Physiological Data*. Springer-Verlag, Berlin, Germany, 1998.
- [464] Peleg S, Naor J, Hartley R, and Avnir D. Multiple-resolution texture analysis and classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(4):518–523, 1984.
- [465] Pentland AP. Fractal-based description of natural scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(6):661–674, 1984.
- [466] Lundahl T, Ohley W, Kay SM, and Siffert R. Fractal Brownian motion: A maximum likelihood estimator and its application to image texture. *IEEE Transactions on Medical Imaging*, 5(3):152–161, 1986.
- [467] Schepers HE, van Beek JHGM, and Bassingthwaite JB. Four methods to estimate the fractal dimension from self-affine signals. *IEEE Engineering in Medicine and Biology Magazine*, 11(2):57–64, June 1992.

- [468] Fortin C, Kumaresan R, Ohley W, and Hoefler S. Fractal dimension in the analysis of medical images. *IEEE Engineering in Medicine and Biology Magazine*, 11(2):65–71, June 1992.
- [469] Geraets WGM and van der Stelt PF. Fractal properties of bone. *Dentomaxillofacial Radiology*, 29:144–153, 2000.
- [470] Chen CC, DaPonte JS, and Fox MD. Fractal feature analysis and classification in medical imaging. *IEEE Transactions on Medical Imaging*, 8(2):133–142, 1989.
- [471] Burdett CJ, Longbotham HG, Desai M, Richardson WB, and Stoll JF. Non-linear indicators of malignancy. In *Proceedings of SPIE Vol. 1905 on Biomedical Image Processing and Biomedical Visualization*, pages 853–860, San Jose, CA, Feb. 1993.
- [472] Wu CM, Chen YC, and Hsieh KS. Texture features for classification of ultrasonic liver images. *IEEE Transactions on Medical Imaging*, 11(2):141–152, 1992.
- [473] Lee WL, Chen YC, and Hsieh KS. Ultrasonic liver tissues classification by fractal feature vector based on M-band wavelet transform. *IEEE Transactions on Medical Imaging*, 22(3):382–392, 2003.
- [474] Yaffe MJ, Byng JW, and Boyd NF. Quantitative image analysis for estimation of breast cancer risk. In Bankman IN, editor, *Handbook of Medical Imaging: Processing and Analysis*, chapter 21, pages 323–340. Academic Press, London, UK, 2000.
- [475] Caldwell CB, Stapleton SJ, Holdsworth DW, Jong RA, Weiser WJ, Cooke G, and Yaffe MJ. Characterization of mammographic parenchymal pattern by fractal dimension. *Physics in Medicine and Biology*, 35(2):235–247, 1990.
- [476] Iftekharuddin KM, Jia W, and Marsh R. Fractal analysis of tumor in brain MR images. *Machine Vision and Applications*, 13:352–362, 2003.
- [477] Chaudhuri BB and Sarkar N. Texture segmentation using fractal dimension. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(1):72–77, 1995.
- [478] Zheng L and Chan AK. An artificial intelligent algorithm for tumor detection in screening mammogram. *IEEE Transactions on Medical Imaging*, 20(7):559–567, 2001.
- [479] Saparin PI, Gowin W, Kurths J, and Felsenberg D. Quantification cancellous bone structure using symbol dynamics and measures of complexity. *Physical Review E*, 58:6449–6459, 1998.
- [480] Jennane R, Ohley WJ, Majumdar S, and Lemineur G. Fractal analysis of bone X-ray tomographic microscopy projections. *IEEE Transactions on Medical Imaging*, 20(5):443–449, 2001.
- [481] Samarabandhu J, Acharya R, Hausmann E, and Allen K. Analysis of bone X-rays using morphological fractals. *IEEE Transactions on Medical Imaging*, 12(3):466–470, 1993.
- [482] Sedivy R, Windischberger Ch, Svozil K, Moser E, and Breiteneker G. Fractal analysis: An objective method for identifying atypical nuclei in dysplastic lesions of the cervix uteri. *Gynecologic Oncology*, 75:78–83, 1999.

- [483] Esgiar AN, Naguib RNG, Sharif BS, Bennett MK, and Murray A. Fractal analysis in the detection of colonic cancer images. *IEEE Transactions on Information Technology in Biomedicine*, 6(1):54–58, 2002.
- [484] Penn AI and Loew MH. Estimating fractal dimension with fractal interpolation function models. *IEEE Transactions on Medical Imaging*, 16(6):930–937, 1997.
- [485] Bankman IN, Spisz TS, and Pavlopoulos S. Two-dimensional shape and texture quantification. In Bankman IN, editor, *Handbook of Medical Imaging: Processing and Analysis*, chapter 14, pages 215–230. Academic Press, London, UK, 2000.
- [486] Jernigan ME and D’Astous F. Entropy-based texture analysis in the spatial frequency domain. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(2):237–243, 1984.
- [487] Liu S and Jernigan ME. Texture analysis and discrimination in additive noise. *Computer Vision, Graphics, and Image Processing*, 49:52–67, 1990.
- [488] Laine A and Fan J. Frame representations for texture segmentation. *IEEE Transactions on Image Processing*, 5(5):771–780, 1996.
- [489] McLean GF. Vector quantization for texture classification. *IEEE Transactions on Systems, Man, Cybernetics*, 23(3):637–649, 1993.
- [490] Bovik AC. Analysis of multichannel narrow-band filters for image texture segmentation. *IEEE Transactions on Signal Processing*, 39(9):2025–2043, 1991.
- [491] Vilnrotter FM, Nevatia R, and Price KE. Structural analysis of natural textures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(1):76–89, 1986.
- [492] He DC and Wang L. Textural filters based on the texture spectrum. *Pattern Recognition*, 24(12):1187–1195, 1991.
- [493] Wang S, Velasco FRD, Wu AY, and Rosenfeld A. Relative effectiveness of selected texture primitive statistics for texture discrimination. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-11(5):360–370, 1981.
- [494] Tomita F, Shirai Y, and Tsuji S. Description of textures by structural analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 4(2):183–191, 1982.
- [495] Turner MR. Texture discrimination by Gabor functions. *Biological Cybernetics*, 55:71–82, 1986.
- [496] Bovik AC, Clark M, and Geisler WS. Multichannel texture analysis using localized spatial filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(1):55–73, 1990.
- [497] Porat M and Zeevi YY. Localized texture processing in vision: Analysis and synthesis in the Gaborian space. *IEEE Transactions on Biomedical Engineering*, 36(1):115–129, 1989.
- [498] Reed TR and Wechsler H. Segmentation of textured images and Gestalt organization using spatial/ spatial-frequency representations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(1):1–12, 1990.

- [499] Reed TR, Wechsler H, and Werman M. Texture segmentation using a diffusion region growing technique. *Pattern Recognition*, 23(9):953–960, 1990.
- [500] Jain AK and Farrokhnia F. Unsupervised texture segmentation using Gabor filters. *Pattern Recognition*, 24(12):1167–1186, 1991.
- [501] Unser M. Texture classification and segmentation using wavelet frames. *IEEE Transactions on Image Processing*, 4:1549–1560, 1995.
- [502] Ravichandran G and Trivedi MM. Circular-Mellin features for texture segmentation. *IEEE Transactions on Image Processing*, 4:1629–1639, 1995.
- [503] Tardif P and Zaccarin A. Multiscale autoregressive image representation for texture segmentation. In *Proceedings of SPIE: Image Processing*, volume 3026, pages 327–337, 1997.
- [504] Unser M and Eden M. Multiresolution feature extraction and selection for texture segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7):717–728, 1989.
- [505] Martins ACG and Rangayyan RM. Complex cepstral filtering of images and echo removal in the Radon domain. *Pattern Recognition*, 30(11):1931–1938, 1997.
- [506] Chambers JM, Mathews MV, and Moore FR. *Auditory Data Inspection*. Report TM 74-122-2, Bell Laboratories, New York, NY, 1974.
- [507] Kramer G, editor. *Auditory Display: Sonification, Audification, and Auditory Interfaces*. Addison Wesley, Reading, MA, 1994.
- [508] Meijer P. An experimental system for auditory image representation. *IEEE Transactions on Biomedical Engineering*, 39(2):112–121, 1992.
- [509] Meijer PBL. *Let's Make Vision Accesible*. <http://www.visualprosthesis.com/voicover.htm>, accessed June 2004.
- [510] Makhoul J. Linear prediction: A tutorial. *Proceedings of the IEEE*, 63(4):561–580, 1975.
- [511] Martins ACG, Rangayyan RM, Portela LA, Amaro Jr. E, and Ruschioni RA. Auditory display and sonification of textured images. In *Proceedings of the Third International Conference on Auditory Display*, pages 9–11, Palo Alto, CA, Nov. 1996.
- [512] Kinoshita SK, de Azevedo Marques PM, Slaets AFF, Marana HRC, Ferrari RJ, and Villela RL. Detection and characterization of mammographic masses by artificial neural network. In Karssemeijer N, Thijssen M, Hendriks J, and van Erning L, editors, *Proceedings of the 4th International Workshop on Digital Mammography*, pages 489–490, Nijmegen, The Netherlands, June 1998.
- [513] Sahiner B, Chan HP, Petrick N, Wei D, Helvie MA, Adler DD, and Goodsitt MM. Classification of mass and normal breast tissue: a convolution neural network classifier with spatial domain and texture images. *IEEE Transactions on Medical Imaging*, 15(5):598–611, 1996.
- [514] Wei D, Chan HP, Helvie MA, Sahiner B, Petrick N, Adler DD, and Goodsitt MM. Classification of mass and normal breast tissue on digital mammograms: multiresolution texture analysis. *Medical Physics*, 22(9):1501–1513, 1995.

- [515] Wei D, Chan HP, Petrick N, Sahiner B, Helvie MA, Adler DD, and Goodsitt MM. False-positive reduction technique for detection of masses on digital mammograms: global and local multiresolution texture analysis. *Medical Physics*, 24(6):903–914, 1997.
- [516] Sahiner B, Chan HP, Petrick N, Helvie MA, and Goodsitt MM. Design of a high-sensitivity classifier based on a genetic algorithm: application to computer-aided diagnosis. *Physics in Medicine and Biology*, 43(10):2853–2871, 1998.
- [517] Kok SL, Brady JM, and Tarassenko L. The detection of abnormalities in mammograms. In Gale AG, Astley SM, Dance DR, and Cairns AY, editors, *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 261–270, York, England, 10-12 July 1994.
- [518] Huo Z, Giger ML, Vyborny CJ, Bick U, Lu P, Wolverton DE, and Schmidt RA. Analysis of spiculation in the computerised classification of mammographic masses. *Medical Physics*, 22(10):1569–1579, 1995.
- [519] Giger ML, Lu P, Huo Z, Bick U, Vyborny CJ, Schmidt RA, Zhang W, Metz CE, Wolverton D, Nishikawa RM, Zouras W, and Doi K. CAD in digital mammography: computerized detection and classification of masses. In Gale AG, Astley SM, Dance DR, and Cairns AY, editors, *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 281–288, York, England, 10-12 July 1994.
- [520] Huo Z, Giger ML, Vyborny CJ, Wolverton DE, Schmidt RA, and Doi K. Computer-aided diagnosis: Automated classification of mammographic mass lesions. In Doi K, Giger ML, Nishikawa RM, and Schmidt RA, editors, *Proceedings of the 3rd International Workshop on Digital Mammography*, pages 207–211, Chicago, IL, 9-12 June 1996.
- [521] Highnam RP, Brady JM, and Shepstone BJ. A quantitative feature to aid diagnosis in mammography. In Doi K, Giger ML, Nishikawa RM, and Schmidt RA, editors, *Proceedings of the 3rd International Workshop on Digital Mammography*, pages 201–206, Chicago, IL, 9-12 June 1996.
- [522] Claridge E and Richter JH. Characterisation of mammographic lesions. In Gale AG, Astley SM, Dance DR, and Cairns AY, editors, *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 241–250, York, England, 10-12 July 1994.
- [523] Sahiner B, Chan HP, Petrick N, Helvie MA, Adler DD, and Goodsitt MM. Classification of masses on mammograms using rubber-band straightening transform and feature analysis. In *Proceedings of SPIE vol. 2710 on Medical Imaging 1996 – Image Processing*, pages 44–50, Newport Beach, CA, 1996.
- [524] Hadjiiski L, Sahiner B, Chan HP, Petrick N, and Helvie MA. Classification of malignant and benign masses based on hybrid ART2LDA approach. *IEEE Transactions on Medical Imaging*, 18(12):1178–1187, 1999.
- [525] Giger ML, Huo Z, Wolverton DE, Vyborny CJ, Moran C, Schmidt RA, Alhallaq H, Nishikawa RM, and Doi K. Computer-aided diagnosis of digital mammographic and ultrasound images of breast mass lesions. In Karssemeijer N, Thijssen M, Hendriks J, and van Erning L, editors, *Proceedings*

- of the 4th International Workshop on Digital Mammography, pages 143–147, Nijmegen, The Netherlands, June 1998.
- [526] Gonzalez RC and Woods RE. *Digital Image Processing*. Addison-Wesley, Reading, MA, 1992.
- [527] Russ JC. *The Image Processing Handbook*. CRC Press, Boca Raton, FL, 1995.
- [528] Alto H. *Computer-aided Diagnosis of Breast Cancer*. PhD thesis, Department of Electrical and Computer Engineering, University of Calgary, Calgary, Alberta, Canada, April 2003.
- [529] Alto H, Rangayyan RM, and Desautels JEL. Content-based retrieval and analysis of mammographic masses. *Journal of Electronic Imaging*, 14:in press, 2005.
- [530] Ojala T, Pietikäinen M, and Harwood D. A comparative study of texture measures with classification based on feature distributions. *Pattern Recognition*, 29(1):51–59, 1996.
- [531] Liu ZQ, Rangayyan RM, and Frank CB. Directional analysis of images in scale-space. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(11):1185–1192, 1991.
- [532] Dziedzic-Goclawska A, Rozycka M, Czyba JC, Sawicki W, Moutier R, Lenczowski S, and Ostrowski K. Application of the optical Fourier transform for analysis of the spatial distribution of collagen fibers in normal and osteopetrotic bone tissue. *Histochemistry*, 74:123–137, 1982.
- [533] Komori M, Minato K, Nakano Y, Hirakawa Y, and Kuwahara M. Automatic measurement system for congenital hip dislocation using computed radiography. In *Proceedings of the SPIE, Volume 914: Medical Imaging II*, pages 665–668, 1988.
- [534] Kurogi S, Jianqiang Y, and Matsuoka K. Measurement of the angle of rotated images using Fourier transform. *Transactions of the Institute of Electronics, Information and Communication Engineers D-II*, J73D-II(4):590–596, April 1990.
- [535] Denslow S, Zhang Z, Thompson RP, and Lam CF. Statistically characterized features for directionality quantitation in patterns and textures. *Pattern Recognition*, 26:1193–1205, 1993.
- [536] Goncharov AB and Gelfand MS. Determination of mutual orientation of identical particles from their projections by the moments method. *Ultramicroscopy*, 25(4):317–328, 1988.
- [537] Abramson SB and Fay FS. Application of multiresolution spatial filters to long-axis tracking. *IEEE Transactions on Medical Imaging*, 9(2):151–158, 1990.
- [538] Kronick PL and Sacks MS. Quantification of vertical-fiber defect in cattle hide by small-angle light scattering. *Connective Tissue Research*, 27:1–13, 1991.
- [539] Sacks MS and Chuong CJ. Characterization of collagen fiber architecture in the canine diaphragmatic central tendon. *Journal of Biomechanical Engineering*, 114:183–190, 1992.

- [540] Petroll WM, Cavanagh HD, Barry P, Andrews P, and Jester JV. Quantitative analysis of stress fiber orientation during corneal wound contraction. *Journal of Cell Science*, 104:353–363, 1993.
- [541] Thackray BD and Nelson AC. Semi-automatic segmentation of vascular network images using a rotating structural element (ROSE) with mathematical morphology and dual feature thresholding. *IEEE Transactions on Medical Imaging*, 12(3):385–392, 1993.
- [542] Rolston WA. Directional image analysis. Master's thesis, Department of Electrical and Computer Engineering, University of Calgary, Calgary, Alberta, Canada, April 1994.
- [543] Rolston WA and Rangayyan RM. Directional analysis of images using multi-resolution Gabor filters. In *Proceedings of the International Conference on Robotics, Vision, and Parallel Processing for Industrial Vision*, pages 307–313, Ipoh, Malaysia, 26-28 May 1994.
- [544] Johnson RW. Characterization of fiber behavior in a nonwoven web through image analysis of tracer fibers. In *TAPPI Proceedings – 1988 Nonwovens Conference*, pages 217–221, Nashville, TN, April 1988. TAPPI Press.
- [545] Villa KM and Buchanan DR. Image analysis and the structure of nonwoven fabrics. In *INDA-TEC: The International Nonwovens Technological Conference*, pages 83–101, Philadelphia, PA, June 1986. Association of the Nonwoven Fabrics Industry, New York, NY.
- [546] Haley CS and Landoll LM. Image analysis of real and simulated nonwoven fabrics. In *INDA-TEC: The International Nonwovens Technological Conference*, pages 65–82, Philadelphia, PA, June 1986.
- [547] Yuhara T, Hasuike M, and Murakami K. Fibre orientation measurement with the two-dimensional power spectrum of a high-resolution soft x-ray image. *Journal of Pulp and Paper Science*, 17(4):J110–J114, 1991.
- [548] Yang CF, Crosby CM, Eusufzai ARK, and Mark RE. Determination of paper sheet fiber orientation by a laser optical diffraction method. *Journal of Applied Polymer Science*, 34:1145–1157, 1987.
- [549] Bresee RR and Donelson DS. Small-angle light scattering for analysis of a single fiber. *Journal of Forensic Sciences*, 25(2):413–422, 1980.
- [550] Embree P and Burg JP. Wide-band velocity filtering — the pie slice process. *Geophysics*, 28:948–974, 1963.
- [551] Treitel S, Shanks JL, and Frasier CW. Some aspects of fan filtering. *Geophysics*, 32:789–806, 1967.
- [552] Bezdova V, Ježek J, and Segeth K. FREDPACK- A program package for linear filtering in the frequency domain. *Computers & Geosciences*, 16(8):1123–1154, 1990.
- [553] Thorarinnsson F, Magnusson SG, and Bjornsson A. Directional spectral analysis and filtering of geophysical maps. *Geophysics*, 53(12):1587–1591, 1988.
- [554] Tashiro H, Yoshikawa K, Nomura T, and Hamabe A. Measurement of position and posture using image processing by projective pattern on an object. *Journal of the Japan Society of Precision Engineering*, 56(7):1286–1291, 1990.

- [555] Kono H. Measurement of angle and side length on the rectangle using directional code. In *Proceedings of the 10th International Conference on Assembly Automation*, pages 413–420, 1989.
- [556] Marra M, Dunlay R, and Mathis D. Terrain classification using texture for the ALV. In *Proceedings of SPIE*, volume 1289, pages 64–70, 1989.
- [557] Jacobberger PA. Mapping abandoned river channels in Mali through directional filtering of thematic mapper data. *Remote Sensing of Environment*, 26(2):161–170, 1988.
- [558] Arsenault HH, Sequin MK, and Brousseau N. Optical filtering of aeromagnetic maps. *Applied Optics*, 13:1013–1017, May 1974.
- [559] Moore GK and Waltz FA. Objective procedures for lineament enhancement and extraction. *Photogrammetric Engineering and Remote Sensing*, 49(5):641–647, 1983.
- [560] Duvernoy J and Chalasinska-Macukow K. Processing measurements of the directional content of Fourier spectra. *Applied Optics*, 20(1):136–144, 1981.
- [561] Duggin M, Rowntree RA, and Odell AW. Application of spatial filtering methods to urban feature analysis using digital image data. *International Journal of Remote Sensing*, 9(3):543–553, 1988.
- [562] Carrere V. Development of multiple source data processing for structural analysis at a regional scale. *Photogrammetric Engineering and Remote Sensing*, 56(5):587–595, 1990.
- [563] Shlomot E, Zeevi Y, and Pearlman WA. The importance of spatial frequency and orientation in image decomposition and coding. In *Proceedings of SPIE, Volume 845: Visual Communication and Image Processing*, pages 152–158, 1987.
- [564] Li H and He Z. Directional subband coding of images. In *International Conference on Acoustics, Speech, and Signal Processing*, volume III, pages 1823–1826, Glasgow, Scotland, May 1989.
- [565] Ikonomopoulos A and Kunt M. Directional filtering, zero crossing, edge detection and image coding. In Schussler HW, editor, *Signal Processing II: Theories and Applications*, pages 203–206. Elsevier, New York, NY, 1983.
- [566] Kunt M, Ikonomopoulos A, and Kocher M. Second-generation image-coding techniques. *Proceedings of the IEEE*, 73(4):549–574, 1985.
- [567] Ikonomopoulos A and Kunt M. High compression image coding via directional filtering. *Signal Processing*, 8:179–203, 1985.
- [568] Kunt M. Recent results in high-compression image coding. *IEEE Transactions on Circuits and Systems*, CAS-34:1306–1336, November 1987.
- [569] Hou HS and Vogel MJ. Detection of oriented line segments using discrete cosine transform. In *Intelligent Robots and Computer Vision: Seventh in a Series*, volume 1002, pages 81–87. Proceedings of SPIE, 1988.
- [570] Mardia KV. *Statistics of Directional Data*. Academic Press, New York, NY, 1972.

- [571] Schiller P, Finlay B, and Volman S. Quantitative studies of single-cell properties of monkey striate cortex. I. Spatiotemporal organization of receptive fields. *Journal of Neurophysiology*, 6:1288–1319, 1976.
- [572] Kass M and Witkin A. Analyzing oriented patterns. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, pages 944–952, Los Angeles, CA, 1985.
- [573] Zucker SW. Early orientation selection: Tangent fields and the dimensionality of their support. *Computer Vision, Graphics, and Image Processing*, 32:74–103, 1985.
- [574] Low KC and Coggins JM. Multiscale vector fields for image pattern recognition. In *Proceedings of SPIE, Volume 1192, Intelligent Robots and Computer Vision VIII: Algorithms and Techniques*, pages 159–168, 1989.
- [575] Allen T, Mead C, Faggin F, and Gribble G. Orientation-selective VLSI retina. In *Proceedings of SPIE, Volume 1001: Visual Communications and Image Processing*, pages 1040–1046, 1988.
- [576] Bigün J, Grandlund GH, and Wiklund J. Multidimensional orientation estimation with applications to texture analysis and optical flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(8):775–790, 1991.
- [577] Bruton LT, Bartley NR, and Stein RA. The design of stable high-quality two-dimensional recursive filters for seismic signal processing. *Advances in Geophysical Data Processing*, 2:233–261, 1985.
- [578] Bamberger RH and Smith MJT. A filter bank for the directional decomposition of images: Theory and design. *IEEE Transactions on Signal Processing*, 40(4):882–893, 1992.
- [579] Ikonomopoulos A and Unser M. A directional filtering approach to texture discrimination. In *Proceedings of the 7th International Conference of Pattern Recognition*, volume 1, pages 87–89, Montréal, Québec, Canada, 1984.
- [580] Wang TX. Three-dimensional filtering using Hilbert transforms. *Chinese Science Bulletin*, 35(2):123–127, January 1990.
- [581] Bonnet C, Brettel H, and Cohen I. Visibility of the spatial frequency components predicts the perceived orientational structure of a visual pattern. In *Proceedings of SPIE, Volume 1077: Human Vision, Visual Processing, and Digital Display*, pages 277–284, 1989.
- [582] Porat B and Friedlander B. A frequency domain algorithm for multiframe detection and estimation of dim targets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(4):398–401, 1990.
- [583] O’Gorman L and Nickerson JV. Matched filter design for fingerprint image enhancement. In *IEEE International Conference on Acoustics, Speech, and Signal Processing ’88*, pages 916–919, New York, NY, April 1988.
- [584] Fowlow TJ and Bruton LT. Attenuation characteristics of three-dimensional planar-resonant recursive digital filters. *IEEE Transactions on Circuits and Systems*, 35(5):595–599, 1988.
- [585] Freeman WT and Adelson EH. The design and use of steerable filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(9):891–906, 1991.

- [586] Marquerre H. Optical preparation of image data with directional filter for automatic inspection. *Optik*, 79(2):47–52, March 1988.
- [587] Bruton LT and Bartley NR. Using nonessential singularities of the second kind in two-dimensional filter design. *IEEE Transactions on Circuits and Systems*, 36:113–116, 1989.
- [588] Goodman D. Some difficulties with the double bilinear transformation in 2-D recursive filter design. *Proceedings of the IEEE*, 66:796–797, 1978.
- [589] Gonzalez RC and Wintz P. *Digital Image Processing*. Addison-Wesley, Reading, MA, 2nd edition, 1992.
- [590] Ning SX, Fan YP, and Tong C. A new smooth filter for directional detection and enhancement. In *9th International Conference on Pattern Recognition*, pages 628–630, Rome, Italy, 1988. IEEE Computer Society Press.
- [591] Otsu N. A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(1):62–66, 1979.
- [592] Illingworth J and Kittler J. A parallel threshold selection algorithm. In *Proceedings of the SPIE, Volume 596: Architectures and Algorithms for Digital Image Processing*, pages 129–133, 1985.
- [593] Clark M, Bovik AC, and Geisler WS. Texture segmentation using Gabor modulation/ demodulation. *Pattern Recognition Letters*, 6(4):261–267, 1987.
- [594] Bovik AC, Clark M, and Geisler WS. Computational texture analysis using localized spatial filtering. In *Proceedings of the IEEE Computer Society Workshop on Computer Vision*, pages 201–206, Miami Beach, FL, November 1987. IEEE.
- [595] Ayres FJ and Rangayyan RM. Characterization of architectural distortion in mammograms. In *CDROM Proceedings of the 25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Cancún, Mexico, September 2003.
- [596] Gabor D. Theory of communication. *Journal of the Institute of Electrical Engineers*, 93:429–457, 1946.
- [597] Zhou YT, Venkateswar V, and Chellappa R. Edge detection and linear feature extraction using a 2-D random field model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(1):84–95, 1989.
- [598] Chen J, Sato Y, and Tamura S. Orientation space filtering for multiple orientation line segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(5):417–429, 2000.
- [599] Rangayyan RM and Rolston A. Directional image analysis with the Hough and Radon transforms. *Journal of the Indian Institute of Science*, 78:3–16, 1998.
- [600] Deans SR. Hough transform from the Radon transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 3:185–188, 1981.
- [601] Leavers VF and Boyce JF. The Radon transform and its application to shape parameterization in machine vision. *Image and Vision Computing*, 5:161–166, 1987.

- [602] Nimni M. Collagen: Its structure and function in normal and pathological connective tissue. *Seminars in Arthritis and Rheumatology*, 4:95–150, 1974.
- [603] Butler DL, Zernicke RF, Grood ES, and Noyes FR. Biomechanics of ligaments and tendons. In Hutton R, editor, *Exercise and Sports Science Review*, pages 125–182. Franklin Institute Press, Hillsdale, NJ, 1978.
- [604] Frank C, Amiel D, Woo SLY, and Akeson W. Normal ligament properties and ligament healing. *Clinical Orthopaedics*, 196:15–25, 1985.
- [605] Forrester JC, Zederfeldt BH, Hayes TL, and Hunt TK. Tape-closed and sutured wounds: A comparison by tensiometry and scanning electron microscopy. *British Journal of Surgery*, 57:729–737, 1970.
- [606] Oegema T, An K, Weiland A, and Furcht L. Injury and repair of the musculoskeletal soft tissues. In Woo SLY and Buckwalter JA, editors, *American Academy of Orthopaedic Surgeons Symposium*, page 355. C.V. Mosby, St. Louis, MO, 1988.
- [607] Arnoczky SP, Marshall JL, and Rubin RM. Microvasculature of the cruciate ligaments and its response to injury. *Journal of Bone Joint Surgery*, 61A:1221–1229, 1979.
- [608] Arnoczky SP, Marshall JL, and Tarvin GB. Anterior cruciate ligament replacement using patellar tendon – An evaluation of graft revascularization in the dog. *Journal of Bone Joint Surgery*, 64A:217–224, 1982.
- [609] Bray RC, Fisher AWF, and Frank CB. Fine vascular anatomy of adult rabbit knee ligaments. *Journal of Anatomy*, 172:69–79, 1990.
- [610] Chaudhuri S. Digital image processing techniques for quantitative analysis of collagen fibril alignment in ligaments. Master's thesis, Department of Electrical and Computer Engineering, University of Calgary, Calgary, Alberta, Canada, 1987.
- [611] Rosenfeld A and de la Torre P. Histogram concavity as an aid to threshold selection. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-13:231–235, 1983.
- [612] Walpole RE and Myers RH, editors. *Probability and Statistics for Engineers and Scientists*. Macmillan, New York, NY, 1985.
- [613] Chimich DD, Bray RC, Frank CB, and Shrive NG. Contralateral knee ligaments may not be 'normal' after opposite knee surgery: A biomechanical study in the adult rabbit MCL complex. In *Transactions of the Canadian Orthopaedic Research Society 24th Annual Meeting*, Vancouver, BC, Canada, June 1990.
- [614] Winsberg F, Elkin M, Macy JJ, Bordaz V, and Weymouth W. Detection of radiographic abnormalities in mammograms by means of optical scanning and computer analysis. *Radiology*, 89:211–215, 1967.
- [615] Ackerman LV and Gose E. Breast lesion classification by computer and xeroradiograph. *Cancer*, 30:1025–1035, 1972.
- [616] Kimme C, O'Loughlin BJ, and Sklansky J. Automatic detection of suspicious abnormalities in breast radiographs. In Klinger A, Fu KS, and Kunii TL, editors, *Data Structures, Computer Graphics, and Pattern Recognition*, pages 427–447. Academic Press, New York, NY, 1977.

- [617] Hand W, Semmlow JL, Ackerman LV, and Alcorn FS. Computer screening of xeromammograms: A technique for defining suspicious areas of the breast. *Computers and Biomedical Research*, 12:445–460, 1979.
- [618] Semmlow JL, Shadagoppan A, Ackerman LV, Hand W, and Alcorn FS. A fully automated system for screening xeromammograms. *Computers and Biomedical Research*, 13:350–362, 1980.
- [619] Lai SM, Li XB, and Bischof WF. On techniques for detecting circumscribed masses in mammograms. *IEEE Transactions on Medical Imaging*, 8(4):377–386, 1989.
- [620] Brzakovic D, Luo XM, and Brzakovic P. An approach to automated detection of tumours in mammograms. *IEEE Transactions on Medical Imaging*, 9(3):233–241, 1990.
- [621] Barman H and Granlund GH. Computer aided diagnosis of mammograms using a hierarchical framework. In Gale AG, Astley SM, Dance DR, and Cairns AY, editors, *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 271–280, York, England, 10-12 July 1994.
- [622] Li HD, Kallergi M, Clarke LP, Jain VK, and Clark RA. Markov random field for tumor detection in digital mammography. *IEEE Transactions on Medical Imaging*, 14(3):565–576, 1995.
- [623] Qian W, Li L, Clarke LP, Mao F, and Clark RA. Adaptive CAD modules for mass detection in digital mammography. In Chang HK and Zhang YT, editors, *Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 1013–1016, Hong Kong, October 1998.
- [624] Chang CM and Laine A. Coherence of multiscale features for enhancement of digital mammograms. *IEEE Transactions on Information Technology in Biomedicine*, 3(1):32–46, 1999.
- [625] Woods KS and Bowyer KW. Computer detection of stellate lesions. In Gale AG, Astley SM, Dance DR, and Cairns AY, editors, *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 221–230, York, England, 10-12 July 1994.
- [626] Cerneaz NJ. *Model-based Analysis of Mammograms*. PhD thesis, Department of Engineering Science, University of Oxford, Oxford, England, 1994.
- [627] Woods KS and Bowyer KW. A general view of detection algorithms. In Doi K, Giger ML, Nishikawa RM, and Schmidt RA, editors, *Proceedings of the 3rd International Workshop on Digital Mammography*, pages 385–390, Chicago, IL, 9-12 June 1996.
- [628] Kok SL, Brady M, and Highnam R. Comparing mammogram pairs for the detection of lesions. In Karssemeijer N, Thijssen M, Hendriks J, and van Erning L, editors, *Proceedings of the 4th International Workshop on Digital Mammography*, pages 103–110, Nijmegen, The Netherlands, June 1998.
- [629] Guissin R and Brady JM. *Iso-intensity contours for edge detection. Technical report OUEL 1935/92*. Department of Engineering Science, Oxford University, Oxford, England, 1992.

- [630] Cerneaz N and Brady M. Enriching digital mammogram image analysis with a description of the curvi-linear structures. In Gale AG, Astley SM, Dance DR, and Cairns AY, editors, *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 297–306, York, England, 10-12 July 1994.
- [631] Lindeberg T. Detecting salient blob-like image structures and their scales with a scale-space primal sketch: a method for focus-of-attention. *International Journal of Computer Vision*, 11(3):283–318, 1993.
- [632] Petrick N, Chan HP, Sahiner B, and Wei D. An adaptive density-weighted contrast enhancement filter for mammographic breast mass detection. *IEEE Transactions on Medical Imaging*, 15(1):59–67, 1996.
- [633] Petrick N, Chan HP, Wei D, Sahiner B, Helvie MA, and Adler DD. Automated detection of breast masses on mammograms using adaptive contrast enhancement and texture classification. *Medical Physics*, 23(10):1685–1696, 1996.
- [634] Kobatake H, Murakami M, Takeo H, and Nawano S. Computerized detection of malignant tumors on digital mammograms. *IEEE Transactions on Medical Imaging*, 18(5):369–378, 1999.
- [635] Kegelmeyer Jr. WP. Evaluation of stellate lesion detection in a standard mammogram data set. *International Journal of Pattern Recognition and Artificial Intelligence*, 7(12):1477–1493, 1993.
- [636] Gupta R and Undrill PE. The use of texture analysis to delineate suspicious masses in mammography. *Physics in Medicine and Biology*, 40(5):835–855, 1995.
- [637] Priebe CE, Lorey RA, Marchette DJ, Solka JL, and Rogers GW. Nonparametric spatio-temporal change point analysis for early detection in mammography. In Gale AG, Astley SM, Dance DR, and Cairns AY, editors, *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 111–120, York, England, 10-12 July 1994.
- [638] Karssemeijer N. Recognition of stellate lesions in digital mammograms. In Gale AG, Astley SM, Dance DR, and Cairns AY, editors, *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 211–220, York, England, 10-12 July 1994.
- [639] Zhang M, Giger ML, Vyborny CJ, and Doi K. Mammographic texture analysis for the detection of spiculated lesions. In Doi K, Giger ML, Nishikawa RM, and Schmidt RA, editors, *Proceedings of the 3rd International Workshop on Digital Mammography*, pages 347–350, Chicago, IL, 9-12 June 1996.
- [640] Parr T, Astley S, and Boggis C. The detection of stellate lesions in digital mammograms. In Gale AG, Astley SM, Dance DR, and Cairns AY, editors, *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 231–240, York, England, 10-12 July 1994.
- [641] Karssemeijer N and te Brake GM. Detection of stellate distortions in mammograms. *IEEE Transactions on Medical Imaging*, 15(10):611–619, 1996.
- [642] te Brake GM and Karssemeijer N. Comparison of three mass detection methods. In Karssemeijer N, Thijssen M, Hendriks J, and van Erning L, editors,

- Proceedings of the 4th International Workshop on Digital Mammography*, pages 119–126, Nijmegen, The Netherlands, June 1998.
- [643] te Brake GM and Karssemeijer N. Single and multiscale detection of masses in digital mammograms. *IEEE Transactions on Medical Imaging*, 18(7):628–639, 1999.
- [644] Kobatake H and Yoshinaga Y. Detection of spicules on mammogram based on skeleton analysis. *IEEE Transactions on Medical Imaging*, 15(3):235–245, 1996.
- [645] Polakowski WE, Cournoyer DA, Rogers SK, DeSimio MP, Ruck DW, Hoffmeister JW, and Raines RA. Computer-aided breast cancer detection and diagnosis of masses using difference of Gaussians and derivative-based feature saliency. *IEEE Transactions on Medical Imaging*, 16(6):811–819, 1997.
- [646] Matsubara T, Fujita H, Hara T, Kasai S, Otsuka O, Hatanaka Y, and Endo T. Development of a new algorithm for detection of mammographic masses. In Karssemeijer N, Thijssen M, Hendriks J, and van Erning L, editors, *Proceedings of the 4th International Workshop on Digital Mammography*, pages 139–142, Nijmegen, The Netherlands, June 1998.
- [647] Parr T, Astley S, Taylor CJ, and Boggis CRM. Model based classification of linear structures in digital mammograms. In Doi K, Giger ML, Nishikawa RM, and Schmidt RA, editors, *Proceedings of the 3rd International Workshop on Digital Mammography*, pages 351–356, Chicago, IL, 9–12 June 1996.
- [648] Parr T, Taylor CJ, Astley S, and Boggis CRM. A statistical representation of pattern structure for digital mammography. In Doi K, Giger ML, Nishikawa RM, and Schmidt RA, editors, *Proceedings of the 3rd International Workshop on Digital Mammography*, pages 357–360, Chicago, IL, 9–12 June 1996.
- [649] Zwiggelaar R, Astley S, and Taylor C. Detecting the central mass of a spiculated lesion using scale-orientation signatures. In Karssemeijer N, Thijssen M, Hendriks J, and van Erning L, editors, *Proceedings of the 4th International Workshop on Digital Mammography*, pages 63–70, Nijmegen, The Netherlands, June 1998.
- [650] Parr T, Taylor CJ, Astley S, and Boggis CRM. Statistical modeling of oriented line patterns in mammograms. In *Proceedings of SPIE, Volume 3034: Medical Imaging — Image Processing*, pages 44–55, 1997.
- [651] Parr T, Zwiggelaar R, Astley S, Boggis C, and Taylor C. Comparison of methods for combining evidence for spiculated lesions. In Karssemeijer N, Thijssen M, Hendriks J, and van Erning L, editors, *Proceedings of the 4th International Workshop on Digital Mammography*, pages 71–78, Nijmegen, The Netherlands, June 1998.
- [652] Zwiggelaar R, Parr TC, Schumm JE, Hutt IW, Taylor CJ, Astley SM, and Boggis CRM. Model-based detection of spiculated lesions in mammograms. *Medical Image Analysis*, 3(1):39–62, 1999.
- [653] Rao AR and Schunck BG. Computing oriented texture fields. *Computer Vision, Graphics, and Image Processing*, 53(2):157–185, 1991.

- [654] Giger ML, Yin FF, Doi K, Metz CE, Schmidt RA, and Vyborny CJ. Investigation of methods for the computerized detection and analysis of mammographic masses. In *Proceedings of SPIE Volume 1233, Medical Imaging IV: Image Processing*, pages 183–184, 1990.
- [655] Yin FF, Giger ML, Doi K, Metz CE, Vyborny CJ, and Schmidt RA. Computerized detection of masses in digital mammograms: Analysis of bilateral subtraction images. *Medical Physics*, 18(5):955–963, 1991.
- [656] Sallam M and Bowyer KW. Registering time sequences of mammograms using a two-dimensional image unwarping technique. In Gale AG, Astley SM, Dance DR, and Cairns AY, editors, *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 121–130, York, England, 10-12 July 1994.
- [657] Stamatakis EA, Cairns AY, Ricketts IW, Walker C, Preece PE, and Thompson AJ. A novel approach to aligning mammograms. In Gale AG, Astley SM, Dance DR, and Cairns AY, editors, *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 355–364, York, England, 10-12 July 1994.
- [658] Nishikawa RM, Giger ML, Doi K, Vyborny CJ, and Schmidt RA. Computer-aided detection and diagnosis of masses and microcalcifications from digital mammograms. In Bowyer KW and Astley S, editors, *State of the Art in Digital Mammographic Image Analysis*, pages 82–102. World Scientific, Singapore, 1994.
- [659] Miller P and Astley S. Automated detection of mammographic asymmetry using anatomical features. In Bowyer KW and Astley S, editors, *State of the Art in Digital Mammographic Image Analysis*, pages 247–261. World Scientific, Singapore, 1994.
- [660] Burrell HC, Sibbering DM, Wilson ARM, Pinder SE, Evans AJ, Yeoman LJ, Elston CW, Ellis IO, Blamey RW, and Robertson JFR. Screening interval breast cancers: Mammographic features and prognostic factors. *Radiology*, 199:811–817, 1996.
- [661] Brzakovic D, Vujovic N, Neskovic M, Brzakovic P, and Fogarty K. Mammogram analysis by comparison with previous screenings. In Gale AG, Astley SM, Dance DR, and Cairns AY, editors, *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 131–140, York, England, 10-12 July 1994.
- [662] Sallam M and Bowyer KW. Detecting abnormal densities in mammograms by comparison to previous screenings. In Doi K, Giger ML, Nishikawa RM, and Schmidt RA, editors, *Proceedings of the 3rd International Workshop on Digital Mammography*, pages 417–420, Chicago, IL, 9-12 June 1996.
- [663] te Brake GM, Karssemeijer N, and Hendriks JHCL. Automated detection of breast carcinomas not detected in a screening program. *Radiology*, 207(2):465–471, 1998.
- [664] Sameti M, Morgan-Parkes J, Ward RK, and Palcic B. Classifying image features in the last screening mammograms prior to detection of a malignant mass. In Karssemeijer N, Thijssen M, Hendriks J, and van Erning L, editors,

- Proceedings of the 4th International Workshop on Digital Mammography*, pages 127–134, Nijmegen, The Netherlands, June 1998.
- [665] Hadjiiski L, Chan HP, Sahiner B, Petrick N, Helvie MA, and Gopal SS. Automated identification of breast lesions in temporal pairs of mammograms for interval change analysis. *Radiology*, 213(P):229–230, 1999.
- [666] Gopal SS, Chan HP, Wilson TE, Helvie MA, Petrick N, and Sahiner B. A regional registration technique for automated interval change analysis of breast lesions on mammograms. *Medical Physics*, 26:2669–2679, 1999.
- [667] Petrick N, Chan HP, Sahiner B, Helvie MA, and Paquerault S. Evaluation of an automated computer-aided diagnosis system for the detection of masses on prior mammograms. In *Proceedings of SPIE Volume 3979, Medical Imaging 2000: Image Processing*, pages 967–973, 2000.
- [668] Hadjiiski L, Chan HP, Sahiner B, Petrick N, Helvie MA, Paquerault S, and Zhou C. Interval change analysis in temporal pairs of mammograms using a local affine transformation. In *Proceedings of SPIE Volume 3979, Medical Imaging 2000: Image Processing*, pages 847–853, 2000.
- [669] Sameti M. *Detection of Soft Tissue Abnormalities in Mammographic Images for Early Diagnosis of Breast Cancer*. PhD thesis, Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, BC, Canada, November 1998.
- [670] Ranganath S. Image filtering using multiresolution representations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(5):426–440, 1991.
- [671] Rezaee MR, van der Zwet PMJ, Lelieveldt BPF, van der Geest RJ, and Reiber JHC. Multiresolution image segmentation technique based on pyramidal segmentation and fuzzy clustering. *IEEE Transactions on Image Processing*, 9(7):1238–1248, 2000.
- [672] Daubechies I. Ten lectures on wavelets. *CBMS, SIAM*, 61:198–202, 1995.
- [673] Shiffman S, Rubin GD, and Napel S. Medical image segmentation using analysis of isolable-contour maps. *IEEE Transactions on Medical Imaging*, 19(11):1064–1074, 2000.
- [674] Brown MB and Engelman L. *BMDP Statistical Software Manual*. University of California, Berkeley, CA, 1988.
- [675] Groshong BR and Kegelmeyer Jr. WP. Evaluation of a Hough transform method for circumscribed lesion detection. In Doi K, Giger ML, Nishikawa RM, and Schmidt RA, editors, *Proceedings of the 3rd International Workshop on Digital Mammography*, pages 361–366, Chicago, IL, 9–12 June 1996.
- [676] Mudigonda NR, Rangayyan RM, and Desautels JEL. Segmentation and classification of mammographic masses. In *SPIE Vol. 3979 Medical Imaging 2000: Image Processing*, pages 55–67, February 2000.
- [677] Schunck BG. *Gaussian filters and edge detection*, Research Publication GMR-5586. Computer Science Department, General Motors Research Laboratories, Detroit, MI, 1986.
- [678] Kass M and Witkin A. Analyzing oriented patterns. *Computer Vision, Graphics, and Image Processing*, 37:362–385, 1987.

- [679] Ayres FJ and Rangayyan RM. Characterization of architectural distortion in mammograms via analysis of oriented texture. *IEEE Engineering in Medicine and Biology Magazine*, 24:in press, January 2005.
- [680] Ayres FJ and Rangayyan RM. Detection of architectural distortion in mammograms using phase portraits. In Fitzpatrick JM and Sonka M, editors, *Proceedings of SPIE Medical Imaging 2004: Image Processing*, volume 5370, pages 587–597, San Diego, CA, February 2004.
- [681] Rangayyan RM, Ferrari RJ, and Frère AF. Detection of asymmetry between left and right mammograms. In *Proceedings of the 7th International Workshop on Digital Mammography*, Chapel Hill, NC, June 2004.
- [682] Yin FF, Giger ML, Doi K, Vyborny CJ, and Schmidt RA. Computerized detection of masses in digital mammograms: Automated alignment of breast images and its effect on bilateral-subtraction technique. *Medical Physics*, 21(3):445–452, 1994.
- [683] Vujovic N and Brzakovic D. Establishing the correspondence between control points in pairs of mammographic images. *IEEE Transactions on Image Processing*, 6(10):1388–1399, 1997.
- [684] Karssemeijer N and te Brake GM. Combining single view features and asymmetry for detection of mass lesions. In Karssemeijer N, Thijssen M, Hendriks J, and van Erning L, editors, *Proceedings of the 4th International Workshop on Digital Mammography*, pages 95–102, Nijmegen, The Netherlands, June 1998.
- [685] Doi K, Giger ML, Nishikawa RM, and Schmidt RA, editors. *Proceedings of the 3rd International Workshop on Digital Mammography*, Chicago, IL, June 1996. Elsevier.
- [686] Karssemeijer N, Thijssen M, Hendriks J, and van Erning L, editors. *Proceedings of the 4th International Workshop on Digital Mammography*, Nijmegen, The Netherlands, June 1998. Kluwer Academic Publishers.
- [687] Yaffe MJ, editor. *Proceedings of the 5th International Workshop on Digital Mammography*, Toronto, ON, Canada, June 2000. Medical Physics Publishing.
- [688] Peitgen HO, editor. *Proceedings of the 6th International Workshop on Digital Mammography*, Bremen, Germany, June 2002. Springer-Verlag.
- [689] Wolfe JN. Risk for breast cancer development determined by mammographic parenchymal pattern. *Cancer*, 37:2486–2492, 1976.
- [690] Matsubara T, Yamazaki D, Fujita H, Hara T, Iwase T, and Endo T. An automated classification method for mammograms based on evaluation of fibroglandular breast tissue density. In Yaffe MJ, editor, *Proceedings of the 5th International Workshop on Digital Mammography*, pages 737–741, Toronto, ON, Canada, June 2000.
- [691] Zhou C, Chan HP, Petrick N, Helvie MA, Goodsitt MM, Sahiner B, and Hadjiiski LM. Computerized image analysis: Estimation of breast density on mammograms. *Medical Physics*, 28(6):1056–1069, 2001.

- [692] Byng JW, Boyd NF, Fishell E, Jong RA, and Yaffe MJ. The quantitative analysis of mammographic densities. *Physics in Medicine and Biology*, 39:1629–1638, 1994.
- [693] Tahoces PG, Correa J, Souto M, Gómez L, and Vidal JJ. Computer-assisted diagnosis: the classification of mammographic breast parenchymal patterns. *Physics in Medicine and Biology*, 40:103–117, 1995.
- [694] Huo Z, Giger ML, Zhong W, and Olopade OI. Analysis of relative contributions of mammographic features and age to breast cancer risk prediction. In Yaffe MJ, editor, *Proceedings of the 5th International Workshop on Digital Mammography*, pages 732–736, Toronto, ON, Canada, June 2000.
- [695] Sivaramakrishna R, Obuchowski NA, Chilcote WA, and Powell KA. Automatic segmentation of mammographic density. *Academic Radiology*, 8(3):250–256, 2001.
- [696] Bassett LW and Gold RH. *Breast Cancer Detection: Mammography and Other Methods in Breast Imaging*. Grune & Stratton, Orlando, FL, 2nd edition, 1987.
- [697] Caulkin S, Astley S, Asquith J, and Boggis C. Sites of occurrence of malignancies in mammograms. In Karssemeijer N, Thijssen M, Hendriks J, and van Erning L, editors, *Proceedings of the 4th International Workshop on Digital Mammography*, pages 279–282, Nijmegen, The Netherlands, June 1998.
- [698] McLachlan GJ and Krishnan T. *The EM Algorithm and Extensions*. Wiley-Interscience, New York, NY, 1997.
- [699] Rissanen J. Modeling by shortest data description. *Automatica*, 14:465–471, 1978.
- [700] Bishop CM. *Neural Networks for Pattern Recognition*. Clarendon Press, Oxford, England, 1995.
- [701] Wolfe JN. Breast parenchymal patterns and their changes with age. *Radiology*, 121:545–552, 1976.
- [702] Rutter CM, Mandelson MT, Laya MB, and Taplin S. Changes in breast density associated with initiation, discontinuation, and continuing use of hormone replacement therapy. *Journal of the American Medical Association*, 285(2):171–176, 2001.
- [703] Dempster AP, Laird NM, and Rubin DB. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, B* 39(1):1–38, 1977.
- [704] Liavas AP and Regalia PA. On the behavior of information theoretic criteria for model order selection. *IEEE Transactions on Signal Processing*, 49(8):1689–1695, 2001.
- [705] Hansen MH and Yu B. Model selection and the principle of minimum description length. *Journal of the American Statistical Association*, 96(454):746–774, 2001.
- [706] Carson C, Thomas M, Belongie S, Hellerstein JM, and Malik J. Blobworld: A system for region-based image indexing and retrieval. In Huijismans DP and

- Smeulders AWM, editors, *Proceedings of the 3rd International Conference on Visual Information and Information Systems*, pages 509–516, Amsterdam, The Netherlands, June 1999.
- [707] Dance DR. Physical principles of breast imaging. In Doi K, Giger ML, Nishikawa RM, and Schmidt RA, editors, *Proceedings of the 3rd International Workshop on Digital Mammography*, pages 427–430, Chicago, IL, June 1996.
- [708] Ueda N and Nakano R. Deterministic annealing EM algorithm. *Neural Networks*, 11:271–282, 1998.
- [709] Masson P and Pieczynski W. SEM algorithm and unsupervised statistical segmentation of satellite images. *IEEE Transactions on Geoscience and Remote Sensing*, 31(3):618–633, 1993.
- [710] Campbell FW and Robson JG. Application of Fourier analysis to the visibility of gratings. *Journal of Physiology*, 197:551–566, 1968.
- [711] Marcelja S. Mathematical description of the response of simple cortical cells. *Journal of the Optical Society of America*, 70(11):1297–1300, 1980.
- [712] Daugman JG. Complete discrete 2-D Gabor transforms by neural networks for image analysis and compression. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 36(7):1169–1179, 1988.
- [713] Malik J and Perona P. Preattentive texture discrimination with early vision mechanisms. *Journal of the Optical Society of America A*, 7(2):923–932, 1990.
- [714] Dunn D, Higgins WE, and Wakeley J. Texture segmentation using 2-D Gabor elementary functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(2):130–149, 1994.
- [715] Chang T and Kuo CCJ. Texture analysis and classification with tree-structured wavelet transform. *IEEE Transactions on Image Processing*, 2(4):429–441, 1993.
- [716] Li L, Mao F, Qian W, and Clarke LP. Wavelet transform for directional feature extraction in medical imaging. In *Proceedings of the IEEE International Conference on Image Processing*, volume 3, pages 500–503, Santa Barbara, CA, 1997.
- [717] Qian W and Clarke LP. Hybrid M-channel wavelet transform methods: adaptive, automatic and digital X-ray sensor independent. *Medical Physics*, 22(6):983–984, 1995.
- [718] Li L, Qian W, and Clarke LP. Digital mammography: CAD method for mass detection using multiresolution and multiorientation wavelet transforms. *Academic Radiology*, 4:724–731, 1997.
- [719] Graham NVS. *Visual Pattern Analyzers*. Oxford University Press, New York, NY, 1989.
- [720] Daubechies I. The wavelet transform, time-frequency localization and signal analysis. *IEEE Transactions on Information Theory*, 36(5):961–1004, 1990.
- [721] Duda RO and Hart PE. *Pattern Classification and Scene Analysis*. Wiley, New York, NY, 1973.

- [722] Schneider MA. Better detection: Improving our chances. In Yaffe MJ, editor, *Digital Mammography: 5th International Workshop on Digital Mammography*, pages 3–6, Toronto, ON, Canada, June 2000. Medical Physics Publishing.
- [723] Bird RE, Wallace TW, and Yankaskas BC. Analysis of cancers missed at screening mammography. *Radiology*, 184:613–617, 1992.
- [724] Baker JA, Rosen EL, Lo JY, Gimenez EI, Walsh R, and Soo MS. Computer-aided detection (CAD) in screening mammography: Sensitivity of commercial CAD systems for detecting architectural distortion. *American Journal of Roentgenology*, 181:1083–1088, 2003.
- [725] van Dijck JAAM, Verbeek ALM, Hendriks JHCL, and Holland R. The current detectability of breast cancer in a mammographic screening program. *Cancer*, 72:1933–1938, 1993.
- [726] Sickles EA. Mammographic features of 300 consecutive nonpalpable breast cancers. *American Journal of Roentgenology*, 146:661–663, 1986.
- [727] Broeders MJM, Onland-Moret NC, Rijken HJTM, Hendriks JHCL, Verbeek ALM, and Holland R. Use of previous screening mammograms to identify features indicating cases that would have a possible gain in prognosis following earlier detection. *European Journal of Cancer*, 39:1770–1775, 1993.
- [728] Kegelmeyer Jr. WP, Pruneda JM, Bourland PD, Hillis A, Riggs MW, and Nipper ML. Computer-aided mammographic screening for spiculated lesions. *Radiology*, 191:331–337, 1994.
- [729] Sampat MP and Bovik AC. Detection of spiculated lesions in mammograms. In *Proceedings of the 25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (CD-ROM)*, pages 810–813, Cancún, Mexico, September 2003.
- [730] Mudigonda NR and Rangayyan RM. Texture flow-field analysis for the detection of architectural distortion in mammograms. In *Proceedings of Bio Vision*, pages 76–81, Bangalore, India, December 2001.
- [731] Matsubara T, Ichikawa T, Hara T, Fujita H, Kasai S, Endo T, and Iwase T. Automated detection methods for architectural distortions around skin-line and within mammary gland on mammograms. In Lemke HU, Vannier MW, Inamura K, Farman AG, Doi K, and Reiber JHC, editors, *International Congress Series: Proceedings of the 17th International Congress and Exhibition on Computer Assisted Radiology and Surgery*, pages 950–955, London, UK, June 2003. Elsevier.
- [732] Burhenne LJW, Wood SA, D’Orsi CJ, Feig SA, Kopans DB, O’Shaughnessy KF, Sickles EA, Tabar L, Vyborny CJ, and Castellino RA. Potential contribution of computer-aided detection to the sensitivity of screening mammography. *Radiology*, 215:554–562, 2000.
- [733] Evans WP, Burhenne LJW, Laurie L, O’Shaughnessy KF, and Castellino RA. Invasive lobular carcinoma of the breast: Mammographic characteristics and computer-aided detection. *Radiology*, 225(1):182–189, 2002.
- [734] Birdwell RL, Ikeda DM, O’Shaughnessy KF, and Sickles EA. Mammographic characteristics of 115 missed cancers later detected with screening mam-

- mography and the potential utility of computer-aided detection. *Radiology*, 219(1):192–202, 2001.
- [735] Wylie CR and Barrett LC. *Advanced Engineering Mathematics*. McGraw-Hill, New York, NY, 6th edition, 1995.
- [736] Rao AR and Jain RC. Computerized flow field analysis: Oriented texture fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(7):693–709, 1992.
- [737] Gershenfeld N. *The Nature of Mathematical Modeling*. Cambridge University Press, Cambridge, UK, 1999.
- [738] Sweet SA. *Data Analysis with SPSS*. Allyn & Bacon, Boston, MA, 1999.
- [739] Louis AK and Natterer F. Mathematical problems of computerized tomography. *Proceedings of the IEEE*, 71(3):379–389, 1983.
- [740] Lewitt RM. Reconstruction algorithms: Transform methods. *Proceedings of the IEEE*, 71(3):390–408, 1983.
- [741] Censor Y. Finite series-expansion reconstruction methods. *Proceedings of the IEEE*, 71(3):409–419, 1983.
- [742] Gordon R. A tutorial on ART (Algebraic Reconstruction Techniques). *IEEE Transactions on Nuclear Science*, 21:78–93, 1974.
- [743] Gordon R and Herman GT. Three-dimensional reconstruction from projections: A review of algorithms. *International Review of Cytology*, 38:111–151, 1974.
- [744] Gordon R and Rangayyan RM. Geometric deconvolution: A meta-algorithm for limited view computed tomography. *IEEE Transactions on Biomedical Engineering*, 30:806–810, 1983.
- [745] Gordon R, Dhawan AP, and Rangayyan RM. Reply to comments on geometric deconvolution: A meta-algorithm for limited view computed tomography. *IEEE Transactions on Biomedical Engineering*, 32:242–244, 1985.
- [746] Soble PJ, Rangayyan RM, and Gordon R. Quantitative and qualitative evaluation of geometric deconvolution of distortion in limited-view computed tomography. *IEEE Transactions on Biomedical Engineering*, 32:330–335, 1985.
- [747] Rangayyan RM, Gordon R, and Dhawan AP. Algorithms for limited-view computed tomography: An annotated bibliography and a challenge. *Applied Optics*, 24(23):4000–4012, 1985.
- [748] Dhawan AP, Rangayyan RM, and Gordon R. Image restoration by Wiener deconvolution in limited-view computed tomography. *Applied Optics*, 24(23):4013–4020, 1985.
- [749] Boulfelfel D, Rangayyan RM, Hahn LJ, and Kloiber R. Three-dimensional restoration of single photon emission computed tomography images. *IEEE Transactions on Nuclear Science*, 41(5):1746–1754, 1994.
- [750] Boulfelfel D, Rangayyan RM, Hahn LJ, Kloiber R, and Kuduvalli GR. Restoration of single photon emission computed tomography images by the Kalman filter. *IEEE Transactions on Medical Imaging*, 13(1):102–109, 1994.

- [751] Boulfefel D, Rangayyan RM, Hahn LJ, and Kloiber R. Pre-reconstruction restoration of single photon emission computed tomography images. *IEEE Transactions on Medical Imaging*, 11(3):336–341, 1992.
- [752] Herman GT, Lent A, and Rowland SW. ART: Mathematics and Applications – A report on the mathematical foundations and on the applicability to real data of the Algebraic Reconstruction Techniques. *Journal of Theoretical Biology*, 42:1–32, 1973.
- [753] Rangayyan RM and Gordon R. Streak preventive image reconstruction via ART and adaptive filtering. *IEEE Transactions on Medical Imaging*, 1:173–178, 1982.
- [754] Kaczmarz MS. Angenäherte auflösung von systemen linearer gleichungen. *Bulletin International de l'Academie Polonaise des Sciences et des Lettres Serie A, Sciences Mathematiques*, pages 355–357, 1937.
- [755] Guan H and Gordon R. Computed tomography using ART with different projection access schemes: a comparison study under practical situations. *Physics in Medicine and Biology*, 41:1727–1743, 1996.
- [756] Lent A. A convergent algorithm for maximum entropy image restoration, with a medical x-ray application. In Shaw R, editor, *Image Analysis and Evaluation*, pages 249–257. Society of Photographic Scientists and Engineers, Washington DC, 1977.
- [757] Gilbert P. Iterative methods for the three-dimensional reconstruction of an object from projections. *Journal of Theoretical Biology*, 36:105–117, 1972.
- [758] Fullerton GD. Fundamentals of CT tissue characterization. In Fullerton GD and Zagzebski JA, editors, *Medical Physics of CT and Ultrasound: Tissue Imaging and Characterization*, pages 125–162. American Association of Physicists in Medicine, New York, NY, 1980.
- [759] Mategrano VC, Petasnick J, Clark J, Bin AC, and Weinstein R. Attenuation values in computed tomography of the abdomen. *Radiology*, 125:135–140, October 1977.
- [760] Cruvinel PE, Cesareo R, Crestana S, and Mascarenhas S. X- and gamma-rays computerized minitomograph scanner for soil science. *IEEE Transactions on Instrumentation and Measurements*, 39(5):745–750, 1990.
- [761] Vaz CMP, Crestana S, Mascarenhas S, Cruvinel PE, Reichardt K, and Stolf R. Using a computed tomography miniscanner for studying tillage induced soil compaction. *Soil Technology*, 2:313–321, 1989.
- [762] Onoe AM, Tsao JW, Yamada H, Nakamura H, Kogure J, Kawamura H, and Yoshimatsu M. Computed tomography for measuring annual rings of a live tree. *Proceedings of the IEEE*, 71(7):907–908, 1983.
- [763] Stock SR. X-ray microtomography of materials. *International Materials Review*, 44(4):141–164, 1999.
- [764] Johnson RH, Karau KL, Molthen RC, Haworth ST, and Dawson CA. Micro-CT image-derived metrics quantify arteial wall distensibility reduction in a rat model of pulmonary hypertension. In *Proceedings SPIE 3978: Medical Imaging 2000 – Physiology and Function from Multidimensional Images*, pages 320–330, San Diego, CA, February 2000.

- [765] Illman B and Dowd B. High-resolution microtomography for density and spatial information about wood structures. In *Proceedings SPIE 3772: Developments in X-ray Tomography*, pages 198–330, Denver, CO, July 1999.
- [766] Shimizu K, Ikezoe J, Ikura H, Ebara H, Nagareda T, Yagi N, Umetani K, Uesugi K, Okada K, Sugita A, and Tanaka M. Synchrotron radiation microtomography of the lung specimens. In *Proceedings SPIE 3977: Medical Imaging 2000 – Physics of Medical Imaging*, pages 196–204, San Diego, CA, February 2000.
- [767] Umetani K, Yagi N, Suzuki Y, Ogasawara Y, Kajiya F, Matsumoto T, Tachibana H, Goto M, Yamashita T, Imai S, and Kajihara Y. Observation and analysis of microcirculation using high-spatial-resolution image detectors and synchrotron radiation. In *Proceedings SPIE 3977: Medical Imaging 2000 – Physics of Medical Imaging*, pages 522–533, San Diego, CA, February 2000.
- [768] Sasov A. High-resolution in-vivo micro-CT scanner for small animals. In *Proceedings SPIE 4320: Medical Imaging 2001 – Physics of Medical Imaging*, pages 705–710, San Diego, CA, February 2001.
- [769] Shaler SM, Keane DT, Wang H, Mott L, Landis E, and Holzman L. Microtomography of cellulosic structures. In *TAPPI Proceedings: Process and Product Quality Conference*, pages 89–96, 1998.
- [770] Machin K and Webb S. Cone-beam x-ray microtomography of small specimens. *Physics in Medicine and Biology*, 39:1639–1657, 1994.
- [771] Boyd SK, Müller R, Matyas JR, Wohl GR, and Zernicke RF. Early morphometric and anisotropic change in periarticular cancellous bone in a model of experimental knee osteoarthritis quantified using microcomputed tomography. *Clinical Biomechanics*, 15:624–631, 2000.
- [772] Boyd SK. *Microstructural Bone Adaptation in an Experimental Model of Osteoarthritis*. PhD thesis, Faculty of Kinesiology, University of Calgary, Calgary, Alberta, Canada, August 2001.
- [773] Alexander F. Neuroblastoma. *Urologic Clinics of North America*, 27(3):383–392, 2000.
- [774] Cotterill SJ, Pearson ADJ, Pritchard J, Foot ABM, Roald B, Kohler JA, and Imeson J. Clinical prognostic factors in 1277 patients with neuroblastoma: Results of The European Neuroblastoma Study Group ‘Survey’ 1982-1992. *European Journal of Cancer*, 36:901–908, 2000.
- [775] Meza MP, Benson M, and Slovis TL. Imaging of mediastinal masses in children. *Radiologic Clinics of North America*, 31(3):583–604, 1993.
- [776] Abramson SJ. Adrenal neoplasm in children. *Radiologic Clinics of North America*, 35(6):1415–1453, 1997.
- [777] Castleberry RP. Neuroblastoma. *European Journal of Cancer*, 33(9):1430–1438, 1997.
- [778] Goodman MT, Gurney JG, Smith MA, and Olshan AF. *Cancer incidence and survival among children and adolescents: United States Surveillance, Epidemiology, and End Results (SEER) Program 1975-1995*.

- Chapter IV Sympathetic nervous system tumors.* National Cancer Institute, <http://seer.cancer.gov/publications/childhood/sympathetic.pdf>, accessed May 2003.
- [779] Gurney JG, Ross JA, Wall DA, Bleyer WA, Severson RK, and Robison LL. Infant cancer in the U.S.: Histology-specific incidence and trends. *Journal of Pediatric Hematology/Oncology*, 19(5):428–432, 1997.
- [780] Parker L and Powell J. Screening for neuroblastoma in infants younger than 1 year of age: Review of the first 30 years. *Medical and Pediatric Oncology*, 31:455–469, 1998.
- [781] Woods WG and Tuchman M. A population-based study of the usefulness of screening for neuroblastoma. *Lancet*, 348(9043):1682–1687, 1998.
- [782] Bousvaros A, Kirks DR, and Grossman H. Imaging of neuroblastoma: An overview. *Pediatric Radiology*, 16:89–106, 1986.
- [783] Brodeur GM, Pritchard J, Berthold F, Carlsen NLT, Castel V, Castleberry RP, de Bernardi B, Evans AE, Favrot M, Hedborg F, Kaneko M, Kemshead J, Lampert F, Lee REJ, Look T, Pearson ADJ, Philip T, Roald B, Sawada T, Seeger RC, Tsuchida Y, and Voute PA. Revisions of the international criteria for neuroblastoma diagnosis, staging, and response to treatment. *Journal of Clinical Oncology*, 11(8):1466–1477, 1993.
- [784] Stark DD, Moss AA, Brasch RC, deLorimier AA, Albin AR, London DA, and Gooding CA. Neuroblastoma: Diagnostic imaging and staging. *Radiology*, 148:101–105, July 1983.
- [785] Kirks DR, Merten DF, Grossman H, and Bowie JD. Diagnostic imaging of pediatric abdominal masses: An overview. *Radiologic Clinics of North America*, 19(3):527–545, 1981.
- [786] Cohen MD, Bugaieski EM, Haliloglu M, Faught P, and Siddiqui AR. Visual presentation of the staging of pediatric solid tumors. *RadioGraphics*, 16(3):523–545, 1996.
- [787] Boechat MI, Ortega J, Hoffman AD, Cleveland RH, Kangaroo H, and Gilsanz V. Computed tomography in Stage III neuroblastoma. *American Journal of Radiology*, 145:1456–1283, December 1985.
- [788] Corbett R, Olliff J, Fairley N, Moyes J, Husband J, Pinkerton R, Carter R, Treleaven J, McElwain T, and Meller S. A prospective comparison between magnetic resonance imaging, meta-iodobenzylguanidine scintigraphy and marrow histology/cytology in neuroblastoma. *European Journal of Cancer*, 27(12):1560–1564, 1991.
- [789] Fletcher BD, Kapiwoda SY, Strandjord SE, Nelson AD, and Pickering SP. Abdominal neuroblastoma: Magnetic resonance imaging and tissue characterization. *Radiology*, 155(3):699–703, 1985.
- [790] Sofka CM, Semelka RC, Kelekis NL, Worawattanakul S, Chung CJ, Gold S, and Fordham LA. Magnetic resonance imaging of neuroblastoma using current techniques. *Magnetic Resonance Imaging*, 17(2):193–198, 1999.
- [791] Chezmar JL, Robbins SM, Nelson RC, Steinberg HV, Torres WE, and Bernardino ME. Adrenal masses: Characterization with T1-weighted MR imaging. *Radiology*, 166(2):357–359, 1988.

- [792] Kornreich L, Horev G, Kaplinsky NZ, and Grunebaum M. Neuroblastoma: Evaluation with contrast enhanced MR imaging. *Pediatric Radiology*, 21:566–569, 1991.
- [793] Foglia RP, Fonkalsrud EW, Feig SA, and Moss TJ. Accuracy of diagnostic imaging as determined by delayed operative intervention for advanced neuroblastoma. *Journal of Pediatric Surgery*, 24(7):708–711, 1989.
- [794] Sonka M and Fitzpatrick JM, editors. *Handbook of Medical Imaging, Volume 2: Medical Image Processing and Analysis*. SPIE Press, Bellingham, WA, 2000.
- [795] Duncan JS and Ayache N. Medical image analysis: Progress over two decades and the challenges ahead. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(1):85–106, January 2000.
- [796] Wheatley JM, Rosenfield NS, Heller G, Feldstein D, and LaQuaglia MP. Validation of a technique of computer-aided tumor volume determination. *Journal of Surgical Research*, 59(6):621–626, 1995.
- [797] Hopper KD, Singapuri K, and Finkel A. Body CT and oncologic imaging. *Radiology*, 215(1):27–40, 2000.
- [798] Ayres FJ. Segmentação e estimação da composição histológica da massa tumoral em imagens de CT de neuroblastomas. Master's thesis, Universidade de São Paulo, São Paulo, Brazil, August 2001.
- [799] Ayres FJ, Zuffo MK, Rangayyan RM, Odone Filho V, and Valente M. Segmentation and estimation of the histological composition of the tumor mass in computed tomographic images of neuroblastoma. In *Proceedings of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Istanbul, Turkey, October 2001.
- [800] Ayres FJ, Zuffo MK, Rangayyan RM, Boag GS, Odone Filho V, and Valente M. Estimation of the tissue composition of the tumor mass in neuroblastoma using segmented CT images. *Medical and Biological Engineering and Computing*, 42:366–377, 2004.
- [801] Rao PS and Gregg EC. Attenuation of monoenergetic gamma rays in tissues. *American Journal of Roentgenology*, 123(3):631–637, 1975.
- [802] Wilson CR. Quantitative computed tomography. In Fullerton GD and Zagzebski JA, editors, *Medical Physics of CT and Ultrasound: Tissue Imaging and Characterization*, pages 163–175. American Association of Physicists in Medicine, New York, NY, 1980.
- [803] Brooks RA. A quantitative theory of the Hounsfield unit and its application to dual energy scanning. *Journal of Computer Assisted Tomography*, 1(4):487–493, 1977.
- [804] Alter AJ. Computerized tomography: A clinical perspective. In Fullerton GD and Zagzebski JA, editors, *Medical Physics of CT and Ultrasound: Tissue Imaging and Characterization*, pages 125–162. American Association of Physicists in Medicine, New York, NY, 1980.
- [805] Williams G, Bydder GM, and Kreel L. The validity and use of computed tomography attenuation values. *British Medical Bulletin*, 36(3):279–287, 1980.

- [806] Duerinckx AJ and Macovski A. Information and artifact in computed tomography image statistics. *Medical Physics*, 7(2):127–134, March-April 1980.
- [807] Pullan BR, Fawcitt RA, and Isherwood I. Tissue characterization by an analysis of the distribution of attenuation values in computed tomography scans: A preliminary report. *Journal of Computer Assisted Tomography*, 2(1):49–54, 1978.
- [808] Kramer RA, Yoshikawa BM, Scheibe PO, and Janetos GP. Statistical profiles in computed tomography. *Radiology*, 125:145–147, October 1977.
- [809] Latchaw RE, Gold LHA, Moore JS, and Payne JT. The nonspecificity of absorption coefficients in the differentiation of solid tumors and cystic lesions. *Radiology*, 125:141–144, October 1977.
- [810] Goodenough DJ. Tomographic imaging. In Beutel J, Kundel H L, and Van Metter R L, editors, *Handbook of Medical Imaging, Volume 1: Physics and Psychophysics*, chapter 8, pages 511–554. SPIE Press, Bellingham, WA, 2000.
- [811] Jain AK, Duin RPW, and Mao J. Statistical pattern recognition: A review. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(1):4–37, January 2000.
- [812] Tanner MA. *Tools for Statistical Inference: Methods for the Exploration of Posterior Distributions and Likelihood Functions*. Springer-Verlag, New York, NY, 3rd edition, 1996.
- [813] Dawant BM and Zijdenbos AP. Image segmentation. In Sonka M and Fitzpatrick JM, editors, *Handbook of Medical Imaging, Volume 2: Medical Image Processing and Analysis*, chapter 2, pages 71–127. SPIE Press, Bellingham, WA, 2000.
- [814] Copsey K and Webb A. Bayesian approach to mixture models for discrimination. In Ferri FJ, Iñesta JM, Amin A, and Pudil P, editors, *Advances in Pattern Recognition, Joint IAPR International Workshops SSPR 2000 and SPR 2000, [8th International Workshop on Structural and Syntactic Pattern Recognition, 3rd International Workshop on Statistical Techniques in Pattern Recognition]*, pages 491–500, Alicante, Spain, August 30 – September 1, 2000. Springer. Lecture Notes in Computer Science, Vol. 1876.
- [815] Richardson S and Green PJ. On Bayesian analysis of mixtures with an unknown number of components. *Journal of the Royal Statistical Society B*, 59(4):731–792, 1997.
- [816] Meng XL and van Dyk D. The EM algorithm – and old folk-song sung to a fast new tune. *Journal of the Royal Statistical Society*, 59(3):511–567, 1997.
- [817] Rangayyan RM, Ferrari RJ, Desautels JEL, and Frère AF. Directional analysis of images with Gabor wavelets. In *Proceedings of SIBGRAPI 2000: XIII Brazilian Symposium on Computer Graphics and Image Processing*, pages 170–177, Gramado, Rio Grande do Sul, Brazil, 17-20 October 2000. IEEE Computer Society Press.
- [818] Goldszal AF and Pham DL. Volumetric segmentation. In Bankman IN, editor, *Handbook of Medical Imaging: Processing and Analysis*, chapter 12, pages 185–194. Academic Press, London, UK, 2000.

- [819] Laidlaw DH, Fleischer KW, and Barr AH. Partial volume segmentation with voxel histograms. In Bankman IN, editor, *Handbook of Medical Imaging: Processing and Analysis*, chapter 13, pages 185–194. Academic Press, London, UK, 2000.
- [820] Chen EL, Chung PC, Chen CL, Tsai HM, and Chang CI. An automatic diagnostic system for CT liver image classification. *IEEE Transactions on Biomedical Engineering*, 45(6):783–794, 1998.
- [821] Srinivasa N, Ramakrishnan KR, and Rajgopal K. Detection of edges from projections. *IEEE Transactions on Medical Imaging*, 11(1):76–80, 1992.
- [822] Stark H, editor. *Image Recovery: Theory and Application*. Academic, Orlando, FL, 1987.
- [823] Sezan MI, editor. *Selected Papers on Digital Image Restoration*. SPIE, Bellingham, WA, 1992.
- [824] Jansson PA, editor. *Deconvolution of Images and Spectra*. Academic, San Diego, CA, 2nd edition, 1984.
- [825] Andrews HC and Hunt BR. *Digital Image Restoration*. Prentice Hall, Englewood Cliffs, NJ, 1977.
- [826] Sondhi MM. Image restoration: The removal of spatially invariant degradations. *Proceedings of the IEEE*, 60:842–853, 1972.
- [827] Sawchuk AA. Space-variant motion degradation and restoration. *Proceedings of the IEEE*, 60:854–861, 1972.
- [828] Robbins GM and Huang TS. Inverse filtering for space-variant imaging systems. *Proceedings of the IEEE*, 60:862–872, 1972.
- [829] Sezan MI and Tekalp AM. A survey of recent developments in digital image restoration. *Optical Engineering*, 29:393–404, 1990.
- [830] Sanz JLC and Huang TS. Unified Hilbert space approach to iterative least-squares linear signal restoration. *Journal of the Optical Society of America*, 73:1455–1465, 1983.
- [831] McGlamery BL. Restoration of turbulence-degraded images. *Journal of the Optical Society of America*, 57(3):293–297, 1967.
- [832] Alonso Jr. M and Barreto AB. Pre-compensation for high-order aberrations of the human eye using on-screen image deconvolution. In *Proceedings of the 25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 556–559, Cancún, Mexico, 2003.
- [833] Haykin S. *Adaptive Filter Theory*. Prentice Hall, Upper Saddle River, NJ, 4th edition, 2002.
- [834] King MA, Doherty PW, and Schwinger RB. A Wiener filter for nuclear medicine images. *Medical Physics*, 10(6):876–880, 1983.
- [835] Honda N, Machida K, Tsukada J, Kaizu H, and Hosoba M. Optimal preprocessing Butterworth-Wiener filter for Tl-201 myocardial SPECT. *European Journal of Nuclear Medicine*, 13:404–407, 1987.
- [836] Tikhonov AN and Arsenin VY. *Solutions of Ill-posed Problems*. VH Winston and Sons, Washington, DC, 1977.

- [837] Frieden BR. Restoring with maximum likelihood and maximum entropy. *Journal of the Optical Society of America*, 62:511–518, 1972.
- [838] Gull SF and Daniell GJ. Image reconstruction from incomplete and noisy data. *Nature*, 272:686–690, 1978.
- [839] Leahy RM and Goutis CE. An optimal technique for constraint-based image restoration and reconstruction. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 34:1629–1642, 1986.
- [840] Biraud Y. A new approach for increasing the resolving power by data processing. *Astronomy and Astrophysics*, 1:124–127, 1969.
- [841] Boas Jr. RP and Kac M. Inequalities for Fourier transforms of positive functions. *Duke Mathematics Journal*, 12:189–206, 1945.
- [842] Webb S, Long AP, Ott RJ, Leach MO, and Flower MA. Constrained deconvolution of SPECT liver tomograms by direct digital image restoration. *Medical Physics*, 12:53–58, 1985.
- [843] Metz CE and Beck RN. Quantitative effects of stationary linear image processing on noise and resolution of structure in radionuclide images. *Journal of Nuclear Medicine*, 15:164–170, 1974.
- [844] King MA, Doherty PW, Schwinger RB, Jacobs DA, Kidder RE, and Miller TR. Fast count-dependent digital filtering of nuclear medicine images: Concise communication. *Journal of Nuclear Medicine*, 24:1039–1045, 1983.
- [845] King MA, Schwinger RB, Doherty PW, and Penney BC. Two-dimensional filtering of SPECT images using the Metz and Wiener filters. *Journal of Nuclear Medicine*, 25:1234–1240, 1984.
- [846] King MA, Schwinger RB, Penney BC, Doherty PW, and Bianco JA. Digital restoration of Indium-111 and Iodine-123 SPECT images with optimized Metz filters. *Journal of Nuclear Medicine*, 27:1327–1336, 1986.
- [847] King MA, Schwinger RB, and Penney BC. Variation of the count-dependent Metz filter with imaging system modulation transfer function. *Medical Physics*, 13(2):139–149, 1986.
- [848] King MA, Glick SJ, Penney BC, Schwinger RB, and Doherty PW. Interactive visual optimization of SPECT prereconstruction filtering. *Journal of Nuclear Medicine*, 28:1192–1198, 1987.
- [849] King MA, Penney BC, and Glick SJ. An image-dependent Metz filter for nuclear medicine images. *Journal of Nuclear Medicine*, 29:1980–1989, 1988.
- [850] Gilland DR, Tsui BMW, McCartney WH, Perry JR, and Berg J. Determination of the optimum filter function for SPECT imaging. *Journal of Nuclear Medicine*, 29:643–650, 1988.
- [851] Boulfelfel D, Rangayyan RM, Hahn LJ, and Kloiber R. Pre-reconstruction restoration versus post-reconstruction restoration of single photon emission computed tomography images. In *Proceedings of the 1990 IEEE Colloquium in South America*, pages 112–118. IEEE, Piscataway, NJ, September, 1990.
- [852] Rabie TF, Paranjape RB, and Rangayyan RM. Iterative method for blind deconvolution. *Journal of Electronic Imaging*, 3(3):245–250, 1994.

- [853] Hunt BR. Digital image processing. In Oppenheim AV, editor, *Applications of Digital Signal Processing*, pages 169–237. Prentice Hall, Englewood Cliffs, NJ, 1978.
- [854] Oppenheim AV and Lim JS. The importance of phase in signals. *Proceedings of the IEEE*, 69(5):529–541, 1981.
- [855] Hayes MH, Lim JS, and Oppenheim AV. Signal reconstruction from phase or magnitude. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 28(6):672–680, 1980.
- [856] Huang TS, Burnett JW, and Deczky AD. The importance of phase in image processing filters. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 23(6):529–542, 1975.
- [857] Behar J, Porat M, and Zeevi YY. Image reconstruction from localized phase. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 40(4):736–743, 1992.
- [858] Oppenheim AV, Hayes MH, and Lim JS. Iterative procedures for signal reconstruction from phase. In *Proceedings of SPIE Vol. 231: International Conference on Optical Computing*, pages 121–129, 1980.
- [859] Hayes MH. The reconstruction of a multidimensional sequence from the phase or magnitude of its Fourier transform. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 30(2):140–154, 1982.
- [860] Kermisch D. Image reconstruction from phase information only. *Journal of the Optical Society of America*, 60(1):15–17, 1970.
- [861] Espy CY and Lim JS. Effects of additive noise on signal reconstruction from Fourier transform phase. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 31(4):894–898, 1983.
- [862] Stockham Jr. TG, Cannon TM, and Ingebretsen RB. Blind deconvolution through digital signal processing. *Proceedings of the IEEE*, 63(4):678–692, 1975.
- [863] Cannon M. Blind deconvolution of spatially invariant image blurs with phase. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 24(1):58–63, 1976.
- [864] Pohlig SC, Lim JS, Oppenheim AV, Dudgeon DE, and Filip AE. New technique for blind deconvolution. In *Proceedings of SPIE Vol. 207: Applications of Digital Image Processing III*, pages 119–124, 1979.
- [865] Lim JS. Image restoration by short space spectral subtraction. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 28(2):191–197, 1980.
- [866] Ghiglia DC and Romero LA. Robust two-dimensional weighted and unweighted phase unwrapping that uses fast transforms and iterative methods. *Journal of the Optical Society of America A*, 11(1):107–117, 1994.
- [867] Ching NH, Rosenfeld D, and Braun M. Two-dimensional phase unwrapping using a minimum spanning tree algorithm. *IEEE Transactions on Image Processing*, 1(3):355–365, 1992.

- [868] Secilla JP, Garcia N, and Carrascosa JL. Evaluation of two-dimensional unwrapped phase averaging for the processing of quasi-periodical noisy images. *Signal Processing IV: Theories and Applications*, pages 255–258, 1988.
- [869] Hedley M and Rosenfeld D. A new two-dimensional phase unwrapping algorithm for MRI images. *Magnetic Resonance in Medicine*, 24:177–181, 1992.
- [870] Taxt T. Restoration of medical ultrasound images using two-dimensional homomorphic deconvolution. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 42(4):543–554, 1995.
- [871] Tribolet JM. Applications of short-time homomorphic signal analysis to seismic wavelet estimation. *Geoexploration*, 16:75–96, 1978.
- [872] Oppenheim AV, editor. *Applications of Digital Signal Processing*. Prentice Hall, Englewood Cliffs, NJ, 1978.
- [873] Dudgeon DE. The computation of two-dimensional cepstra. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 25(6):476–484, 1977.
- [874] Bandari E and Little JJ. Visual echo analysis. In *Proceedings of the 4th International Conference on Computer Vision*, pages 202–225, Berlin, Germany, May 1993.
- [875] Skoneczny S. Homomorphic 2-D filtering in computer simulation of image degradation. *AMSE Review*, 14(2):31–40, 1990.
- [876] Lee JK, Kabrisky M, Oxley ME, Rogers SK, and Ruck DW. The complex cepstrum applied to two-dimensional images. *Pattern Recognition*, 26(10):1579–1592, 1993.
- [877] Yeshurun Y and Schwartz EL. Cepstral filtering on a columnar image architecture: a fast algorithm for binocular stereo segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7):759–767, 1989.
- [878] Wahl FM. *Digital Image Signal Processing*. Artech House, Norwood, MA, 1987.
- [879] Biemond J, Lagendijk RL, and Mersereau RM. Iterative methods for image deblurring. *Proceedings of the IEEE*, 78(5):856–883, 1990.
- [880] Welch PD. The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio and Electroacoustics*, AU-15:70–73, 1967.
- [881] Bingham C, Godfrey MD, and Tukey JW. Modern techniques of power spectrum estimation. *IEEE Transactions on Audio and Electroacoustics*, AU-15(2):56–66, 1967.
- [882] Rader CM. An improved algorithm for high speed autocorrelation with applications to spectral estimation. *IEEE Transactions on Audio and Electroacoustics*, AU-18(4):439–441, 1970.
- [883] Angel ES and Jain AK. Restoration of images degraded by spatially varying point spread functions by a conjugate gradient method. *Applied Optics*, 17:2186–2190, 1978.
- [884] Strickland RN. Transforming images into block stationary behavior. *Applied Optics*, 22(10):1462–1473, May 1983.

- [885] Trussell HJ and Hunt BR. Image restoration of space-variant blurs by sectioned methods. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(6):608–609, 1978.
- [886] Rajala SA and De Figueiredo RJP. Adaptive nonlinear image restoration by a modified Kalman filtering approach. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 29(5):1033–1042, 1981.
- [887] Kalman RE. A new approach to linear filtering and prediction problems. *Transactions of the American Society of Mechanical Engineers: Journal of Basic Engineering*, 82:35–45, 1960.
- [888] Kalman RE and Bucy RS. New results in linear filtering and prediction theory. *Transactions of the American Society of Mechanical Engineers: Journal of Basic Engineering*, 83:95–108, 1961.
- [889] Sage AP and Melsa JL. *Estimation Theory with Applications to Communications and Control*. McGraw-Hill, New York, NY, 1971.
- [890] Grewal MS and Andrews AP. *Kalman Filtering: Theory and Practice Using MATLAB*. Wiley Interscience, New York, NY, 2nd edition, 2001.
- [891] Woods JW and Radewan CH. Kalman filtering in two dimensions. *IEEE Transactions on Information Theory*, IT-23(4):473–482, 1977.
- [892] Woods JW. Correction to “Kalman filtering in two dimensions”. *IEEE Transactions on Information Theory*, IT-23(5):628, 1979.
- [893] Woods JW and Ingle VK. Kalman filtering in two dimensions: Further results. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP-29(2):188–197, 1981.
- [894] Tekalp AM, Kaufman H, and Woods J. Edge-adaptive Kalman filtering for image restoration with ringing suppression. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37:892–898, 1989.
- [895] Tekalp AM and Pavlovic G. Space-variant and color image restoration using Kalman filtering. In *Proceedings of IEEE International Symposium on Circuits and Systems*, pages 1029–1031, 1989.
- [896] Tekalp AM, Kaufman H, and Woods J. Model-based segmentation and space-variant restoration of blurred images by decision-directed filtering. *Signal Processing*, 15:259–269, 1988.
- [897] Adams R. The scintillation gamma camera. In Williams LE, editor, *Nuclear Medical Physics*, pages 89–154. CRC Press, Boca Raton, FL, 1987.
- [898] Rosenthal MS, Cullom J, Hawkins W, Moore SC, Tsui BMW, and Yester M. Quantitative SPECT imaging: A review and recommendations by the Focus Committee of the Society of Nuclear Medicine Computer and Instrumentation Council. *Journal of Nuclear Medicine*, 36:1489–1513, 1995.
- [899] Brown BH, Smallwood RH, Barber DC, Lawford PV, and Hose DR. *Medical Physics and Biomedical Engineering*. Institute of Physics Publishing, Bristol, UK, 1999.
- [900] Todd-Pokropek A. Image processing in nuclear medicine. *IEEE Transactions on Nuclear Science*, 27:1080–1094, 1980.

- [901] Pizer SM and Todd-Pokropek A. Improvement of scintigrams by computer processing. *Seminars in Nuclear Medicine*, VIII(2):125–146, 1978.
- [902] Jaszczak RJ, Coleman RE, and Lim CB. SPECT: Single-photon emission computed tomography. *IEEE Transactions on Nuclear Science*, 27:1137–1153, 1980.
- [903] Todd-Pokropek AE, Zurowski S, and Soussaline F. Non-uniformity and artifact creation in emission tomography. *Journal of Nuclear Medicine*, 21:38–45, 1980.
- [904] Sorenson JA and Phelps ME. *Physics in Nuclear Medicine*. Grune & Stratton, New York, NY, 1987.
- [905] Jaszczak RJ, Coleman RE, and Whitehead FR. Physical factors affecting quantitative measurements using camera-based single photon computed tomography (SPECT). *IEEE Transactions on Nuclear Science*, 28:69–80, 1981.
- [906] Wicks R and Blau M. Effects of spatial distortion on Anger camera field-uniformity correction. *Journal of Nuclear Medicine*, 20:252–265, 1979.
- [907] Muehllehner G, Colsher JG, and Stoub EW. Correction for field nonuniformity in scintillation cameras through removal of spatial distortion. *Journal of Nuclear Medicine*, 21:771–779, 1980.
- [908] Chandra R. *Introductory Physics of Nuclear Medicine*. Lea and Febiger, Philadelphia, PA, 1987.
- [909] Croft BY. *Single-Photon Emission Computed Tomography*. Year Book Medical Publishers, Chicago, IL, 1986.
- [910] Jaszczak RJ, Greer KL, and Floyd Jr. CE. Improved SPECT quantitation using compensation for scattered photons. *Journal of Nuclear Medicine*, 25:893–906, 1984.
- [911] Egbert SD and May RS. An integral-transport method for Compton scatter correction in emission computed tomography. *IEEE Transactions on Nuclear Science*, 27:543–551, 1980.
- [912] Chang T. A method for attenuation correction in radionuclide computed tomography. *IEEE Transactions on Nuclear Science*, 25:638–643, 1978.
- [913] Axelsson B, Msaki P, and Israelsson A. Subtraction of Compton scattered photons in single-photon emission computerized tomography. *Journal of Nuclear Medicine*, 25:490–494, 1984.
- [914] Floyd Jr. CE, Jaszczak RJ, and Harris CC. Monte Carlo evaluation of Compton scatter compensation by deconvolution in SPECT. *Journal of Nuclear Medicine*, 25:71, 1984.
- [915] Floyd Jr. CE, Jaszczak RJ, Greer KL, and Coleman RE. Deconvolution of Compton scatter in SPECT. *Journal of Nuclear Medicine*, 26:403–408, 1985.
- [916] Buvat I, Benali H, Todd-Pokropek A, and Di Paola R. Scatter correction in scintigraphy: the state of the art. *European Journal of Nuclear Medicine*, 21:675–694, 1994.
- [917] Ljungberg M, King MA, Hademenos GJ, and Strand SE. Comparison of four scatter correction methods using Monte Carlo simulated source distributions. *Journal of Nuclear Medicine*, 35:143–151, 1994.

- [918] Rogers L and Clinthorne NH. Single photon emission computed tomography (SPECT). In Williams LE, editor, *Nuclear Medical Physics*, pages 1–48. CRC Press, Boca Raton, FL, 1987.
- [919] Tsui ET and Budinger TF. A stochastic filter for transverse section reconstruction. *IEEE Transactions on Nuclear Science*, 26:2687–2690, 1979.
- [920] Kay DB and Keyes Jr. JW. First-order corrections for absorption and resolution compensation in radionuclide Fourier tomography. *Journal of Nuclear Medicine*, 16:540–551, 1975.
- [921] Sorenson JA. Quantitative measurement of radiation in vivo by whole body counting. In Hine GJ and Sorenson JA, editors, *Instrumentation in Nuclear Medicine*, pages 311–365. Academic, New York, NY, 1984.
- [922] Walters TE, Simon W, Chesler DA, and Correia J. Attenuation correction in gamma emission computed tomography. *Journal of Computer Assisted Tomography*, 5:89–102, 1981.
- [923] Budinger TF and Gullberg GT. Three-dimensional reconstruction in nuclear medicine emission imaging. *IEEE Transactions on Nuclear Science*, 21:2–16, 1974.
- [924] Gullberg GT and Budinger TF. The use of filtering methods to compensate for constant attenuation in single-photon emission computed tomography. *IEEE Transactions on Biomedical Engineering*, 28:142–153, 1981.
- [925] Moore SC, Brunelle JA, and Kirsch CM. An iterative attenuation correction for a single photon scanning multidetector tomography system. *Journal of Nuclear Medicine*, 22:65–76, 1981.
- [926] Faber TL, Lewis MH, and Corbett JR. Attenuation correction for SPECT: An evaluation of hybrid approaches. *IEEE Transactions on Medical Imaging*, 3:101–109, 1984.
- [927] Censor Y, Gustafson DE, Lent A, and Tuy H. A new approach to the emission computed tomography problem: Simultaneous calculation of attenuation and activity coefficients. *IEEE Transactions on Nuclear Science*, 26:146–154, 1979.
- [928] King MA, Tsui BMW, Pan TS, Glick SJ, and Soares EJ. Attenuation compensation for cardiac single-photon emission computed tomographic imaging: Part 2. Attenuation compensation algorithms. *Journal of Nuclear Cardiology*, 3:55–63, 1996.
- [929] Boardman AK. Constrained optimization and its application to scintigraphy. *Physics in Medicine and Biology*, 24:363–371, 1979.
- [930] Madsen MT and Park CH. Enhancement of SPECT images by Fourier filtering the projection image set. *Journal of Nuclear Medicine*, 26:395–402, 1985.
- [931] Yanch JC, Flower MA, and Webb S. A comparison of deconvolution and windowed subtraction techniques for scatter compensation in SPECT. *IEEE Transactions on Medical Imaging*, 7:13–20, 1988.
- [932] Miller TR and Rollins ES. A practical method of image enhancement by interactive digital filtering. *Journal of Nuclear Medicine*, 26:1075–1080, 1985.

- [933] Cordier S, Biraud Y, Champaville A, and Voutay M. A study of the application of a deconvolution method to scintigraphy. *Physics in Medicine and Biology*, 24:577–582, 1979.
- [934] Maeda J and Murata K. Digital restoration of scintigraphic images by a two-step procedure. *IEEE Transactions on Medical Imaging*, 6:320–324, 1987.
- [935] Rangayyan RM, Boulfelfel D, Hahn LJ, and Kloiber R. Two-dimensional and three-dimensional restoration of SPECT images. *Medical and Life Sciences Engineering (Journal of the Biomedical Engineering Society of India)*, 14:82–94, 1997.
- [936] Penney BC, Glick SJ, and King MA. Relative importance of the error sources in Wiener restoration of scintigrams. *IEEE Transactions on Medical Imaging*, 9:60–70, 1990.
- [937] Iwata S, Yoshida C, and Nakajima M. Correction method for collimator effect of ECT. *Systems and Computers in Japan*, 17:43–50, 1986.
- [938] Raff U, Stroud DN, and Hendee WR. Improvement of lesion detection in scintigraphic images by SVD techniques for resolution recovery. *IEEE Transactions on Medical Imaging*, 5:35–44, 1986.
- [939] Stritzke P, King MA, Vaknine R, and Goldsmith SL. Deconvolution using orthogonal polynomials in nuclear medicine: A method for forming quantitative functional images from kinetic studies. *IEEE Transactions on Medical Imaging*, 9:11–23, 1990.
- [940] Doré S, Kearney RE, and de Guise J. Quantitative assessment of CT PSF isotropicity and isoplanicity. In *Proceedings of the Canadian Medical and Biological Engineering Conference*, pages 31–32. Canadian Medical and Biological Engineering Society, Winnipeg, MB, 1990.
- [941] Doré S, Kearney RE, and de Guise J. Experimental determination of CT point spread function. In *Proceedings of the 11th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 620–621, Seattle, WA, 1989.
- [942] Glick SJ, King MA, and Penney BC. Characterization of the modulation transfer function of discrete filtered backprojection. *IEEE Transactions on Medical Imaging*, 8:203–213, 1989.
- [943] Glick SJ, King MA, and Knesaurek K. An investigation of the 3D modulation transfer function used in 3D post-reconstruction restoration filtering of SPECT imaging. In Ortendahl DA and Llacer J, editors, *Information Processing in Medical Imaging*, pages 107–122. Wiley-Liss, New York, NY, 1991.
- [944] Msaki P, Axelsson B, Dahl CM, and Larsson SA. Generalized scatter correction method in SPECT using point scatter distribution functions. *Journal of Nuclear Medicine*, 28:1861–1869, 1987.
- [945] Coleman M, King MA, Glick SJ, Knesaurek K, and Penney BC. Investigation of the stationarity of the modulation transfer function and the scatter fraction in conjugate view SPECT restoration filtering. *IEEE Transactions on Nuclear Science*, 36:969–971, 1989.

- [946] King MA, Coleman M, Penney BC, and Glick SJ. Activity quantitation in SPECT: A study of prereconstruction Metz filtering and use of the scatter degradation factor. *Medical Physics*, 18(2):184–189, 1991.
- [947] Glick SJ, King MA, Knesaurek K, and Burbank K. An investigation of the stationarity of the 3D modulation transfer function of SPECT. *IEEE Transactions on Nuclear Science*, 36:973–977, 1989.
- [948] Boulfelfel D, Rangayyan RM, Hahn LJ, and Kloiber R. Three-dimensional restoration of single photon emission computed tomography images using the Kalman-Metz filter. In *Computerized Tomography: Proceedings of the Fourth International Symposium, Novosibirsk, Siberia, Russia, 10-14 August 1993*, pages 98–105. VSP BV, Utrecht, The Netherlands, 1995.
- [949] Sawchuk AA. Space-variant image restoration by coordinate transformations. *Journal of the Optical Society of America*, 64:138–144, 1974.
- [950] Pizer SM and Todd-Pokropek AE. Noise character in processed scintigrams. In *Proceedings of 4th International Conference on Information Processing in Scintigraphy*, pages 1–16, 1975.
- [951] Fredric HJ. On the use of windows for harmonic analysis with the discrete Fourier transform. *Proceedings of the IEEE*, 66:51–83, 1978.
- [952] Frieden BR. Image enhancement and restoration. In Huang TS, editor, *Picture Processing and Digital Filtering*, volume 6, pages 88–93. Springer-Verlag, Berlin, Germany, 1978.
- [953] Aubert G and Kornprobst P. *Mathematical Problems in Image Processing*. Springer, New York, NY, 2002.
- [954] Optical Society of America. *Digest of Topical Meeting on Signal recovery and synthesis with incomplete information and partial constraints*. Optical Society of America, Incline Village, NV, 1983.
- [955] Optical Society of America. *Signal Recovery: Journal of the Optical Society of America, Volume 73, Number 11*. Optical Society of America, New York, NY, November 1983.
- [956] *SPIE International Symposium on Medical Imaging: PACS and Imaging Informatics*. <http://www.spie.org/conferences/calls/05/mi/>, accessed June 2004.
- [957] MathWorld, www.mathworld.wolfram.com/GrayCode.html. *Gray Code*, accessed February 2004.
- [958] American Standard Code for Information Interchange, www.asciitable.com/ascii2.html. *ASCII Table and Description*, accessed February 2004.
- [959] Huffman DA. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*, 40(10):1098–1101, 1952.
- [960] Langdon Jr. GG. An introduction to arithmetic coding. *IBM Journal of Research and Development*, 28(2):135–149, 1984.
- [961] Rissanen J and Langdon Jr. GG. Arithmetic coding. *IBM Journal of Research and Development*, 23(2):149–162, 1979.

- [962] Langdon Jr. GG and Rissanen J. Compression of black-white images with arithmetic coding. *IEEE Transactions on Communications*, COM-29:858–867, 1981.
- [963] Witten IH, Neal RM, and Cleary JG. Arithmetic coding for data compression. *Communications of the ACM*, 30(6):520–540, 1987.
- [964] Pennebaker WB, Mitchell JL, Langdon Jr. GG, and Arps RB. An overview of the basic principles of the Q-coder adaptive binary arithmetic coder. *IBM Journal of Research and Development*, 32:717–726, 1988.
- [965] Rabbani M and Jones PW. Image compression techniques for medical diagnostic imaging systems. *Journal of Digital Imaging*, 4(2):65–78, 1991.
- [966] Ziv J and Lempel A. A universal algorithm for sequential data compression. *IEEE Transactions on Information Theory*, IT-23(3):337–343, 1977.
- [967] Welch TA. A technique for high-performance data compression. *IEEE Computer*, pages 8–19, June 1984.
- [968] Lo SC, Krasner B, and Mun SK. Noise impact on error-free image compression. *IEEE Transactions on Medical Imaging*, 9(2):202–206, 1990.
- [969] Dainty JC. *Image Science*. Academic, London, UK, 1974.
- [970] Ahmed N, Natarajan T, and Rao KR. Discrete cosine transform. *IEEE Transactions on Computer*, C-23:90–93, 1974.
- [971] Ahmed N and Rao KR. *Orthogonal Transforms for Digital Signal Processing*. Springer-Verlag, New York, NY, 1975.
- [972] Kuduvalli GR and Rangayyan RM. Error-free transform coding by maximum-error-limited quantization of transform coefficients. In *Proceedings of SPIE on Visual Communications and Image Processing*, volume 1818, pages 1458–1461. SPIE, 1992.
- [973] Berger T. *Rate Distortion Theory*. Prentice Hall, Englewood Cliffs, NJ, 1971.
- [974] Helstrom CW. *Probability and Stochastic Processes for Engineers*. Macmillan, New York, NY, 1991.
- [975] Reininger RC and Gibson JD. Distribution of the two-dimensional DCT coefficients for images. *IEEE Transactions on Communications*, COM-31:835–839, 1983.
- [976] Cox JR, Moore SM, Blaine GJ, Zimmerman JB, and Wallace GK. Optimization of trade-offs in error-free image transmission. In *Proceedings of SPIE Vol. 1091. Medical imaging III: Image capture and display*, pages 19–30, 1989.
- [977] Wang L and Goldberg M. Progressive image transmission by transform coefficient residual error quantization. *IEEE Transactions on Communications*, 36:75–76, 1988.
- [978] Roos P, Viergever MA, VanDijke MCA, and Peters JH. Reversible intraframe compression of medical images. *IEEE Transactions on Medical Imaging*, 7:328–336, December 1988.
- [979] Strobach P. *Linear Prediction Theory: A Mathematical Basis of Adaptive Systems*. Springer-Verlag, New York, NY, 1990.

- [980] Maragos MA, Schafer RW, and Mersereau RM. Two-dimensional linear prediction and its application to adaptive predictive coding of images. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 32:1213–1229, 1984.
- [981] Wax M and Kailath T. Efficient inversion of Toeplitz-block Toeplitz matrix. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 31:1218–1221, 1983.
- [982] Kalouptsidis N, Carayannis C, and Manolakis D. Fast algorithms for block Toeplitz matrices with Toeplitz entries. *Signal Processing*, 6:77–81, 1984.
- [983] Therrien CW and El-Shaer HT. A direct algorithm for computing 2-D AR spectrum estimates. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37:1795–1797, 1989.
- [984] Therrien CW and El-Shaer HT. Multichannel 2-D AR spectrum estimation. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37:1798–1800, 1989.
- [985] Therrien CW. Relations between 2-D and multichannel linear prediction. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 29:454–457, 1981.
- [986] Strand ON. Multichannel complex maximum entropy (autoregressive) spectral analysis. *IEEE Transactions on Automatic Control*, 22:634–640, 1977.
- [987] Levinson N. The Wiener rms (root mean square) error criterion in filter design and prediction. *Journal of Mathematical Physics*, 25:261–278, 1947.
- [988] Wiggins RA and Robinson EA. Recursive solution to the multichannel filtering problem. *Journal of Geophysical Research*, 70:1885–1891, 1965.
- [989] Graham A. *Kronecker Products and Matrix Calculus with Applications*. Wiley, Chichester, UK, 1981.
- [990] Burg JP. *A new analysis technique for time series data*. NATO Advance Study Institute on Signal Processing, Enschede, The Netherlands, 1968.
- [991] Cioffi JM and Kailath T. Fast recursive-least-squares transversal filters for adaptive filtering. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 32:304–337, 1984.
- [992] Boutalis YS, Kollias SD, and Carayannis G. A fast multichannel approach to adaptive image estimation. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37:1090–1098, 1989.
- [993] Scott KE. *Adaptive Equalization and Antenna Diversity in Digital Radio*. PhD thesis, Department of Electrical and Computer Engineering, University of Calgary, Calgary, Alberta, Canada, May 1991.
- [994] Ohki M and Hashiguchi S. Two-dimensional LMS adaptive filters. *IEEE Transactions on Consumer Electronics*, 37(1):66–72, 1991.
- [995] Alexander ST and Rajala SA. Image compression results using the LMS adaptive algorithm. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 33:712–714, 1985.

- [996] Aiazzi B, Alparone L, and Baronti S. Trends in lossless image compression by adaptive prediction. *Electronic Imaging: SPIE's International Technical Group Newsletter*, 13(1):5,9, January 2003.
- [997] Peano G. Sur une courbe, qui remplit toute une aire plane. *Mathematische Annalen*, 36:157–160, 1890.
- [998] Hilbert D. Ueber die stetige abbildung einer linie auf ein flächenstück. *Mathematische Annalen*, 38:459–460, 1891.
- [999] Provine JA. Peanoscanning for image compression. Master's thesis, Department of Electrical and Computer Engineering, University of Calgary, Calgary, Alberta, Canada, December 1992.
- [1000] Provine JA and Rangayyan RM. Lossless compression of Peanoscanned images. *Journal of Electronic Imaging*, 3(2):176–181, 1994.
- [1001] Zhang YQ, Loew MH, and Pickholtz RL. A methodology for modeling the distributions of medical images and their stochastic properties. *IEEE Transactions on Medical Imaging*, 9(4):376–382, 1990.
- [1002] Zhang YQ, Loew MH, and Pickholtz RL. A combined-transform coding (CTC) scheme for medical images. *IEEE Transactions on Medical Imaging*, 11(2):196–202, 1992.
- [1003] Lempel A and Ziv J. Compression of two-dimensional data. *IEEE Transactions on Information Theory*, IT-32(1):2–8, 1986.
- [1004] Bially T. Space-filling curves: Their generation and their application to bandwidth reduction. *IEEE Transactions on Information Theory*, IT-15(6):658–664, 1969.
- [1005] Moore EH. On certain crinkly curves. *Transactions of the American Mathematical Society*, 1:72–90, 1900.
- [1006] Witten IH and Neal RM. Using Peano curves for bilevel display of continuous-tone images. *IEEE Computer Graphics*, pages 47–52, May 1982.
- [1007] Witten IH and Wyvill B. On the generation and use of space-filling curves. *Software – Practice and Experience*, 13:519–525, 1983.
- [1008] International Telegraph and Telephone Consultative Committee (CCITT). Progressive bi-level image compression. *Recommendation T.82*, 1993.
- [1009] International Organization for Standards/ International Electrotechnical Commission (ISO/IEC). Progressive bi-level image compression. *International Standard 11544*, 1993.
- [1010] Hampel H, Arps RB, Chamzas C, Dellert D, Duttweiler DL, Endoh T, Equitz W, Ono F, Pasco R, Sebestyen I, Starkey CJ, Urban SJ, Yamazaki Y, and Yoshida T. Technical features of the JBIG standard for progressive bi-level image compression. *Signal Processing: Image Communication*, 4:103–111, 1992.
- [1011] JPEG and JBIG website. <http://www.jpeg.org/>, accessed May 2004.
- [1012] Wallace GK. The JPEG still picture compression standard. *Communications of the ACM*, 34(4):30–44, 1991.
- [1013] Sikora T. MPEG video webpage. <http://www.bs.hhi.de/mpeg-video/>, accessed May 2004.

- [1014] Chiariglione L. MPEG and multimedia communications. *IEEE Transactions on Circuits and Systems for Video Technology*, 7(1):5–18, 1997.
- [1015] National Electrical Manufacturers Association, Rosslyn, VA. *Digital Imaging and Communications in Medicine (DICOM) PS 3.1-2003*, available at <http://medical.nema.org/>, accessed May 2004.
- [1016] American College of Radiology and National Electrical Manufacturers Association, Washington, DC. *ACR/NEMA Standards Publication PS 2: Data Compression Standard*, 1989.
- [1017] American College of Radiology and National Electrical Manufacturers Association, Washington, DC. *ACR/NEMA Standards Publication No. 300: Digital Imaging and Communications*, 1989.
- [1018] Wang Y, Best DE, Hoffman JG, Horii SC, Lehr JL, Lodwick GS, Morse RR, Murphy LL, Nelson OL, Perry J, Thompson BG, and Zielonka JS. ACR-NEMA digital imaging and communications standards: minimum requirements. *Radiology*, 166:529–532, February 1988.
- [1019] Rao KR and Huh Y. JPEG 2000. In *Proceedings of IEEE Region 8 International Symposium on Video/Image Processing and Multimedia Communications*, pages 1–6, Zadar, Croatia, June 2002.
- [1020] Sung MM, Kim HJ, Kim EK, Kwak JY, Yoo JK, and Yoo HS. Clinical evaluation of JPEG2000 compression algorithm for digital mammography. *IEEE Transactions on Nuclear Science*, 49(3):827–832, 2002.
- [1021] Fossbakk E, Manzanares P, Yago JL, and Perkis A. An MPEG-21 framework for streaming media. In *Proceedings of IEEE 4th Workshop on Multimedia Signal Processing*, pages 147–152, October 2001.
- [1022] Vaisey J and Gersho A. Image compression with variable block size segmentation. *IEEE Transactions on Signal Processing*, 40(8):2040–2060, August 1992.
- [1023] Leou FC and Chen YC. A contour-based image coding technique with its texture information reconstructed by polyline representation. *Signal Processing*, 25:81–89, 1991.
- [1024] Lopes RD and Rangayyan RM. Lossless volumetric data compression via decomposition based upon region growing. In *Proceedings SPIE 3658: Medical Imaging 1996 – Physics of Medical Imaging*, pages 427–435, San Diego, CA, February 1999.
- [1025] *The bzip2 and libbzip2 official home page*. <http://sources.redhat.com/bzip2/>, accessed May 2004.
- [1026] Acha B, Serrano C, Rangayyan RM, and Roa LM. Lossless compression algorithm for colour images. *Electronics Letters*, 35(3):214–215, 4 February 1999.
- [1027] Serrano C, Acha B, Rangayyan RM, and Roa LM. Segmentation-based lossless compression of burn wound images. *Journal of Electronic Imaging*, 10(3):720–726, 2001.
- [1028] Shen L and Rangayyan RM. Lossless compression of continuous-tone images by combined inter-bit-plane decorrelation and JBIG coding. *Journal of Electronic Imaging*, 6(2):198–207, 1997.

- [1029] Arps RB and Truong TK. Comparison of international standards for lossless still image compression. *Proceedings of the IEEE*, 82(6):889–899, 1994.
- [1030] Wang Y. A set of transformations for lossless image compression. *IEEE Transactions on Image Processing*, 4(5):677–679, 1995.
- [1031] Rabbani M and Melnychuck PW. Conditioning contexts for the arithmetic coding of bit planes. *IEEE Transactions on Signal Processing*, 40(1):232–236, 1992.
- [1032] RICOH California Research Center. *CREW lossless/lossy image compression – Contribution to ISO/IEC JTC 1.29.12*, June 1995.
- [1033] Zandi A, Allen JD, Schwartz EL, and Boliek M. CREW: Compression with reversible embedded wavelets. In *Proceedings of IEEE Data Compression Conference*, pages 212–221, Snowbird, UT, March 1995.
- [1034] Rabbani M and Jones PW. *Digital Image Compression Techniques*. SPIE Optical Engineering Press, Bellingham, WA, 1991.
- [1035] Bacharin GP. On a statistical estimate for the entropy of a sequence of independent random variables. *Theory of Probability and its Applications*, 4:333–336, 1959.
- [1036] Andrus WS and Bird KT. Teleradiology: evolution through bias to reality. *Chest*, 62:655–657, 1972.
- [1037] Carey LS. Teleradiology: part of a comprehensive telehealth system. *Radiologic Clinics of North America*, 23(2):357–362, 1985.
- [1038] House M. Use of telecommunications to meet health needs of rural, remote, and isolated communities. In Rangayyan RM, editor, *Telecommunication for Health Care: Telemetry, Teleradiology, and Telemedicine; Proceedings of SPIE Vol. 1355*, pages 2–9, Bellingham, WA, 1990. SPIE.
- [1039] Brown JHU. *Telecommunication for Health Care*. CRC Press, Boca Raton, FL, 1985.
- [1040] Rangayyan RM, editor. *Telecommunication for Health Care: Telemetry, Teleradiology, and Telemedicine; Proceedings of SPIE Vol. 1355*. SPIE, Bellingham, WA, 1990.
- [1041] Kuduvalli GR, Rangayyan RM, and Desautels JEL. High-resolution digital teleradiology: A perspective. *Journal of Digital Imaging*, 4(4):251–261, 1991.
- [1042] Allman R. *Potential contribution of teleradiology to the management of military. A technical report: Radiologist resources in 'Military Medicine'*. Radiological Society of North America (RSNA), Chicago, IL, December, 1983.
- [1043] Goeringer F, Mun SK, and Kerlin BD. Digital medical imaging: implementation strategy for the defense medical establishment. In *Proceedings of SPIE Vol. 1093. Medical Imaging III: PACS system design and evaluation*, pages 429–437, 1989.
- [1044] Drew PG. *Market Study for a High-performance Teleradiology (in Alberta, Canada)*. Drew Consultants, Carlisle, MA, 1989.
- [1045] Webber MM and Corbus HF. Image communications by telephone. *Journal of Nuclear Medicine*, 13:379–381, 1972.

- [1046] Jelasco DV, Southworth G, and Purcell LH. Telephone transmission of radiographic images. *Radiology*, 127:147–149, 1978.
- [1047] Barnes GT, Johnson GA, and Staab EV. *Teleradiology: Fundamental considerations and clinical applications*. Unpublished lecture notes, 1989.
- [1048] Arenson RL, Sheshadri SB, Kundel HA, DeSimone D, Van der Voorde F, Geftter WB, Epstein DM, Miller WT, Aronchick JM, Simson MB, Lancken PN, Khalsa S, Brikman I, Davey M, and Brisbon N. Clinical evaluation of a medical image management system for chest images. *American Journal of Roentgenology*, 150:55–59, January, 1988.
- [1049] Carey LS, O'Connor BD, Bach DB, Hobbs BB, Hutton LC, Lefcoe MS, Lynos RO, Munro TG, Paterson RG, Rankin RN, and Rutt BK. Digital teleradiology: Seaforth – London network. *Journal of Canadian Association of Radiology*, 40:71–74, 1989.
- [1050] Gershon-Cohen J and Cooley AG. Telegnosis. *Radiology*, 55:582–587, 1950.
- [1051] Jutras A. Teleroentgen diagnosis by means of videotape recording (Editorial). *American Journal of Roentgenology*, 82:1099–1102, 1959.
- [1052] Andrus WS, Dreyfuss JR, Jaffer F, and Bird KT. Interpretation of roentgenograms via interactive television. *Radiology*, 116:25–31, 1975.
- [1053] Murphy RL, Barber D, Broadhurst A, and Bird KT. Microwave transmission of chest roentgenograms. *American Review of Respiratory Diseases*, 102:771–777, 1972.
- [1054] Webber MM, Wilk S, Pirruccello R, and Aiken J. Telecommunication of images in the practice of diagnostic radiology. *Radiology*, 109:71–74, 1973.
- [1055] Lester RG, O'Foghludha F, Porter F, Friedman DS, and Pedolsky HR. Transmission of radiologic information by satellite. *Radiology*, 109:731–732, 1973.
- [1056] Carey LS, Russell ES, Johnson EE, and Wilkins WW. Radiologic consultation to a remote Canadian hospital using Hermes spacecraft. *Journal of Canadian Association of Radiology*, 30:12–20, 1979.
- [1057] James JJ, Grabowski W, and Mangelsdorff AD. The transmission and interpretation of emergency department radiographs. *Annals of Emergency Medicine*, 11:404–408, 1982.
- [1058] Steckel RJ. Daily x-ray rounds in a large teaching hospital using high-resolution closed-circuit television. *Radiology*, 116:25–31, 1975.
- [1059] Page G, Gregoire A, Galand C, Sylvestre J, Chahlaoul J, Faut eux P, Dussault R, Seguin R, and Roberge F. Teleradiology in Northern Québec. *Radiology*, 140:361–366, 1981.
- [1060] Kretz F and Nasse D. Digital television: transmission and coding. *Proceedings of the IEEE*, 73:575–591, 1985.
- [1061] Kuni CC. *Introduction to Computers and Digital Processing in Medical Imaging*. Year Book Medical Publishers, Chicago, IL, 1988.
- [1062] Nudelman S. Historical perspectives on photoelectronic digital radiology. In James AE, Anderson JH, and Higgins CB, editors, *Digital Image Processing in Radiology*, pages 1–27. Williams & Wilkins, Baltimore, MD, 1985.

- [1063] Gayler BW, Gitlin JN, Rappaport W, Skinner FL, and Cerva J. Teleradiology: An evaluation of a microcomputer-based system. *Radiology*, 140:355–360, 1981.
- [1064] Kagetsu NJ, Zulauf DRP, and Ablow RC. Clinical trial digital teleradiology in the practice of emergency room radiology. *Radiology*, 165:551–554, 1987.
- [1065] Gitlin JN. Teleradiology. *Radiologic Clinics of North America*, 24:55–68, 1986.
- [1066] Curtis DJ, Gayler BW, Gitlin JN, and Harrington MB. Teleradiology: results of a field trial. *Radiology*, 149:415–418, 1983.
- [1067] Rasmussen W, Stevens I, Gerber FH, and Kuhlman JA. Teleradiology via the naval Remote Medical Diagnosis System (RMDS). In *Proceedings of SPIE Vol. 318 (Part I). Picture Archiving and Communication Systems (PACS) for Medical Applications*, pages 174–181, 1982.
- [1068] Skinner FL, Cerva J, Kerlin B, and Millstone T. The teleradiology field demonstration. In *Proceedings of SPIE Vol. 318 (Part I). Picture Archiving and Communication Systems (PACS) for Medical Applications*, pages 168–173, 1982.
- [1069] Gordon R, Rangayyan RM, Wardrop DH, and Beeman TM. Improving image quality in teleradiology and tele-computed tomography. In *Proceedings of the IEEE Systems, Man, and Cybernetics Society Conference*, pages 908–913, Bombay and New Delhi, India, December/January 1983/84.
- [1070] Rangayyan RM and Gordon R. Computed tomography from ordinary radiographs for teleradiology. *Medical Physics*, 10:687–690, 1983.
- [1071] Rangaraj MR and Gordon R. Computed tomography for remote areas via teleradiology. In *Proceedings of SPIE Vol. 318 on Picture Archival and Communication Systems for Medical Applications*, pages 182–185, Newport Beach, CA, January 1982.
- [1072] DiSantis DJ, Cramer MS, and Scatarige JC. Excretory urography in the emergency department: utility of teleradiology. *Radiology*, 164:363–364, 1987.
- [1073] Cox GG, Cook LT, McMillan JH, Rosenthal SJ, and Dwyer III SJ. *High-resolution 2560 × 2048 × 12 bit digital displays for chest radiography - A comparison with conventional film and digital hardcopy*. University of Kansas Medical Center, Kansas City, KA, 1988.
- [1074] Batnitzky S, Rosenthal SJ, Siegel EL, Wetzel LH, Murphey MD, Cox GG, McMillan JH, Templeton AW, and Dwyer III SJ. Teleradiology: an assessment. *Radiology*, 177:11–17, 1990.
- [1075] Gillespy T, Staab EV, Staab EW, and Lawrence E. Electronic imaging in a teaching hospital intensive care unit: evaluation of the clinical review system. *Journal of Digital Imaging*, 3:124–128, 1990.
- [1076] Lo SC, Gaskill JW, Mun SK, and Krasner BH. Contrast information of digital imaging in laser film digitizer and display monitor. *Journal of Digital Imaging*, 3:119–123, May 1990.

- [1077] Yip K, Lubinsky AR, Whiting BR, Muka E, and Cocher TE. Performance analysis of medical x-ray film digitizers. In *Proceedings of SPIE Vol. 1231. Medical Imaging IV: Image Formation*, pages 508–525, 1990.
- [1078] Slasky BS, Gur D, Costa-Greco MA, and Harris KM. Receiver operating characteristic analysis of chest image interpretation with conventional, laser-printed, and high-resolution workstation images. *Radiology*, 174:775–780, 1990.
- [1079] Proakis JG. *Digital Communications*. McGraw-Hill, New York, NY, 1989.
- [1080] Telesat Canada, Gloucester, ON, Canada. *Satellite Delay and Response Times (Product Literature)*, 1990.
- [1081] Kim Y and Horii SC, editors. *Handbook of Medical Imaging, Volume 3: Display and PACS*. SPIE Press, Bellingham, WA, 2000.
- [1082] Society for Computer Applications in Radiology (SCAR), Harrisburg, PA. *Understanding Teleradiology*, 1994.
- [1083] Society for Computer Applications in Radiology (SCAR), Harrisburg, PA. *Understanding PACS: Picture Archiving and Communications Systems*, 1994.
- [1084] Ackerman LV, Mucciardi AN, Gose EE, and Alcorn FS. Classification of benign and malignant breast tumours on the basis of 36 radiographic properties. *Cancer*, 31:342–352, 1973.
- [1085] Alto H, Rangayyan RM, Paranjape RB, Desautels JEL, and Bryant H. An indexed atlas of digital mammograms for computer-aided diagnosis of breast cancer. *Annales des Télécommunications*, 58(5-6):820–835, 2003.
- [1086] Tou JT and Gonzalez RC. *Pattern Recognition Principles*. Addison-Wesley, Reading, MA, 1974.
- [1087] Fukunaga K. *Introduction to Statistical Pattern Recognition*. Academic, San Diego, CA, 2nd edition, 1990.
- [1088] Johnson RA and Wichern DW. *Applied Multivariate Statistical Analysis*. Prentice Hall, Englewood Cliffs, NJ, 3rd edition, 1992.
- [1089] Schürmann J. *Pattern Classification – A unified view of statistical and neural approaches*. Wiley, New York, NY, 1996.
- [1090] Micheli-Tzanakou E. *Supervised and Unsupervised Pattern Recognition*. CRC Press, Boca Raton, FL, 2000.
- [1091] Neter J, Kutner MH, Nachtsheim CJ, and Wasserman W. *Applied Linear Statistical Models*. Irwin, Chicago, IL, 4th edition, 1990.
- [1092] SPSS Inc., Chicago, IL. *SPSS Advanced Statistics User's Guide*, 1990.
- [1093] SPSS Inc., Chicago, IL. *SPSS Base System User's Guide*, 1990.
- [1094] Pao YH. *Adaptive Pattern Recognition and Neural Networks*. Addison-Wesley, Reading, MA, 1989.
- [1095] Lippmann RP. An introduction to computing with neural nets. *IEEE Signal Processing Magazine*, pages 4–22, April 1987.
- [1096] Nigrin A. *Neural Networks for Pattern Recognition*. MIT, Cambridge, MA, 1993.

- [1097] Devijver PA and Kittler J. *Pattern Recognition: A Statistical Approach*. Prentice Hall, London, UK, 1982.
- [1098] Shen L, Rangayyan RM, and Desautels JEL. A knowledge-based position matching technique for mammographic calcifications. In *Proceedings of the 14th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 1936–1937, Paris, France, October 1992.
- [1099] Metz CE. Basic principles of ROC analysis. *Seminars in Nuclear Medicine*, VIII(4):283–298, 1978.
- [1100] Metz CE. ROC methodology in radiologic imaging. *Investigative Radiology*, 21:720–733, 1986.
- [1101] Swets JA and Pickett RM. *Evaluation of diagnostic systems: Methods from signal detection theory*. Academic, New York, NY, 1982.
- [1102] Dorfman DD and Alf E. Maximum likelihood estimation of parameters of signal detection theory and determination of confidence intervals – rating method data. *Journal of Mathematical Psychology*, 6:487–496, 1969.
- [1103] Fleiss JL. *Statistical Methods for Rates and Proportions*. Wiley, New York, NY, 2nd edition, 1981.
- [1104] Zar JH. *Biostatistical Analysis*. Prentice Hall, Englewood Cliffs, NJ, 2nd edition, 1984.
- [1105] Fukunaga K and Hayes RR. Effects of sample size in classifier design. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(8):873–885, 1989.
- [1106] Raudys SJ and Jain AK. Small sample size effects in statistical pattern recognition: Recommendations for practitioners. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(3):252–264, 1991.
- [1107] Sahiner BS, Chan HP, Petrick N, Wagner RF, and Hadjiiski L. Feature selection and classifier performance in computer-aided diagnosis: The effect of finite sample size. *Medical Physics*, 27(7):1509–1522, 2000.
- [1108] Swain PH. Fundamentals of pattern recognition in remote sensing. In Swain PH and Davis SM, editors, *Remote Sensing: The Quantitative Approach*, pages 136–187. McGraw-Hill, New York, NY, 1978.
- [1109] Nab HW, Karssemeijer N, van Erning LJTHO, and Hendriks JHCL. Comparison of digital and conventional mammography: A ROC study of 270 mammograms. *Medical Informatics*, 17:125–131, 1992.
- [1110] Nishikawa RM, Giger ML, Doi K, Metz CE, Yin FF, Vyborny CJ, and Schmidt RA. Effect of case selection on the performance of computer-aided detection schemes. *Medical Physics*, 21:265–269, 1994.
- [1111] Altman DG. *Practical Statistics for Medical Research*. Chapman and Hall, London, UK, 1991.
- [1112] Morrow WM and Rangayyan RM. Implementation of adaptive neighborhood image processing algorithm on a parallel supercomputer. In Pelletier M, editor, *Proceedings of the Fourth Canadian Supercomputing Symposium*, pages 329–334, Montreal, PQ, Canada, 1990.

- [1113] Paranjape RB, Rolston WA, and Rangayyan RM. An examination of three high performance computing systems for image processing operations. In *Proceedings of the Supercomputing Symposium*, pages 208–218, Montreal, PQ, Canada, 1992.
- [1114] Alto H, Gavrilov D, and Rangayyan RM. Parallel implementation of the adaptive neighborhood contrast enhancement algorithm. In *Proceedings of SPIE on Parallel and Distributed Methods for Image Processing*, volume 3817, pages 88–97, 1999.
- [1115] Pohlman SK, Powell KA, Obuchowski N, Chilcote W, and Grundfest-Broniatowski S. Classification of breast lesions based on quantitative measures of tumour morphology. In *IEEE Engineering in Medicine and Biology Society 17th Annual International Conference, Montreal, Canada*, page 2.4.2.3, 1995.
- [1116] Kilday J, Palmieri F, and Fox MD. Classifying mammographic lesions using computerized image analysis. *IEEE Transactions on Medical Imaging*, 12(4):664–669, 1993.
- [1117] Bruce LM and Kallergi M. Effects of image resolution and segmentation method on automated mammographic mass shape classification. In *Proceedings of SPIE vol. 3661 on Medical Imaging 1999: Image Processing*, pages 940–947, San Diego, CA, Feb. 1999.
- [1118] Bruce LM and Adhami RR. Classifying mammographic mass shapes using the wavelet transform modulus-maxima method. *IEEE Transactions on Medical Imaging*, 18(12):1170–1177, 1999.
- [1119] Tarassenko L, Hayton P, Cerneaz NJ, and Brady M. Novelty detection for the identification of masses in mammograms. In *Proceedings of the 4th International Conference on Artificial Neural Networks*, pages 442–447, Cambridge, UK, 26–28 June 1995.
- [1120] Alto H, Rangayyan RM, and Desautels JEL. An indexed atlas of digital mammograms. In *Proceedings of the 6th International Workshop on Digital Mammography*, pages 309–311, Bremen, Germany, June 2002.
- [1121] Alto H, Rangayyan RM, Solaiman B, Desautels JEL, and MacGregor JH. Image processing, radiological and clinical information fusion in breast cancer detection. In *Proceedings of SPIE on Sensor Fusion: Architectures, Algorithms, and Applications VI*, volume 4731, pages 134–144, 2002.
- [1122] Yoshitaka A and Ichikawa T. A survey on content-based retrieval for multimedia databases. *IEEE Transactions on Knowledge and Data Engineering*, 11(1):81–93, 1999.
- [1123] Gudivada VN and Raghavan VV. Content-based image retrieval systems. *IEEE Computer*, 28(9):18–22, 1995.
- [1124] Mehrotra R and Gray JE. Similar-shape retrieval in shape data management. *IEEE Computer*, 28(9):57–62, 1995.
- [1125] Trimeche M, Cheikh FA, and Gabbouj M. Similarity retrieval of occluded shapes using wavelet-based shape features. In *Proceedings of SPIE 4210: International Symposium on Internet Multimedia Management Systems*, pages 281–289, 2000.

- [1126] Safar M, Shahabi C, and Sun X. Image retrieval by shape: a comparative study. In *IEEE International Conference on Multimedia and Expo (ICME)*, volume 1, pages 141–144, 2000.
- [1127] Ivarinen J and Visa A. Shape recognition of irregular objects. In *Proceedings of SPIE 2904: Intelligent Robots and Computer Vision XV*, pages 25–32, 1996.
- [1128] Heesch D and Ruger S. Combining features for content-based sketch retrieval – a comparative evaluation of retrieval performance. In Crestani F, Girolarni M, and van Rijsbergen CJ, editors, *24th BCS-IRSG European Colloquium on IR Research*, volume LNCS 2291, pages 41–52, 2002.
- [1129] Flickner M, Sawhney H, Niblack W, Ashley J, Huang Q, Som B, Gorkani A, Hafner J, Lee D, Petkovic D, Steele D, and Yanker P. Query by image and video content: the QBIC system. *IEEE Computer*, 28(9):23–32, 1995.
- [1130] Srihari RK. Automatic indexing and content-based retrieval of captioned images. *IEEE Computer*, 28(9):49–56, 1995.
- [1131] Smith JR. Image retrieval evaluation. In *IEEE Workshop on Content-based Access of Image and Video Libraries*, pages 112–113, 1998.
- [1132] Ivarinen J and Visa A. An adaptive texture and shape based defect classification. In *Proceedings of the International Conference on Pattern Recognition*, pages 117–123, 1998.
- [1133] Cauvin JM, Le Guillou C, Solaiman B, Robaskiewicz M, Le Beux P, and Roux C. Computer-assisted diagnosis system in digestive endoscopy. *IEEE Transactions on Information Technology in Biomedicine*, 7(4):256–262, 2003.
- [1134] Moghaddam B, Tian Q, and Huang TS. Spatial visualization for content-based image retrieval. In *International Conference on Multimedia and Expo (ICME '01)*, 2001.
- [1135] Zheng L, Wetzel AW, Gilbertson J, and Becich MJ. Design and analysis of a content-based pathology image retrieval system. *IEEE Transactions on Information Technology in Biomedicine*, 7(4):249–255, 2003.
- [1136] Wooldridge M. Agent-based software engineering. *IEE Proceedings on Software Engineering*, 144:26–37, 1997.
- [1137] Paranjape RB and Smith KD. Mobile software agents for Web-based medical image retrieval. *Journal of Telemedicine and Telecare*, 6(2):53–55, 2000.
- [1138] Wooldridge M and Jennings NR. Intelligent agents: theory and practice. *Knowledge Engineering Review*, 10(2):115–152, 1997.
- [1139] Huhns MN and Singh MP. Agents on the Web. *IEEE Internet Computing*, pages 80–82, May/June 1997.
- [1140] Oshuga A, Nagai Y, Irie Y, Hattori M, and Honiden S. PLANGENT: An approach to making mobile agents intelligent. *IEEE Internet Computing*, pages 50–57, July/August 1997.

Index

- abdomen, imaging of, 35
- ACR/ NEMA standard, 1050
- active contour model, 444, 456
- acuity, 139
- acutance, 139, 1093, 1096, 1166, 1167, 1178
- adaptive filters, 228
- adaptive-neighborhood
 - contrast enhancement, 338
 - filter, 241
 - histogram equalization, 311
 - LLMMSE filter, 251
 - mean filter, 242
 - median filter, 242
 - noise subtraction, 244
 - region growing, 242, 249
 - restoration, 894
- additive noise, 115, 131, 170
- agreement, 602
- algebraic reconstruction techniques, 813
 - additive, 820
 - constrained, 821
 - multiplicative, 821
- alpha-trimmed mean filter, 181
- Anger camera, *see* gamma camera
- angular moments, 642
- anisotropy, 726
- architectural distortion, 775
- arithmetic coding, 969
- artifact, 151
 - block, 313
 - edge, 322, 893
 - grid, 21, 199
 - motion, 287
 - other, 166
 - periodic, 199
 - physiological, 40, 62, 165
 - ringing, 197, 316, 326
- association, 602
- asymmetry, 596, 704, 742
- attenuation coefficient, 15
- attenuation correction, 923
- audification, 625
- auditory display of images, 625
- autocorrelation, 155, 205, 1006
- averaging
 - geometric, 926
 - of confocal microscopy images, 270
 - of nuclear medicine images, 926
 - synchronized, 171, 271, 277
- back-propagation algorithm, 1127
- backprojection, 801
 - convolution, 811
 - filtered, 804, 811
- backward prediction, 1014
- bandpass filter, 646, 648, 660
- Bayes classifier, 1116, 1118, 1125
- Bayes rule, 840, 1116
- Beer's law, 15
- Bessel function, 105, 197
- Bezier curves, 444
- binarization, 291, 364, 456, 649, 690
- blind deblurring, 875, 877
- block-by-block processing, 167, 311, 313, 893, 998, 1001, 1008, 1010
- blood vessels
 - imaging of, 286, 684
 - in ligaments, 679
 - in the brain, 286
- blur, 16, 90
- bone, imaging of, 831
- box-counting method, 610
- brain, imaging of, 32, 49, 940
- breast boundary
 - detection of, 451, 720
- breast cancer, 3, 6, 22

- pattern classification, 429, 434, 1090, 1097, 1109, 1120, 1162, 1166, 1167
 - screening, 27, 1143
- breast masses, *see also* tumors, 1092, 1093, 1098, 1099
- content-based retrieval, 1166, 1171
- detection of, 699
- pattern classification, 1171
- shape analysis, 578, 1097, 1162, 1166
- texture analysis, 627, 1166
- breast, imaging of, 3, 25
- Burg algorithm, 1019, 1021
- Butterworth filter, 197, 325, 648, 807, 811
- CAD, 53, 55
- calcifications
 - detection of, 410, 414
 - enhancement of, 344
 - shape analysis, 575
- caliper method, 608
- Canny's method, 390
- causality, 92
- central limit theorem, 160
- centroid, 530, 563, 641
- centroidal angle, 642
- cepstrum, 623, 885
- chain coding, 530
- channel capacity, 959
- chest, imaging of, 17
- chord-length statistics, 560
- circle function, 105, 197
- circles
 - detection of, 440, 450
- circulant matrix, 215
 - block-, 221
 - diagonalization, 218
- city-block distance, 1105
- cluster seeking, 1105
 - K-means, 1108
 - maximin-distance, 1108
- coding, 955, 960
 - Huffman, 961
 - interpolative, 1001
 - predictive, 1004
 - run-length, 969
 - source, 961
 - transform, 984
- coherence, 726, 729, 734
- collagen fibers, 10, 14, 105, 110, 390, 442, 679
- collimator, 39
- color
 - pseudo-, 827
- comb filter, 890
- compactness, 578, 1160, 1162
- complexity, 543, 551, 555, 558, 560, 575, 578, 601, 608, 610
- compression, 955
 - lossy versus lossless, 957
 - segmentation based, 1051
- Compton scattering, 920
- computed radiography, 17, 287
- computed tomography, 29
 - display of images, 825
 - with diffracting sources, 825
- computer-aided diagnosis, *see* CAD
- concavity, 537, 569, 578, 691, 1096, 1160, 1162
- confocal microscopy, 270
- constrained least-squares restoration, 872
- content-based image retrieval, 1169
- contingency table, 1138
- contour coding, 977
- contrast, 75, 129, 600, 734
 - for X-ray imaging, 286, 839
- contrast enhancement, 344
 - adaptive-neighborhood, 338
 - clinical evaluation, 1145
 - of mammograms, 350
- contrast histogram, 346
- contrast-to-noise ratio, 137
- convolution, 91, 117, 220, 224, 315
 - backprojection, 811
 - circular, 213
 - illustration, 215
 - linear, 213, 215
 - periodic, 213
- convolution mask, 176, 314
- correlation, 118, 601
- correlation coefficient, 119, 168
- covariance, 168, 205, 1007
- covariance matrix, 1106
- CREW coding, 1066
- cross-correlation, 169

- CT, *see* computed tomography
- cyclo-stationary signal, 167
- cyst, imaging of, 3, 6, 43, 46
- deblurring, *see* also restoration
 blind, 877
 motion, 875
- decision function, 1095, 1101, 1118, 1119
- decision making, 1091
- deconvolution, *see* also restoration
 homomorphic, 623, 885, 886
- decorrelation, 980
- delta function, 78, 90, 95, 798
 sifting property, 90
- density slicing, 292, 710, 827
- detection, *see* also segmentation
 of architectural distortion, 775
 of asymmetry, 742
 of breast boundary, 451, 720
 of breast masses, 699
 of calcifications, 410
 of circles, 440, 450
 of edges, 604
 of isolated lines, 365
 of isolated points, 365
 of lines, 780
 of pectoral muscle, 481
 of ripples, 604
 of spots, 604
 of straight lines, 437
 of the fibroglandular disc, 743
 of the spinal canal, 449
 of tumors, 417, 500
 of waves, 604
- diagnostic accuracy, 1132
- diagnostic decision, 56, 1089, 1090, 1182
- diagonalization of a circulant matrix, 218
- difference of Gaussians, 376
- differentiation, 119, 367–369, 983
 in matrix representation, 224
- difficulties in image acquisition, 61
- digital radiography
 MTF, 127
- digital subtraction angiography, 286
- directional distribution
 measures of, 641
- directional filtering, 644
- directional pattern analysis, 639
- discrete cosine transform, 987
- discriminant analysis, 1095
 of breast masses, 1162
- distance, 1105
 between probability density functions, 1141
 city-block, 1105
 Euclidean, 643, 1101, 1105, 1175
 Hamming, 959
 Jeffries–Matusita, 1142
 Mahalanobis, 1105
 Manhattan, 643, 1105
- distance function, 1097
- distortion measure, 959
- divergence, 1141
- dominant angle, 641
- dot product, 1124
 normalized, 1106
- DPCM, 984, 998, 1005, 1046, 1049
- dual-energy imaging, 287
- dynamic range, 73
- dynamic system, 165, 167
- eccentricity, 551
- ECG signal, 280
- echo removal, 623, 886
- echocardiography, 44
- edge detection, 367, 493, 604
- edge enhancement, 315, 322
- edge function, 96
- edge linking, 392, 493, 495
- edge spread function, 93, 131, 139
- edge-flow propagation, 493
- eight-connected neighborhood, 176
- electron microscopy, 10
- energy-subtraction imaging, 287
- enhancement
 in breast cancer screening, 1143
- ensemble averages, 155, 167, 172
- entropy, 84, 601, 643, 681, 687, 697, 734, 957, 983, 1070
 conditional, 89
 higher-order, 89
 joint, 88
 Markov, 1071
 mutual information, 90
- ergodic process, 167, 176

- error measures, 138
- Euclidean distance, 643, 1101, 1105, 1175
- expectation maximization, 745, 754, 842
- F1 transform, 1063
- F2 transform, 1064
- false negative, 1133
- false positive, 1133
- fan filter, 646, 647, 651
- feature selection, 1141
- feature vector, 1091
- fetal imaging, 46
- fibroblasts, image of, 7
- fibroglandular disc, 743
- fidelity criteria, 959
- film-grain noise, 170, 254
- filter
 - adaptive, 228
 - adaptive 2D LMS, 235
 - adaptive rectangular window, 237
 - adaptive-neighborhood, 241
 - adaptive-neighborhood noise subtraction, 244
 - alpha-trimmed mean, 181
 - band-reject, 199
 - bandpass, 648, 660
 - blind deblurring, 877
 - Butterworth, 197, 325, 648, 807, 811
 - comb, 199, 890
 - comparative analysis, 867
 - constrained least-squares restoration, 872
 - directional, 644
 - fan, 646, 647, 651
 - finite impulse response, 213
 - for periodic artifacts, 199
 - frequency domain, 193
 - Gabor, 657
 - generalized linear, 328
 - high-emphasis, 325
 - highpass, 325
 - homomorphic, 328, 335, 885, 886
 - ideal, 194, 325, 647, 648
 - infinite impulse response, 213
 - inverse, 858, 867
 - Kalman, 898, 943
 - L, 181
 - LMMSE, 225
 - local LMMSE, 228
 - local-statistics, 174
 - lowpass, 194
 - max, 181
 - mean, 176
 - median, 177
 - Metz, 874, 940
 - min, 181
 - min/max, 181
 - motion deblurring, 875
 - moving average, 121
 - multiresolution, 660
 - noise-updating repeated Wiener, 234
 - nonstationary, 231, 235
 - notch, 199, 890
 - optimal, 224, 865, 872, 898
 - order-statistic, 177, 181
 - power spectrum equalization, 860, 867, 877, 935
 - ramp, 648, 806
 - sector, 646, 647, 651
 - space-invariant, 857
 - space-variant, 231, 235, 891
 - three-dimensional, 947
 - Wiener, 225, 233, 863, 897, 935
- filtered backprojection, 804, 811
- finite impulse response filter, 213
- fixed operators
 - limitations of, 323
- fluoroscopy, 17
- focal spot, 19, 92, 127
- four-connected neighborhood, 175
- Fourier descriptors, 562
- Fourier slice theorem, 798
- Fourier transform, 99, 122, 206
 - convolution property, 117
 - coordinates of, 105
 - display of, 116
 - folding of, 116
 - inverse, 115
 - linearity, 115
 - matrix representation, 206
 - multiplication property, 118
 - of a circle, 105
 - of a line, 646
 - of a rectangle, 104, 105

- periodicity, 116
- properties, 110
- rotation, 117
- scaling property, 117
- separability, 115
- shift property, 116
- shifting of, 116
- symmetry, 115, 116
- fractal analysis of texture, 605
- fractals, 1035
- fractional Brownian motion, 609
- frame averaging, *see* synchronized averaging
- Freeman chain coding, 530, 977, 978
- frequency-domain filters, 193
- functional imaging, 6, 36, 40, 43, 44, 49
- fusion operator, 426, 500
- fuzzy connectivity, 449
- fuzzy fusion, 426, 500
- fuzzy region growing, 429
- fuzzy segmentation, 421
- fuzzy sets, 417
- Gabor function, 622, 657, 659, 755, 757, 780
- Gabor wavelets, 487, 663, 757, 780
- gamma, 73
- gamma camera, 38
 - spread function, 97
- gamma correction, 294, 345
- gated blood-pool imaging, 168, 277
- Gaussian mixture model, 744, 840
- Gaussian noise, 154, 172, 180, 194, 254
- Gaussian PDF, 159, 1110, 1118, 1120
- generalized linear filtering, 328
- geometric distortion, 804, 813, 822
- geometric mean, 868, 940
 - of planar images, 925, 926
- geometric transformation, 291
- global operations
 - limitations of, 310
- gold standard, 1134, 1138
- gradient analysis, 632
 - of breast masses, 627, 702
- gradients, 367
- Gray code, 961, 969, 1055, 1062
- gray-level co-occurrence matrix, 597
- gray-scale windowing, 292, 827
- grid artifact, 21, 199
- ground-truth, 1109
- Hadamard matrices, 211
- Hamming distance, 959
- Hamming window, 809, 894
- Haralick's measures of texture, 600
- head, imaging of, 32
- heart, imaging of, 41, 44, 277, 834, 937
- hexadecimal code, 961
- high-emphasis filter, 325
- high-frequency emphasis, 315, 322
- high-frequency noise, 194
- highpass filter, 325
- histogram, 78
 - of contrast, 346
- histogram concavity, 691
- histogram equalization, 301, 478, 501
 - adaptive-neighborhood, 311
 - local-area, 310
- histogram specification, 305
- homogeneity, 601
- homomorphic deconvolution, 623, 885, 886
- homomorphic filtering, 328, 335
- Hough transform, 435, 450, 481
- Hounsfield units, 825, 826
- Huffman coding, 961
- human-instrument system, 53
- Hurter-Driffield curve, 73
- ideal filter, 194, 325, 647, 648
- impulse function, 78, 90
- impulse noise, 177
- indexed atlas, 1172, 1177
- infinite impulse response filter, 213
- inflection
 - points of, 534, 542
- information content, 61
 - joint, 87, 88
 - mutual, 90
- information theory, 956
- infrared imaging, 3
- inhomogeneity, 596
- integration, 121
- interference, 151
 - physiological, 165

- power-line, 164
- interpolation, 66
- interpolative coding, 1001
- interval cancer, 1145, 1146
- intrinsic orientation angle, 726
- inverse filter, 858
- isointensity contours, 710, 712, 725

- JBIG, 1046
 - enhanced, 1062
- Jeffries–Matusita distance, 1142
- JM distance, 1142
- joint information, 88
- JPEG, 1049
- just-noticeable difference, 76, 343, 405

- K-means clustering, 1108
- k-NN method, 1104
 - applied to breast masses, 1171, 1177
- Kaczmarz method, 813
- Kalman filter, 898, 943
- Karhunen–Loève transform, 763, 769, 989
- kidney, imaging of, 10
- knee, imaging of, 49
- kurtosis, 560, 596

- L filter, 181
- Laplacian, 120, 138, 316
 - subtracting, 316
- Laplacian MSE, 138
- Laplacian of Gaussian, 376
- Laplacian PDF, 163, 984, 1001
- Laws' measures of texture, 603
- leave-one-out method, 1125
- left-most-looking rule, 977
- Lempel–Ziv coding, 974
- Levinson algorithm, 1014, 1019, 1023
- ligament tissue, imaging of, 7, 10, 14, 680, 684
- likelihood function, 745, 840, 1110
- limitations of
 - fixed operators, 323
 - global operations, 310
- line detection, 780
- line function, 95
- line spread function, 93, 127, 131
 - gamma camera, 97
 - line, Fourier transform of, 646
- linear discriminant analysis, 1095, 1097, 1104, 1166
- linear prediction, 1005
- linear regression, 692
- live wire, 449
- liver, imaging of, 947
- Lloyd–Max quantizer, 68
- LMMSE filter, 225
- LMS filter, 235
- local LMMSE filter, 228
- local-area histogram equalization, 310
- local-statistics filters, 174
- logarithmic transformation, 334, 456, 462
- logistic regression, 429, 434, 1120
- loss function, 1110
- lossless compression, 957
- lowpass filter, 194

- magnetic resonance imaging, *see* MRI
- magnification, 27
- Mahalanobis distance, 1105
 - applied to breast masses, 1167
- mammograms
 - analysis of, 410, 575, 578, 627, 699, 742, 775, 1160, 1166
 - enhancement of, 335, 344, 350, 1143
- mammography, 22
- Manhattan distance, 643, 1105
- Markov entropy, 1071
- Markov source, 1071
- matched filter, 867
- matrix representation, 69, 202
 - differentiation in, 224
 - of convolution, 212
 - of images, 203
 - of transforms, 206
- max filter, 181
- maximin-distance clustering, 1108
- maximum likelihood, 842, 1124, 1138, 1147
- McNemar's test of symmetry, 1138
 - in breast cancer screening, 1147
- mean, 152, 204, 1118
 - geometric, 868, 925, 926, 940
 - of planar images, 925
- mean filter, 176

- adaptive-neighborhood, 242
- mean-squared error, 68, 138
- mean-squared value, 152
- measure of fuzziness, 426
- medial axis, 548
- median filter, 177
 - adaptive-neighborhood, 242
- Metz filter, 874, 940
- microscopy
 - confocal, 270
 - electron, 10
 - light, 6
- microtomography, 831
- min filter, 181
- min/max filter, 181
- minimum-description length, 746, 843
- mobile agents, 1177
- modulation transfer function, *see* MTF
- moments, 601
 - angular, 642
 - first-order, 152
 - second-order, 152
- motion artifact, 287
- motion deblurring, 875
- moving-average filter, 121
- moving-window processing, 167, 174, 662
- MPEG, 1050
- MRI, 47
- MTF, 122, 129, 131, 140
 - digital radiography, 127
 - screen-film system, 127
- multichannel linear prediction, 1009
- multiframe averaging, *see* synchronized averaging
- multiplication of images, 118, 328
- multiplicative noise, 170
- multiresolution analysis, 660, 707, 762
- multiscale analysis, 380, 390, 491, 666, 700
- mutual information, 90
- myocyte, image of, 7
- nearest-neighbor rule, 1104
- negative predictive value, 1134
- neighborhood
 - eight-connected, 176
 - four-connected, 175
 - shapes, 175
- neural networks, 1126
- neuroblastoma, 834
- NMR, 47
- noise, 131, 151
 - additive, 115, 131, 170
 - film-grain, 170, 254
 - Gaussian, 154, 172, 180, 194, 254
 - high-frequency, 194
 - impulse, 177
 - in nuclear medicine images, 271
 - multiplicative, 170
 - PDFs, 159
 - photon detection, 21
 - Poisson, 169, 180, 254, 271
 - salt-and-pepper, 166, 177, 180, 181, 259
 - signal-dependent, 169, 248
 - SNR, 131
 - speckle, 43, 170, 254
 - structured, 40, 164
 - uniformly distributed, 254
- noiseless coding theorem, 957
- nonstationary filter, 231, 235
- nonstationary process, 166, 230, 414
- normal equations, 1006, 1029
- normalized error, 138
- normalized MSE, 138
- notch filter, 890
- nuclear magnetic resonance, 47
- nuclear medicine images
 - restoration of, 919, 934
- nuclear medicine imaging, 36
 - attenuation correction, 923
 - noise reduction in, 271
 - quality control, 922
 - scatter compensation, 922
- objectives of image analysis, 53
- octal code, 961
- optical density, 72
- optical transfer function, 122
- optimal filter, 224, 865, 872, 898
- order-statistic filters, 177, 181
- oriented pattern analysis, 639
- Otsu's method of thresholding, 649
- outliers, 181
- parabolic modeling of contours, 543
- parallel-ray geometry, 798

- Parseval's theorem, 115, 993
- pattern classification, 1089, 1091
 - reliability, 1140
 - supervised, 1095
 - test set, 1095
 - test step, 1125
 - training set, 1095, 1140
 - training step, 1125
 - unsupervised, 1104
- pectoral muscle
 - detection of, 481
- perceptron, 1127
- PET, 41
- phantom, 64, 92, 97, 199, 271, 934, 935, 943
- phase, 104, 116
 - linear component, 116, 886
 - unwrapping, 886
 - use in blind deblurring, 878
- phase portraits, 779
- photomultiplier tube, 39
- photon detection noise, 21
- physiological imaging, *see* functional imaging
- physiological interference, 165
- planar imaging, 15, 32, 40, 41
- point spread function, 91, 802, 805
- Poisson noise, 169, 180, 254, 271
- Poisson process, 15, 160
- polygonal modeling of contours, 537
- positive predictive value, 1134
- positivity, 203, 874
- positron emission tomography, *see* PET
- power law, 609
- power spectral density, 159
- power spectrum equalization, 860, 877, 935
- power-line interference, 164
- predictive coding, 1004, 1049
- Prewitt operators, 368
- principal axis, 641
- principal-component analysis, 769, 989, 991
- probabilistic models, 1110
 - Bayes rule, 1116
 - conditional probability, 1116
 - likelihood function, 1116
 - posterior probability, 1110
 - prior probability, 1110
- probability density function, 80, 151
 - divergence, 1141
 - Gaussian, 159, 1118
 - Laplacian, 163
 - normalized distance, 1141
 - Poisson, 160
 - Rayleigh, 163
 - uniform, 160
- projection
 - fan-beam, 31
 - geometry, 797
 - image reconstruction from, 797
 - imaging, 15, 32, 40, 41
 - parallel-ray, 31, 797
- projection theorem, 798
- prototype, 157, 1097, 1123
- pseudo-color, 827
- pyramidal decomposition, 707, 720
- quality of images, 61
 - characterization of, 64
- quantitative analysis, 56
- quantization, 66
- quasistationary process, 167, 176
- radiographic imaging, 15
- Radon transform, 798, 886
- ramp filter, 648, 806
- random noise, 151
- random variable, 20, 87, 151
- ray integral, 798
- Rayleigh PDF, 163
- receiver operating characteristics, 1135
 - applied to breast masses, 1147, 1162, 1167
 - free-response, 791
 - in breast cancer screening, 1145
- reconstruction
 - algebraic reconstruction techniques, 813, 821
 - additive, 820
 - constrained, 821
 - multiplicative, 821
 - backprojection, 801
 - convolution backprojection, 804, 811
 - filtered backprojection, 804, 811
 - Fourier method, 97, 801
 - from projections, 797

- simultaneous iterative reconstruction technique, 822
 - with limited data, 804, 813, 822
- rectangle function, 104
- region growing, 340, 393, 421, 429, 1052
 - adaptive-neighborhood, 242, 249
 - splitting and merging, 397
- region-based segmentation, 396
- relative dispersion, 608
- resolution, 99
 - recovery, 924
- restoration, 857
 - comparison of filters, 867
 - information required, 875
- ringing artifact, 197, 316, 326
- ripple detection, 604
- risk-benefit analysis, 54
- Roberts operator, 369
- root mean-squared value, 152
- rose diagram, 641, 678, 682, 683, 686, 696, 771, 772
- rotation, 105, 117, 655
- run-length coding, 969
- Rutherford-Appleton threshold, 691
- salt-and-pepper noise, 166, 177, 180, 181, 259
- sampling, 65
- scale-space, 380
- scaling, 117
- scanning geometry, 31
- scatter compensation, 922
- scattering, 40, 43, 92
 - Compton, 920
- screen-film system, 16, 92
 - MTF, 127
- screening, 1132
 - breast cancer, 1143
- sectioned filtering, 893, 897
- sector filter, 646, 647, 651
- segmentation, 393, 396, *see also* detection
 - improvement of, 444
 - of contours, 534
 - of texture, 622
- segmentation-based coding, 1051
- sensitivity, 1132
- sequency, 210
- shape analysis, 529
 - of breast masses, 1097, 1162
 - of calcifications, 575, 1128
 - of tumors, 578
- shape factors, 549
- sharpness, 139
- short-time analysis, 167, 662
- signal-dependent noise, 169, 248
 - transformation of, 171
- signal-to-noise ratio, *see* SNR
- signature of a contour, 530, 562
- simultaneous iterative reconstruction technique, 822
- sinc function, 104, 646
- single-photon emission computed tomography, *see* SPECT
- skeletonization, 548, 692
- skewness, 560, 596
- snakes, 444, 456
- SNR, 131
- Sobel operators, 369, 604
- sonification, 625
- source coding, 961
- space-invariant filter, 857
- space-variant filters, 231, 235, 891
- space-variant systems
 - transformation to space-invariant systems, 893
- spatial statistics, 157
- specificity, 1132
- speckle noise, 43, 170, 254
- SPECT, 32, 40
 - restoration of, 919
 - spread function, 99
- spectrum, 99
- spiculation index, 570, 578, 1160, 1162
- spinal canal
 - detection of, 449
- splines, 444
- spot detection, 604
- spread functions, 90
- standard deviation, 152
- stationary process, 166
 - block-wise, 167
 - cyclo-, 167
 - ergodic, 167
 - non-, 167
 - quasi-, 167, 176
- statistical decision, 1110

- statistical separability, 1141
- statistically independent processes, 155
- step function, 96
- straight line, detection of, 437
- streaking, 804, 813
- structural analysis of texture, 621
- structured noise, 164
- subtracting Laplacian, 316
- subtraction, 287
 - angiography, 286
 - energy, 287
 - temporal, 291
- synchronized averaging, 171, 277
 - in confocal microscopy, 271
- telemedicine, 1177
- teleradiology, 1079
 - analog, 1080
 - digital, 1082
 - high-resolution, 1084
- temperature, 2, 4, 5
- template matching, 119, 169
- temporal statistics, *see* time averages
- temporal subtraction, 291
- texton, 583, 621, 891
- texture, 583
 - Fourier spectrum, 612
 - fractal measures, 605
 - Haralick's measures, 600
 - in liver image, 43
 - Laws' measures, 603
 - model for generation, 584
 - of breast masses, 627, 699, 701, 1166
 - ordered, 589
 - oriented, 590, 639
 - periodic, 891
 - random, 589
 - statistical analysis, 596
 - structural analysis, 621
- texture energy, 603
- texture flow-field, 726
- thermal imaging, 2
- thermography, 3
- thinning, 548
- three-dimensional filtering, 947
- thresholding, 291, 364, 690, 722
 - optimal, 395
 - Otsu's method, 649
- Rutherford-Appleton, 691
- time averages, 157, 167
- tissue characterization, 838
- Toeplitz matrix, 213, 1008
- tomography, 27
- transform coding, 984
- transillumination, 6
- true negative, 1132
- true positive, 1132
- tumor
 - in neuroblastoma, 834
- tumors, *see also* breast masses
 - detection of, 417, 500, 699
 - shape analysis, 578
- ultrasonography, 43
- uncertainty principle, 659
- uniform PDF, 160, 254
- uniformity, 596
- unit step function, 163
- unsharp masking, 314, 322, 345
- variance, 135, 152, 560, 596, 601, 1118
- varicocele, detection of, 3
- vector representation, *see* matrix representation
- Walsh functions, 209
- Walsh-Hadamard transform, 209, 211
- warping, 291
- wave detection, 604
- wavelets, 659, 663, 707, 756, 757
 - Gabor, 487, 663, 757, 780
- Weber's law, 76, 343, 405
- Wiener filter, 225, 233, 863, 897, 935
 - noise-updating repeated, 234
- window
 - Hamming, 809, 894
- windowing of gray scale, 292, 827
- X ray
 - beam hardening, 20
 - contrast medium, 286, 839
 - dual-energy imaging, 287
 - energy, 19
 - energy-subtraction imaging, 287
 - exposure, 19
 - grids, 20
 - imaging, 15

scatter, 20, 25

stopping, 22

xeromammography, 25

Yule–Walker equations, 1006, 1014

zero padding, 101, 193, 215

zero-crossings, 370, 380