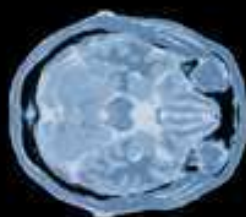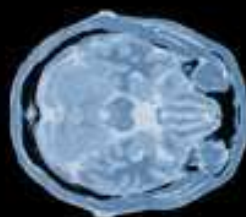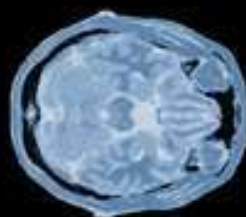# MEDICAL IMAGING SYSTEMS TECHNOLOGY

Methods in Diagnosis Optimization

Cornelius T Leondes

editor

MEDICAL
IMAGING SYSTEMS
TECHNOLOGY
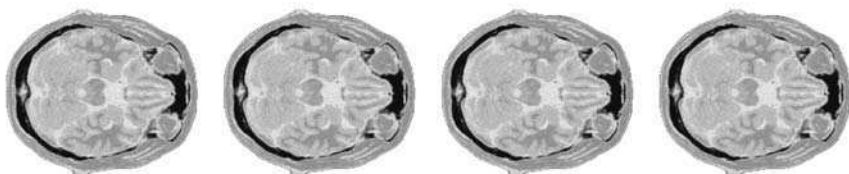*Methods in Diagnosis Optimization*

This page intentionally left blank

MEDICAL IMAGING SYSTEMS TECHNOLOGY

Methods in Diagnosis Optimization

A 5-Volume Set

editor

Cornelius T Leondes
University of California, Los Angeles, USA

W **World Scientific**

MEDICAL IMAGING SYSTEMS TECHNOLOGY
A 5-Volume Set
Methods in Diagnosis Optimization

Preface

Because of the availability of powerful computational techniques, new modality techniques such as Computer-Aided Tomography (CAT), Magnetic Resonance Imaging (MRI) and others, and because of the new techniques of imaging processing (machine vision), the lives of many patients will be saved, and the quality of all our lives improved. This marriage of powerful computer technology and medical imaging has spawned a new and growing generation of young dynamic doctors who hold PhDs in physics and/or computer science, along with their MDs. In addition, technologists and computer scientists, with their superb skills, are also deeply involved in this area of major significance.

This volume covers the subject of medical imaging systems — methods in diagnosis optimization, by leading contributors on the international scene. This is one of the 5 volumes on medical imaging systems technology, and together they collectively constitute an MRW (Major Reference Work). An MRW is a comprehensive treatment of a subject requiring multiple authors and a number of distinctly-titled and well-integrated volumes. Each volume treats a specific subject area of fundamental importance in medical imaging. The titles of the respective 5 volumes which compose this MRW are:

- Medical Imaging Systems — Analysis & Computational Methods
- Medical Imaging Systems — Modalities
- Medical Imaging Systems — Methods in General Anatomy
- Medical Imaging Systems — Methods in Diagnosis Optimization
- Medical Imaging Systems — Methods in Cardiovascular & Brain Systems

Each volume is self-contained and stands alone for those interested in a specific volume. However, collectively this 5-volume set evidently constitutes the first multi-volume comprehensive reference dedicated to the multi-discipline area of medical imaging.

There are over 130 coauthors of this notable work and they come from 25 countries. The chapters are clearly written, self-contained, readable and comprehensive with helpful guides including introduction, summary, extensive figures and examples with in-depth reference lists. Perhaps the most valuable feature of this work is the breadth and depth of the topics covered.

This volume on "Medical Imaging Systems — Methods in Diagnosis Optimization" includes essential subjects like:

(a) Robust techniques for enhancement of micro-calcifications in digital mammography
(b) Techniques in the detection of micro-calcification clusters in digital mammograms
(c) Fuzzy region growing and fusion methods for the segmentation of masses in mammograms
(d) ROC methodology and its application in breast cancer diagnosis
(e) Parametric shape reconstruction in inverse problems: Fundamental performance bounds and algorithms
(f) Wavelet techniques in region-based digital data compression and their application in digital mammography
(g) Techniques in segmenting images with anisotropic spatial resolution and for tracking temporal image
(h) Functional MRI activity characterization: An estimation and decision theoretic approach
(i) Techniques for detection of spectral signatures in MR images and their applications

The contributors of this volume clearly reveal the effectiveness of the techniques available and the essential role that they will play in the future. I hope that practitioners, research workers, computer scientists, and students will find this set of volumes to be a unique and significant reference source for years to come.

Contents

Chapter 7
Techniques in Segmenting 3D Images with Anisotropic Spatial
Resolution and for Tracking Temporal Image Sequences and

Chapter 8
Functional MRI Activity Characterization: An Estimation and

Chapter 9
Techniques for Detection of Spectral Signatures in MR Images

Chapter 10
Studying Anatomy and Disease in Medical Images Using

CHAPTER 1

ROBUST TECHNIQUES FOR ENHANCEMENT OF MICROCALCIFICATIONS IN DIGITAL MAMMOGRAPHY

PETER HEINLEIN

Image diagnost International GmbH
Balanstraße 57, 81541 München, Germany
E-mail: post(at)peter-heinlein.de

We present methods for the enhancement of microcalcifications in breast radiographs. The result of such an algorithm is a manipulated image where calcifications appear more clearly. By making these very subtle signs more visible, the radiologist can be aided in detecting a cancer. This can help to improve both efficiency and accuracy of the diagnosis.

Our main contribution to enhancement of microcalcifications is the improvement of the underlying wavelet methods. Based on the continuous wavelet decomposition we apply a novel method to discretize the continuous wavelet transform. By choosing discrete wavelet scales that match the size of microcalcifications we obtain contrast improvements that outperform hitherto used wavelet transforms. We present new results which indicate that the algorithm is very robust with respect to different mammographic imaging modalities.

The flexibility of our approach is also demonstrated in applications as detection of microcalcification and as artifact free image denoising.

Keywords: Robust techniques; digital mammography; microcalcifications.

1. Introduction

Easily interpretable x-ray mammograms have been a clinicians dream for quite a long time. The recent introduction of dedicated digital review workstations for mammography starts the departure from conventional film-based reading. This development opens up new opportunities for medical image processing. Our emphasis is on enhancement of features that are relevant for diagnosis.

Breast cancer is the main cause of death for women between the ages of 35 to 55. Early detection and treatment of breast cancer are the most effective methods of reducing mortality. Microcalcifications are among the earliest signs of a breast carcinoma. Actually, as radiologists point out,[17] microcalcifications can be the only mammographic sign of non-palpable breast disease. Due to the subtle nature of these microcalcifications these are often overseen in the mammogram. Some retrospective studies state that in up to 40% of the cases unambiguous signs of a cancer were missed by the reader, with in some cases fatal consequences for the

patient. Thus the reliable detection and classification of microcalcifications plays a very important role in early breast cancer detection.

Computer assisted detection (CAD), which is reality today, was a first step to simplify the detection of malignant lesions. Up to now, however, CAD systems just put markers on suspicious regions. They do not generate a processed image that might show relevant features more clearly. This restriction is due to two reasons. First, most mammograms are still film-based, and are read using a lightbox. Thus commercial CAD systems digitize the film and present markers on a small display or on a separate printout. Second, for the purpose of detection the image is just decomposed to generate features for a classifier. The task of enhancement is more complex, as it also requires an image reconstruction. Only a very sophisticated reconstruction algorithm can provide images that are still suitable for diagnostic reading.

Our focus in this article is on providing algorithms for the enhancement of microcalcifications in breast radiographs. The result of such an algorithm is a manipulated image where calcifications appear more clearly. By making these very subtle cancer signs more visible, the radiologist can be aided in detecting a cluster of calcifications. Further, if the features of the calcifications are made more evident, the radiologist might be aided in determining whether the calcifications originate from a benign or a malignant lesion. This can help to improve both efficiency and accuracy of the diagnosis. Thus enhancement of microcalcifications can efficiently complement the markers generated by a standard CAD system.[10]

Considerable research has been undertaken in the development of automated image analysis methods to assist radiologists in the identification of abnormalities. In a recent overview article, H. D. Cheng *et al.*[8] gathered more than 120 citations on the topic of microcalcification detection and classification. Among these techniques, there are many promising wavelet-based approaches.

Our contribution to enhancement of microcalcifications is the improvement of the underlying wavelet methods. Based on the continuous wavelet decomposition we apply a novel method to discretize the continuous wavelet transform. This discretization permits arbitrary scales and orientations while maintaining the reconstruction property of the continuous transform. It allows to adapt the wavelet filterbank optimally to the size and shape of microcalcifications. This is the basis for our microcalcification enhancement algorithm.

We start with a short introduction to the medical background in Sec. 2. It includes an overview of the available mammographic imaging modalities and introduces the relevant aspects of x-ray image formation.

An overview of known methods for enhancement of microcalcifications is given in Sec. 3. We explain the role of wavelets for enhancement and we sketch the basic concept of transform based enhancement.

Section 4 introduces a novel discrete wavelet transform, so called integrated wavelets. Integrated wavelets provide flexible discrete versions of the continuous wavelet transform (CWT). This idea firstly appeared in an article by

M. Duval-Destin, M. A. Muschietti and B. Torrésani[15] on affine wavelets. We have generalized the concept and have introduced directional integrated wavelets.[20,21] We show that these integrated wavelets can be specifically tailored for analysis applications where an image reconstruction is needed.

Results of enhancement are presented in Sec. 5. We study in detail the ingredients of wavelet based enhancement methods. We derive an adapted wavelet from a model for microcalcifications. Our enhancement operator on wavelet coefficients is a refinement of known approaches by A. Laine *et al.*[26,27] With integrated wavelets, we can chose the discrete scales to match the size of the interesting structure. This yields contrast improvements that outperform hitherto used wavelet transforms by about 5 to 10%. Further, we present new results that indicate that the algorithm is very robust with respect to different mammographic imaging modalities.

Finally, Sec. 6 sketches further applications. We present how in the setting of integrated wavelets directional information can be used to distinguish microcalcifications from similar structures. A second application is uncorrelated denoising. Based on the Morlet reconstruction we introduce a denoising procedure that is not correlated in space. For example, applied to mass detection in mammograms, we find that this denoising method does not disturb the local image statistics that we use for feature extraction.

2. Preliminaries

At first glance, given just a single mammogram, the task of detecting microcalcifications seems almost easy. But each breast is different in anatomy. Further, there exist various different microcalcification patterns. To complicate matters even more, there are several x-ray imaging technologies that each have their own characteristics.

Thus, from the point of view of image analysis, the main challenge of mammography is diversity. An algorithm must be reliable under all the varying biometric and technical parameters. This section provides some background that helps to understand the challenges an algorithm for enhancement of microcalcifications has to face.

2.1. *Medical background*

The earlier a breast cancer is diagnosed, the easier it is to cure it. Microcalcifications are an early sign of a breast carcinoma. In fact, as S. Heywang-Köbrunner *et al.*[23] point out, microcalcifications can be the only mammographic sign of non-palpable breast disease. Microcalcifications account for over 50% of all the non-palpable lesions detected using mammography. For this reason, the reliable detection of microcalcifications is one of the major goals of mammographic image processing.[24]

The main reason for x-ray mammography is the fact that other imaging methods such as magnetic resonance imaging or ultrasound imaging cannot capture and
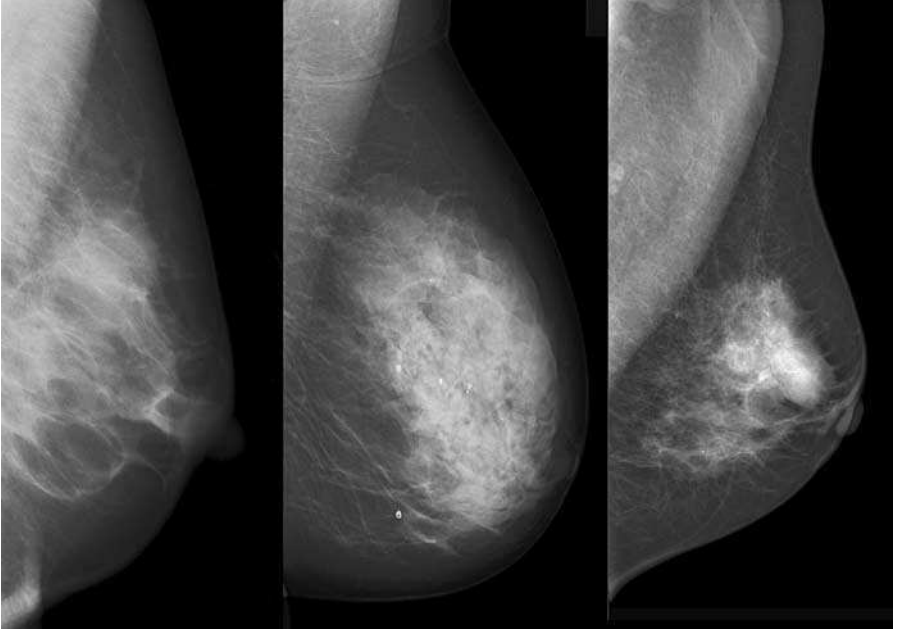
Fig. 1.   Examples of x-ray mammograms from different modalities. A digitized film with high spatial resolution but low contrast (left), a digital luminescence radiograph that exhibits strong noise (center) and a direct digital mammogram with high contrast and low spatial resolution (right).

visualize microcalcifications. Figure 1 shows mammograms from different imaging modalities.

Calcification shows up on a mammogram when calcium deposits have collected in the ducts of the breast. In many cases, calcification can occur without there being a cancer there. If there is a breast cancer in that area, then the pattern of the calcium that shows up on the mammogram can have a particular look to it that a specialist in reading mammograms will recognize.

It is a difficult task to detect microcalcifications in a mammogram with the naked eye. Studies indicate, compare an overview by U. Bick,[7] that in up to 30% of the cases analyzed the radiologist oversees definite signs of a cancer. In a mammogram, microcalcifications appear as small spots that are brighter than the surrounding tissue. They can be very small, actually at the limit of the spatial resolution of the imaging system. Imaging noise and dense breast tissue can further occlude theses signs. In a screening scenario the problem of detection is increased by low incidence. In 100.000 screening cases there are typically no more than 200 cases of cancer.

Basic breast anatomy

The breast is a mound of glandular, fatty and fibrous tissue located over the pectoralis muscles of the chest wall and attached to these muscles by fibrous strands.

A layer of fat surrounds the breast glands and extends throughout the breast. The actual breast is composed of fat, glands with the capacity for milk production, blood vessels and milk ducts to transfer the milk from the glands to the nipples.

Normal, non-fat breast tissue is water dense and appears light. Fatty tissue is practically radio-transparent and appears very dark in a mammogram. The dynamic range in mammography is large, since there are large variations in breast anatomy. Individual breast appearance is influenced by the volume of a woman's breast tissue and fat, her age, a history of previous pregnancies and lactation, her heredity, the quality and elasticity of her breast skin, and the influence of hormones. Normal anatomy on a mammogram will image differently depending on a woman's weight, age, presence of surgical scars and presence of superficial or sub-muscular implants, as well as the amount of fatty tissue in her breasts.

Microcalcifications

Calcifications are small deposits of calcium. They occur in the breast as a result of secretions within structures that have become thickened and dried. The diameter of microcalcifications is about 0.05 mm up to 1 mm.[23,24] On a digital image with a pixel spacing of 0.1 mm this equals an area of less than ten pixels. In comparison, the overall image matrix is about 4.500×5.500 pixels.

When microcalcifications indicate breast cancer, they are most frequently present in clusters of 10 to 100 single findings. Relevant for diagnosis are clusters with at least four findings per square cm. Single findings in a cluster can vary significantly in shape and size. Microcalcifications that originate from malignant lesions are typically very small, irregular in shape and size. The average distance of findings in a cluster is below 1 mm. There are also calcifications that originate from normal phenomena, for example in blood vessels, or as large deposits of calcium with diameters of several millimeters. Figure 2 shows a classification for clusters based on different morphology of individual findings as it was identified by Le Gal.[19] This classification helps deciding whether the cluster belongs to a malignant lesion or a benign process.

2.2. *Mammographic imaging modalities*

Mammographic images pose a tough challenge for image processing. The images have poor signal-to-noise ratio. They are blurred by scattered x-ray photon

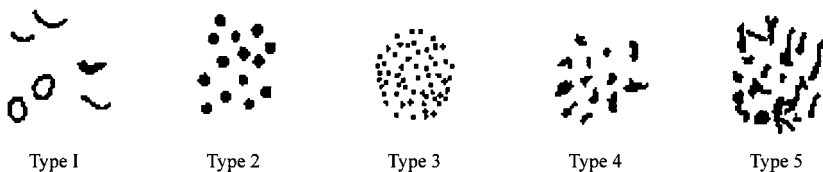|				
Type I	Type 2	Type 3	Type 4	Type 5

Fig. 2. Classification of microcalcifications by Le Gal. Classification of microcalcifications in five classes by Le Gal.[19] The probability of a malignant tumor is growing from 0% for Type 1, to 19%, 39%, 59% up to 96% for Type 5. Figure from M. Brady.[24]

radiation and by intensifying screen glare. In addition, there is a compromise between radiation dose and image quality. An x-ray mammogram is a two-dimensional projection of a three-dimensional anatomical structure. Thus the image inevitably superimposes tissue regions that are in fact not connected.

Today, there are three different x-ray mammography imaging technologies available:

- Full field direct digital mammography (FFDM)
- Digital luminescence radiography, also called Computed Radiography (CR)
- Secondary capture of analog film using a film digitizer (Laser or CCD).

Figure 1 shows examples of mammograms from these different technologies.

Due to the fact that direct digital systems are still significantly more expensive than conventional film-screen based systems, there will be a coexistence of these three technologies for at least the next decade. This is a grand challenge for image analysis techniques designed for mammography: For an algorithm to be widely applicable, it has to perform well on all these image types.

To obtain a digital image from a conventional film, the film is digitized using a high quality laser scanner with optical densities of 4.0. Recently, novel CCD-based scanners also provide an acceptable digital image quality. The spacial resolution of film is about 13 to 15 line pairs per millimeter (LP/mm). Mammography scanners can digitize at a spacial resolution of 10 LP/mm. This corresponds to 50 microns. A film-based system has as a non-linear transfer function, due to film sensitivity depending on the dose.

Since structures indicating cancer might be very tiny, image quality requirements in digital mammography are extremely high. This was the main reason that for a long time there has been no adequate digital alternative to conventional film-screen mammography. Actually, film-screen technology used in conventional mammography is the last remaining non-digital imaging modality used in radiology. First digital systems were based on the storage of phosphor, so called digital luminescence or computed radiography (CR). However, at normal dose these systems have a relatively poor signal-to noise ratio. Higher contrast is an advantage of CR systems.

Novel direct digital systems for mammography based on silicium or selenium detectors have been introduced in the last three years. The main advantage of full field direct digital mammography (FFDM) is high quanta efficiency. This allows for higher contrast-detail resolution at reduced dose. Further there is a linear relationship between dose and detector signal. This simplifies image processing based on dose models. Further there is significantly lower noise, because of less processing steps. A problem is still relatively low spacial resolution of only 5 up to 8 LP/mm. Thus very small microcalcifications are blurred or can not be detected at all.

Figure 3 shows an example of microcalcifications in digitized film versus a direct digital mammogram. The example hints the robustness that is required for an enhancement algorithm to be able to handle both types of images.
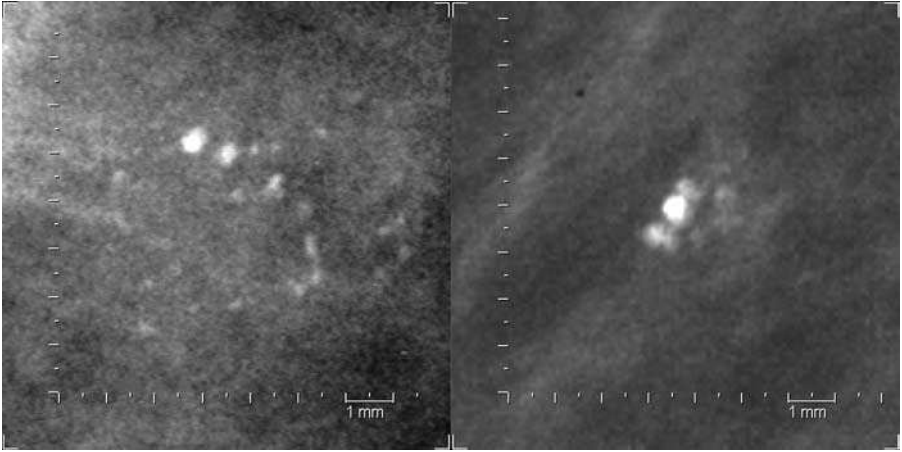
Fig. 3. Spacial resolution. Examples of clearly visible microcalcifications in film digitized at 10 LP/mm (a) and in a direct digital mammogram at 5 LP/mm (b). While detail resolution is better in the digitized mammogram, there is less noise, especially no film grain, in the direct digital image.

2.3. *Image formation*

In this section, we give a short explanation of how a digital mammographic image is generated. From this we develop a model that allows for the relevant imaging effects relevant for microcalcifications. These effects are noise and unsharp projection as described by K. Even.[16] In Sec. 5.1 we apply this model to design an adapted filter.

Originating from a radiation source, an x-ray beam passes through the compressed breast. An intensifying screen is used as a signal amplifier before the dose is measured by a flat detector. The x-ray beam is attenuated depending on the absorption properties of the breast tissue. Calcium absorbs tens of times more x-ray photons than does other breast tissue. Thus x-ray mammography can make calcifications visible in the image.

This capability to visualize calcifications at a high spacial resolution is the main reason why clinicians use x-ray mammography for early breast cancer detection. Other modalities such as ultrasound and especially MR are better suited to distinguish malign tumors from healthy tissue. But they can not visualize microcalcifications, the earliest signs of cancer.

Main sources for noise are the limited number of x-ray quanta which form a Poisson-distributed process, film granularity, varying absorption due to changing energy of the photons on their path through the breast, and random inhomogeneities in or on the intensifying screen.

Classical x-ray image formation generates geometric distortions. The main effect is the projection of a three-dimensional anatomical structure onto image plane. This leads to superimposed tissue regions which can be a difficult source for artifacts. A magnification $V := D/O$ results from the quotient of the distance from the focal

spot to the film O and the distance from the focal spot to the object D. This magnification is about 1.1 to 1.4 for standard systems. The source of the x-ray beam has a diameter B. For mammography system the focal spot size is typically $0,36$ to $0,6$ mm,[16,24] the area Br is about 0.01 to 0.2 mm.[2] As a result there is a blur $U := B(V-1)$ of about $0,001$ mm and $0,08$ mm. microcalcifications have a diameter of about 0.05 to 1 mm, this blur is of almost the same order as the size of microcalcifications.

3. Method

In the last years, huge effort was spent on the enhancement and detection of microcalcifications.

It became clear quite early that mammograms cannot be enhanced by "from the book" global or fix-neighborhood techniques due to their lack of adaptiveness. Most conventional enhancement techniques enhance not only the microcalcifications but also the background and noise. H. D. Cheng, X. Cai, X. Chen, L. Hu and X. Lou[8] give a good overview of this history.

3.1. *Known approaches*

Microcalcifications appear at a range of varying sizes. Thus it is natural to approach this problem by multiscale techniques. Among the more recently developed techniques, most share a feature based approach based on multiscale filterbank decompositions. The most successful methods apply filterbanks that are variations of the standard discrete wavelet transform decomposition.

All authors used digitized film to evaluate their results. Thus evaluation of robustness was limited to that imaging technology. We go one step further, and derive a method that we evaluated on film as well as on images from direct digital modalities.

H. Yoshida *et al.*[42] apply a discrete wavelet transform (DWT) with dyadic scales. They multiply every wavelet scale by a weight factor. Then they reconstruct an image by applying the inverse transform. The weights are determined by supervised learning, using a set of training cases. This approach results in an overall enhancement of edges and structures. There is no coefficient selection scheme in wavelet domain. Further, the DWT is not translation covariant. Thus if the origin of the image is shifted, then the result is inherently different.

R. Strickland *et al.*[38] use the discrete wavelet transform (DWT) with biorthogonal spline filters. To overcome the restriction of dyadic scales and to adapt the transform better to microcalcifications they abandon the reconstruction property. They compute four dyadic scales plus two additional interpolating scales (voices). On every wavelet scale a binary threshold-operator is applied. The responses of the individual wavelet scales are then combined by the rule of probability summation. The output is used as a feature for detection of microcalcifications. Despite being a very simple algorithm, the detection results of R. Strickland *et al.* demonstrate the power of a wavelet-based approach. However, due to the thresholding, there

is no way to reconstruct an image from the filterbank used. R. Strickland *et al.* also develop a matched filter model that we apply in Sec. 5.1 to determine wavelet functions that are appropriate for detection of microcalcifications.

L. P. Clarke *et al.*[9,34,35] apply a denoising step to the image and then take the highpass scale of a discrete wavelet transform using spline wavelets. This results in a general edge detector which finds calcifications but also several other structures such as film artifacts or lines.

J. K. Kim *et al.*[25] apply several directional Sobel operators. The results are weighted, depending on local variance, and then are recombined. An important advantage of this approach is the application of directional features with the capability to discriminate line-like structures from spot-like structures. Again, here the common approach is applied, consisting of applying a set of filters, modifying the coefficients, and recombining the result to an image. Unfortunately, the Sobel operators do not allow a perfect reconstruction. Thus this method, as is, is not suitable for enhancement of isolated structures such as microcalcifications. Due to the fact that the Sobel filters detect edges and are not specific with respect to microcalcifications, a preprocessing step to remove film artifacts is needed. This is done by setting pixels that differ from the neighborhood by values larger than a given threshold to the neighborhood average. Results of a receiver operating characteristic (ROC) analysis show, that this preprocessing step is an essential improvement to the method. The enhancement method of J. K. Kim *et al.* works fine on a phantom which consists of pixel size structures. It fails with larger structures, due to the fact that the Sobel filters are not scalable.

A. Laine *et al.*[26,27] apply several wavelet-type filterbank decompositions such as the dyadic wavelet transform, also called Mallat-Algorithm, as described in S. Mallat and W. L. Hwang.[29] An adaptive enhancement operator is defined on the wavelet coefficient scales. A. Laine *et al.* carefully design the adaptive enhancement operator in a very robust manner. Evaluating the method on phantoms they obtain good contrast improvement values for irregular structures such as microcalcifications. The enhancement operation is defined for each scale separately. However, A. Laine *et al.* do not exploit the correlation of wavelet coefficients over scales. Furthermore, the dyadic wavelet transform does not allow flexibility in the choice of discrete scales. In Sec. 5.2 we apply a variation of their adaptive enhancement operator which will be tailored specifically to detect microcalcifications.

There are extensions to the discrete wavelet transform, that allow interpolating scales, also called voices, as for example described by A. Aldroubi and M. Unser.[1] There, the original dyadic wavelet scheme is repeated with auxiliary dilated wavelet functions. Thereby, for every dyadic scale a_0 computed, one obtains additional scales $a_i = a_0 \cdot 2^{i/M}$ with $i = 0, \ldots, M - 1$. The drawback of such an approach is the calculation of a rising number of irrelevant scales: in an iterative scheme one cannot omit computing intermediate scales. In 1-D signal analysis this may not be a serious problem. Applied to images such as mammograms, where each scale costs 40 MB of storage, such an approach would be very inefficient.

3.2. *The role of wavelets*

For local image enhancement, wavelet-based multiscale methods appear among the most successful techniques; an observation not limited to the enhancement of mammograms. Why is this the case?

One basic aim in image analysis is to decompose an image into well localized, interpretable components that make local features in the image easily accessible. Especially for mammograms, as was pointed out in Sec. 2.3, we need a feature extraction method that is robust with respect to images coming from different image sources such as digitized film or direct digital mammography.

A local decomposition of an image can be accomplished by the two-dimensional continuous wavelet transform (CWT) over the Euclidean group with dilation $G := \mathbb{R}^2 \rtimes (\mathbb{R}^+ \times SO(2))$. This transform was investigated in depth by J.-P. Antoine, P. Carette, R. Murenzi and B. Piette.[4] The CWT provides an image decomposition into coefficients that describe the local geometry of the image in terms of scale and orientation. Further, the continuous wavelet decomposition is flexible with respect to image resolution as well as the size of the objects of interest.

However, it is not trivial to translate the continuous model into a discrete algorithm. Analyzing the approaches cited above, we find that in most cases limitations arise from the discrete wavelet transform applied. How can we fix this problem? Obviously, there is need for a discrete wavelet filterbank that is less restrictive and provides more flexibility to support model based approaches to enhancement of microcalcifications.

To discretize means to compute the wavelet coefficient on a discrete subset of the continuous group only. The resulting discrete transform should retain the interesting properties of the continuous wavelet transform. For image enhancement in digital mammography the most important properties are invertibility, translation covariance and flexibility in the choice of discrete scales.

Invertibility of the transform provides a means to reconstruct the image from its discrete wavelet coefficient. This implies that there is no information lost in the decomposition step.

Translation covariance means, that if the original image is shifted, the wavelet coefficient should also only be shifted. This leads to equidistant sampling of translations, independent of scales, and forces a redundant transform, compare A. Teolis.[39]

Details of image structures often exist between commonly used dyadic scales. For precise localization in scale we thus need flexibility to discretize scale and orientation depending on the specific problem. The standard discrete wavelet transform (DWT) only fulfills the first requirement, i.e. invertibility. The dyadic wavelet transform or Mallat-Algorithm is also translation covariant. But none of the classic discrete versions of the continuous wavelet transform provide flexibility with respect to scale and choice of wavelet.

However, this problem can be solved: There is a suitable method, so called integrated wavelets, that was proposed first by M. Duval-Destin, M. A. Muschietti
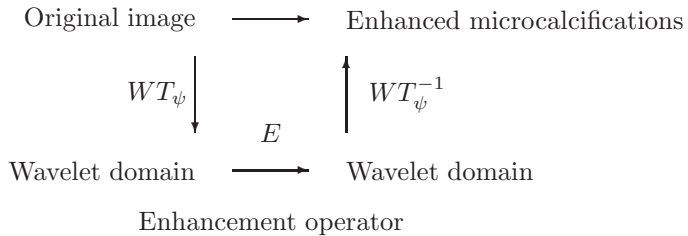
and B. Torrésani[15] for affine wavelets. We extended their original approach to wavelet transforms over semidirect product groups $G = \mathbb{R}^m \rtimes H$ and provided an implementation.[20,21] An introduction to integrated wavelets is given in Sec. 4.

3.3. *Transform based enhancement*

It is a classical approach in image processing to transform an image into a representation that is more suited for solving a specific problem than working with its representation on the original coordinate grid. For enhancement of microcalcifications in digital mammograms we propose the use of wavelets, especially the integrated wavelet transform.

The method consists of the following steps: At first, we compute a multiresolution decomposition of the image into wavelet coefficients, using an appropriate wavelet transform WT. Then we define a local enhancement operator E and apply it to the wavelet coefficients. Finally, we reconstruct the enhanced image using perfect reconstruction filters generated by a reconstruction operator WT^{-1}.

Our aim is enhancement of microcalcifications in digital mammograms without introducing artifacts. We show in Sec. 5, that using the integrated wavelet transform we can improve other methods that were based on dyadic wavelets. The wavelet plays the role of a matched filter and determines the magnitude of the wavelet coefficients, especially the points of local extrema. A model based approach leads us to a Gaussian wavelet that matches the unsharp projection of a microcalcification on the film plane in a noisy background. The enhancement operation E on the wavelet coefficients is defined by a heuristic rule. A. Laine *et al.*[27] introduced so called spot enhancement. As operation on wavelet coefficients we use a refined version.[22]

$$
\begin{array}{ccc}
\text{Original image} & \longrightarrow & \text{Enhanced microcalcifications} \\
WT_\psi \Big\downarrow & & \Big\uparrow\, WT_\psi^{-1} \\
& E & \\
\text{Wavelet domain} & \longrightarrow & \text{Wavelet domain}
\end{array}
$$

Enhancement operator

The main components of our enhancement procedure are:

(1) A discrete wavelet transform WT. We apply the integrated wavelet transform which is a refinement of wavelet based analysis. It allows a flexible and adapted discretization of scales. This transform is discussed in Sec. 4.

(2) The wavelet ψ. It plays the role of a matched filter and determines the magnitude of the wavelet coefficients, especially the points of local extrema. We discuss design of the wavelet in Sec. 5.1.

(3) The enhancement operation E on wavelet coefficients. In Sec. 5.2 we review the enhancement operator by A. Laine and present refinements using additional knowledge of the behavior of wavelet coefficients.

The quality of the enhancement result depends on all these ingredients. We discuss the influence of each ingredient separately.

4. Integrated Wavelet Transform

For analysis we are interested in a decomposition of the image into well-localized pieces, that make local features in the image easily accessible. The CWT is well known to be a useful tool for this problem.[1,4,5] For example the two-dimensional continuous wavelet transform of the Euclidean group with dilation $\mathbb{R}^2 \rtimes (\mathbb{R}^+ \times SO(2))$, studied in detail by R. Murenzi,[31] provides the natural affine operations on images that we use in order to describe the geometry of images. In this section, we present an invertible discretization of the CWT that provides the flexibility to pick scales where the image features are located.

The question of reconstruction in signal analysis often seems to be a purely theoretical problem. Signal decompositions are typically used to compute features and no reconstruction is performed. But for the problem of image enhancement reconstruction is crucial. While the CWT provides a model to design such enhancement operators, it is the discrete version of the wavelet transform used for implementation that determines the result of this procedure.

We introduce an approach to discretize the scale parameter by local averaging of wavelet coefficients. This yields so called integrated wavelets. Integrated wavelets provide flexible discrete versions of the continuous wavelet transform (CWT). This idea firstly appeared in an article by M. Duval-Destin, M. A. Muschietti and B. Torrésani[15] on affine wavelets. We extended the technique to wavelet transforms over semidirect product groups[20,21] $G = \mathbb{R}^m \rtimes H$. Here we show how to apply the general theory to images and extend our method to the construction of directional integrated wavelets.

Our starting point is the continuous wavelet transform defined over the Euclidean group with dilation $G := \mathbb{R}^2 \rtimes H$ with dilation group $H := \mathbb{R}_+^* \times SO(2)$ in two dimensions, compare J.-P. Antoine, P. Carette, R. Murenzi and B. Piette.[4] $SO(2)$ denotes the group of rotations in the plane. Let $g = (b, \rho, a) \in G$, where $b \in \mathbb{R}^2$, $a \in \mathbb{R}_+^*$ and $\rho \in SO(2)$. For $f \in L^2(\mathbb{R}^2)$ we introduce operators for translation $T_b f(x) := f(x - b)$, dilation $\mathcal{D}_a f(x) := a^{-1} f(a^{-1} x)$, rotation $R_\rho f(x) := f(\rho^{-1}(x))$ and involution $\tilde{\psi}(x) := \bar{\psi}(-x)$. We write $\mathcal{U}_g$ for $T_b R_\rho \mathcal{D}_a$. A function $\psi \in L^2(\mathbb{R}^2)$ is admissible (or a wavelet), iff

$$0 < \int_H \left| \hat{\psi}(\omega_0 h) \right|^2 d\mu_H(h) =: c_\psi < \infty,$$

for some $\omega_0 \in \widehat{\mathbb{R}^2} \setminus \{0\}$. Here $d\mu_H = a^{-1}\, da\, d\rho$ denotes the left Haar-measure of H. The continuous wavelet transform of $f \in L^2(\mathbb{R}^2)$ with respect to a wavelet $\psi \in L^2(\mathbb{R}^2)$ then is defined by

$$\mathcal{WT}_\psi f(g) := \langle f, \mathcal{U}_g \psi \rangle.$$

Important properties of the continuous wavelet transform are

(i) isometry, i.e. energy conservation:

$$c_\psi^{-1} \|\mathcal{WT}_\psi f\|_{L^2(G)}^2 = \|f\|_{L^2(\mathbb{R}^2)}^2,$$

(ii) a reconstruction formula

$$f(x) = \frac{1}{c_\psi} \int_{\mathbb{R}^2} \int_0^\infty \int_{SO(2)} \mathcal{WT}_\psi f(g)\, \mathcal{U}_g \psi(x)\, d\rho\, \frac{da}{a^3}\, db, \quad \text{and}$$

(iii) covariance under translation and dilation,

$$\mathcal{WT}_\psi \circ \mathcal{U}_g = \pi_g \circ \mathcal{WT}_\psi, \quad g \in G,$$

i.e. if we shift or scale the image, the wavelet coefficient is shifted or scaled but does not change its structure.[a] Here π denotes the left regular representation of G. More details on the construction of continuous wavelet transforms can be found in H. Führ[18] and the book of S. T. Ali, J.-P. Antoine and J.-P. Gazeau.[2]

The classical approach to construct discrete transforms from the continuous wavelet transform is by restriction to a discrete subset of G. It leads to a discrete transform such as the dyadic wavelet transform which is based on pairs of wavelet filters for analysis and synthesis. The drawback of this approach is that explicit examples of such pairs of filters are only known for dyadic or equidistant scales. In the case of arbitrarily chosen discrete scales, there exist only slow iterative schemes to compute the associated reconstruction filters, such as the extrapolated Richardson or conjugate gradient methods, compare S. Mallat.[28]

4.1. *Integrated wavelets*

For precise localization in scale we need freedom to compute scales depending on the image and the features we are interested in. Geometric scales, as used in pyramidal decompositions, often are not a natural choice. For example, interesting details may to be concentrated on a small range of scales. This is exactly the case for microcalcifications.

---

[a]Note that covariance is different from invariance — a property also often demanded from features for image analysis. An invariant operation S would not change the outcome at all, i.e. $S \circ \mathcal{WT} = \mathcal{WT}$. However, there are several authors who use the term "translation invariance" instead of "translation covariance".

The integrated wavelet transform is a discrete version that provides the flexibility to pick scales where the signal features are located. It is a translation covariant and invertible discretization of the CWT. Thus there exists a reconstruction formula. Moreover, we show that integrated wavelets also provide directional sensitivity. This feature can be exploited more easily and accurately than compared to just sampling the CWT.

We discretize the scale parameter by averaging over the wavelet coefficient. We interpret the continuous wavelet transform as a continuous partition of unity in Fourier domain generated by the wavelet, formally

$$\hat{f}(\omega) = \frac{1}{c_\psi} \left(\int_H |\hat{\psi}(\omega h)|^2 d\mu_H(h) \right) \cdot \hat{f}(\omega) \quad \text{for a. e. } \omega \in \mathbb{R}^2. \tag{1}$$

The idea of integrated wavelets is to decompose this partition into discrete blocks. Therefore we split the integral over H in the admissibility condition into parts. Let $(a_j)_{j\in\mathbb{Z}}$ be a decreasing sequence in $\mathbb{R}_+^*$ with $\lim_{j\to-\infty} a_j = \infty$ and $\lim_{j\to\infty} a_j = 0$. Further, let $(K_l)_{l\in L}$ be a measurable partition of $SO(2)$, with L a discrete index set. Let $\psi \in L^2(\mathbb{R}^2)$ be a wavelet. We define the integrated wavelet $\Psi^{j,l} \in L^2(\mathbb{R}^2)$ in Fourier domain by

$$\left|\widehat{\Psi^{j,l}}(\omega)\right|^2 := \frac{1}{c_\psi} \int_{a_{j+1}}^{a_j} \int_{K_l} \left|\widehat{D_a R_\rho \psi}(\omega)\right|^2 d\rho \frac{da}{a}, \quad j \in \mathbb{Z}, \ l \in L, \ \omega \in \mathbb{R}^2. \tag{2}$$

Then the integrated wavelet transform is defined by

$$WT_\psi^I : L^2(\mathbb{R}^2) \to l^2(L^2(\mathbb{R}^2 \times \mathbb{Z} \times L)),$$
$$f \mapsto \langle f, T_b \Psi^{j,l}\rangle = f * \widetilde{\Psi^{j,l}}(b).$$

By construction of the integrated wavelet we have $\sum_{j\in\mathbb{Z},l\in L} |\widehat{\Psi^{j,l}}(\omega)|^2 = 1$ almost everywhere. Thus we obtain the isometry

$$\sum_{j\in\mathbb{Z},l\in L} \left\|\langle f, T_\bullet \Psi^{j,l}\rangle\right\|^2 = \|f\|^2.$$

This isometry describes a semi-discrete tight frame. It thus provides a reconstruction of f in L^2-norm by

$$f(x) = \sum_{j\in\mathbb{Z},l\in L} WT_\psi^I f(x,j,l) * \widetilde{\Psi^{j,l}}. \tag{3}$$

Compared to results on the construction of frames by sampling, see for example S. Mallat,[28] we observe that for every wavelet we obtain a reconstruction.[20] Moreover, we do not depend on a special choice for the partition of scales.

Examples

As a first example we study the Gaussian wavelet on $\mathbb{R}^2$. It is given by $\psi(x) := \frac{1}{\sqrt{2\pi}}(1 - \|x\|^2)e^{-\|x\|^2}$, in Fourier domain by $\hat{\psi}(\omega) = \|\omega\|^2 e^{-\|\omega\|^2/2}$. The associated
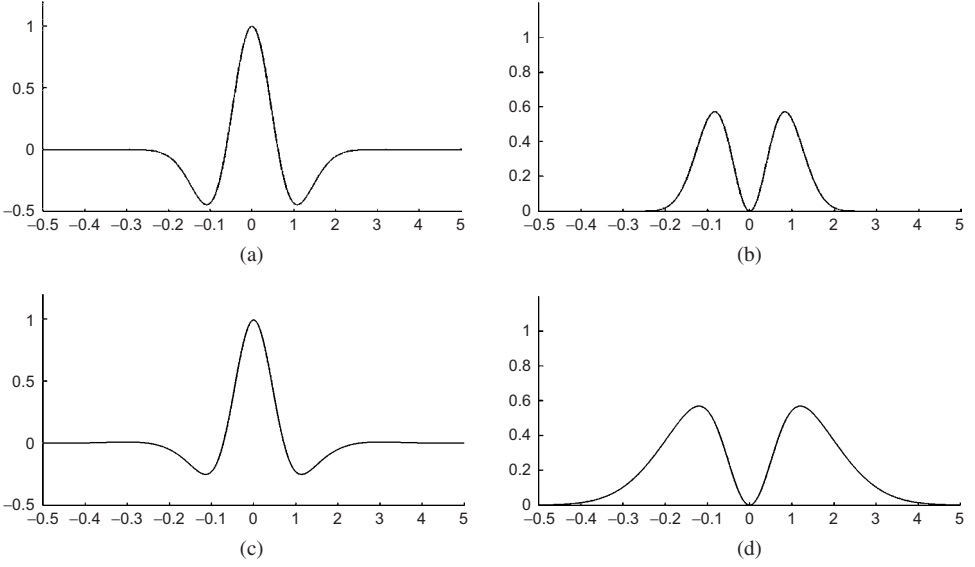
Fig. 4. Gaussian wavelet ψ in time domain (a) and Fourier domain (b). Integrated Gaussian wavelet Ψ^0, $a_0 = 1$, $a_1 = 2$, in time domain (c) and Fourier domain (d).

integrated wavelet, as given in formula (2), is

$$\left|\widehat{\Psi^j}(\omega)\right|^2 = \frac{1}{2}\left(((a_j\|\omega\|)^2 + 1)e^{-(a_j\|\omega\|)^2}\right.$$
$$\left. - ((a_{j+1}\|\omega\|)^2 + 1)e^{-(a_{j+1}\|\omega\|)^2}\right). \tag{4}$$

We observe that if we solve the integral in formula (2) explicitly once, then we can generate integrated wavelets for arbitrary scales from (4) by just changing the value of a_j. Figure 4 illustrates the Gaussian wavelet and a corresponding integrated wavelet.

As a second example we take the Poisson wavelet on $\mathbb{R}$. It is given by $\psi(x) := (2\pi)^{-2}\left((1+ix)^{-2} + (1-ix)^{-2}\right)$; in Fourier domain $\hat{\psi}(\omega) = |\omega|e^{-|\omega|}$. The associated integrated wavelet is

$$\left|\widehat{\Psi^j}(\omega)\right|^2 = \left(\frac{1}{2}a_{j+1}|\omega| + \frac{1}{4}\right)e^{-2a_{j+1}|\omega|} - \left(\frac{1}{2}a_j|\omega| + \frac{1}{4}\right)e^{-2a_j|\omega|}. \tag{5}$$

Figure 5 illustrates the Poisson wavelet and a corresponding integrated wavelet.

4.2. *Directional wavelets*

Theoretically we can choose the wavelet arbitrarily. Unfortunately, for some directional wavelets it might be difficult or impossible to solve the double integration over scale and orientation in the definition of the integrated wavelet (2) explicitly. We use wavelets which are separable in polar coordinates, since they allow explicit
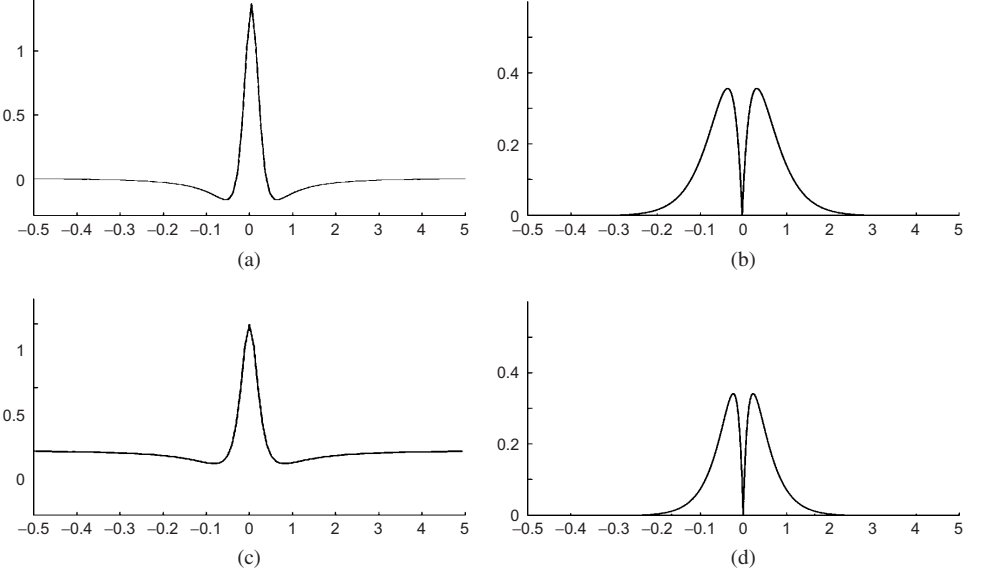
Fig. 5. Poisson wavelet ψ in time domain (a) and Fourier domain (b). Integrated Poisson wavelet Ψ^0, $a_0 = 1$, $a_1 = 2$, in time domain (c) and Fourier domain (d).

expressions for the integrated wavelet. Let the wavelet $\psi \in L^2(\mathbb{R}^2)$ be given by

$$\hat{\psi}(\omega) = \mu(|\omega|)\eta(\arg(\omega)), \quad \forall \omega \in \widehat{\mathbb{R}^2},$$

with functions $\mu : \mathbb{R} \to \mathbb{C}$ and $\eta : S_1 \to \mathbb{C}$. We obtain the corresponding integrated wavelet $\Psi^{j,l} \in L^2(\mathbb{R}^2)$, $(j,l) \in \mathbb{Z} \times L$, by

$$\left|\widehat{\Psi^{j,l}}(\omega)\right|^2 = \frac{1}{c_\psi} \int_{a_{j+1}}^{a_j} \left|\mu(|\omega a|)\right|^2 \frac{da}{a} \int_{K_l} \left|R_\rho \eta(\arg(\omega))\right|^2 d\rho, \quad \omega \in \widehat{\mathbb{R}^2}. \qquad (6)$$

Thus we can solve the integrals over scale and orientation separately. Since this class of wavelets is separable in polar coordinates it is better adapted to the geometric structure of images than functions separable in cartesian coordinates. The latter are more adapted to the tensor-product model of images, which is more vulnerable to producing artifacts.

As an example we construct a polar separable directed integrated wavelet that is similar to the 2-D directional Cauchy wavelets analyzed by J.-P. Antoine et al.[3,6] Starting with the Poisson wavelet, we define the directional wavelet by

$$\hat{\psi}(\omega) := \|\omega\| e^{-\|\omega\|} \eta(\arg(\omega)), \quad \omega \in \widehat{\mathbb{R}^2}.$$

The oriented component η is given by

$$\eta(\rho) := \begin{cases} \cos^2(\rho/\alpha) & \text{if } \rho \in [-\alpha\frac{\pi}{2}, \alpha\frac{\pi}{2}], \\ 0 & \text{if } \rho \in [-\pi, \pi] \backslash [-\alpha\frac{\pi}{2}, \alpha\frac{\pi}{2}]. \end{cases}$$
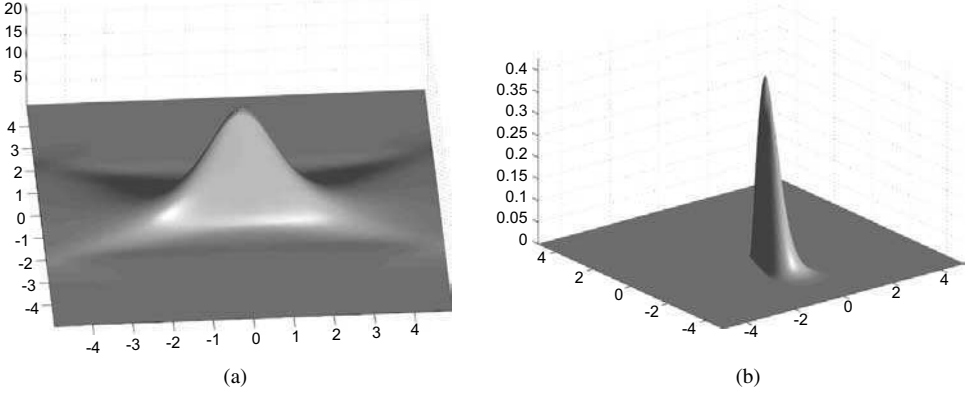
Fig. 6. Directional wavelet separable in polar coordinates. Absolute value in time domain (a) and in Fourier domain (b).

with $0 < \alpha < 1$. Thus the support of the wavelet in Fourier domain is restricted to a convex cone with its top at the origin and opening angle $\alpha\pi$. Let $K_l = [l_1, l_2]$ be an interval on the circle S_1. The directional component of the integrated wavelet then is given by

$$\left|\widehat{\Psi^l}(\rho)\right|^2 = \Xi(l_1) - \Xi(l_2),$$

with

$$\Xi(\rho) = \begin{cases} 0 & \text{if } \rho \leq -\alpha\pi/2, \\ \alpha\sin(4\rho/\alpha)/32 + \alpha\sin(2\rho/\alpha)/4 + 3\rho/8 + 3\pi\alpha/16 & \text{if } -\alpha\pi/2 \\ & \quad\quad\quad < \rho < \alpha\pi/2, \\ 1 & \text{if } \rho \geq \alpha\pi/2. \end{cases}$$

We obtain the integrated directional wavelet on $\mathbb{R}^2$ by $\widehat{\Psi^{j,l}}(\omega) = \widehat{\Psi^j}(|\omega|)\,\widehat{\Psi^l}(\arg(\omega))$, $\omega \in \widehat{\mathbb{R}^2}$. Figure 6 shows the directional wavelet in time and Fourier domain. This wavelet, like the Cauchy wavelet, fulfills the minimality condition given in J.-P. Antoine *et al.*[6] [Proposition 4.3] and thus has minimal uncertainty with respect to angular selectivity. We see clearly from Fig. 6(a) the directional selectivity in time domain. Figure 6(b) displays the optimal angular selectivity in frequency domain that is steerable by the parameter α.

    Our approach emphasizes the wavelet as a template that has directional sensitivity, while the transform is just a generic discrete version of the CWT. One can simplify the construction, compromising the role of the wavelet and the role of the transform. Following J.-P. Antoine *et al.*,[4] let us start with an isotropic wavelet ψ, which generates an isotropic integrated wavelet Ψ^j. Take a partition of unity on the circle given by $\eta_x(\rho) := \cos^2(\rho)$ and $\eta_y(\rho) := \sin^2(\rho)$. Then oriented wavelets $\Psi^{j,l}$ are given in Fourier domain by

$$\widehat{\Psi^{j,l}}(\omega) := \widehat{\Psi^j}(|\omega|)\eta_l(\arg(\omega)), \quad \forall\omega \in \widehat{\mathbb{R}^2},\ l \in \{x, y\}.$$
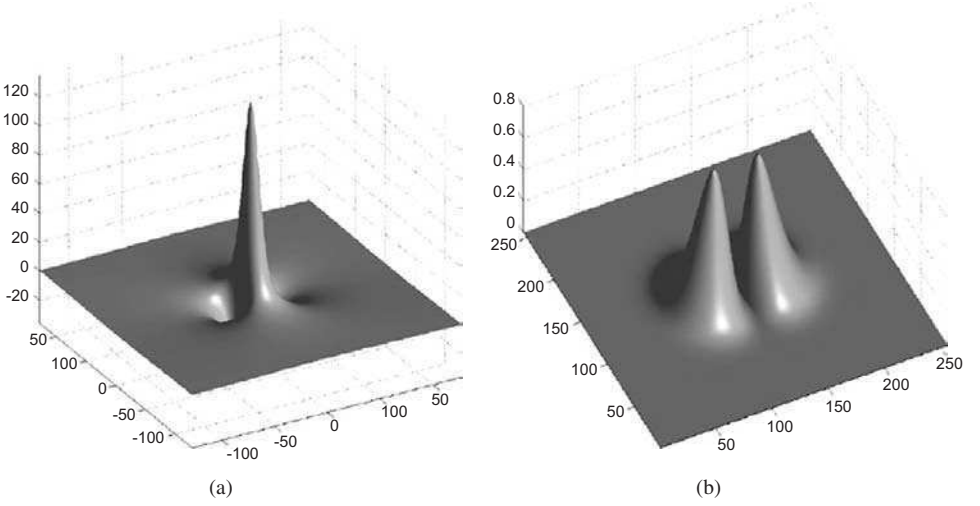
Fig. 7.   Integrated directional filter. (a) Directional Gaussian wavelet in space domain and (b) in Fourier domain.

These functions are separable in polar coordinates, and thus better adapted to the geometric structure of images than functions separable in cartesian coordinates. Figure 7 shows the directional wavelet constructed from the integrated Gaussian wavelet from Formula (4).

4.3. *Discrete algorithm*

The key to reconstruction from integrated wavelet coefficients is the partition of unity in Fourier domain from Eq. (1). Thus it is no surprise that our algorithm for the integrated wavelet transform is based on computation in Fourier domain.

An algorithm can compute only a finite number of scales. If the integral in formula (2) also exists over the interval $[a_j, \infty[$, we can define a scaling function that acts as a low-pass filter. Let $\psi \in L^2(\mathbb{R}^2) \cap L^1(\mathbb{R}^2)$ be a wavelet. Under appropriate conditions,[21] we can define a scaling function $\phi \in L^2(\mathbb{R}^2)$ which is defined by its Fourier transform as

$$\left|\hat{\phi}(\omega)\right|^2 := \frac{1}{c_\psi} \int_{\|\omega\|}^{\infty} \int_{SO(2)} \left|\widehat{D_a R_\phi \psi}(\omega)\right|^2 d\rho \frac{da}{a}. \tag{7}$$

From this we obtain a simplified partition of unity

$$\left|\widehat{D_{a_k}\phi}(\omega)\right|^2 + \sum_{j=k}^{\infty} \left|\widehat{\Psi^j}(\omega)\right|^2 = 1, \quad \text{for all } \omega \in \widehat{\mathbb{R}^2}. \tag{8}$$

Up to now, we thought of the signal as a continuous function of $\mathbb{R}^2$. For an implementation we have to work with discrete functions from a space $l^2(\mathbb{Z}_n^2)$ with some $n \in \mathbb{N}$. ($l^2(\mathbb{Z}_n^2)$ can be identified with the space $\mathbb{C}^{n \times n}$.) Let us define a sampling operator $S : L^2(\mathbb{R}^2) \to l^2(\mathbb{Z}_n^2)$ by $S(f)(k) := f(k - \lfloor n/2 \rfloor)$, $0 \le k < n$.

Thus we identify Z_n^2 with a regular subset of $\mathbb{R}^2$ centered around 0. We denote $H_{j,l} := [a_{j+1}, a_j] \times K_l \subset H$ and $H_j := [a_{j+1}, a_j] \times SO(2)$. We now define the discrete integrated wavelet $\Psi^{j,l} \in l^2(\mathbb{Z}_n^2)$ by equidistant sampling of its continuous version in Fourier domain. Let F denote the discrete Fourier transform on $\mathbb{Z}_n^2$. Fix $j_0 \in \mathbb{Z}$. The discrete integrated wavelet transform WT^D of $f \in l^2(\mathbb{Z}_n^2)$ with respect to an admissible vector $\psi \in L^2(\mathbb{R}^2) \cap L^1(\mathbb{R}^2)$ is defined by

$$WT_\psi^D f(b, 0) := w_0 \, F^{-1}\left(Ff \cdot \overline{S\widehat{D_{j_0}\phi}} \right)(b), \quad b \in \mathbb{Z}_n^2,$$

$$WT_\psi^D f(b, j, l) := w_{j,l} \, F^{-1}\left(Ff \cdot \overline{S\widehat{\Psi^{j,l}}} \right)(b), \quad b \in \mathbb{Z}_n^2, \; j \geq j_0, \; l \in L.$$

Here the continuous integrated wavelets $\Psi^{j,l}$ are sampled in the Fourier domain. For the scales to approximate the continuous transform, we introduce weights $w_{j,l} := \mu_H(H_{j,l})^{-1/2}$. By linearity of the discrete Fourier transform F we obtain a reconstruction formula

$$f(x) = w_0^{-1}\left\langle WT_\psi^D f(\bullet, 0), T_x F^{-1} S\widehat{\phi} \right\rangle_{l^2(\mathbb{Z}_n^2)}$$
$$+ \sum_{j \geq j_0} \sum_{l \in L} w_{j,l}^{-1}\left\langle WT_\psi^D f(\bullet, j, l), T_x F^{-1} S\widehat{\Psi^{j,l}} \right\rangle_{l^2(\mathbb{Z}_n^2)}, \quad x \in \mathbb{Z}_n^2.$$

This gives a finite filter bank, consisting of (i) the low pass filter which is the scaling function $\widehat{D_{a_{j_0}}\phi}$, (ii) integrated wavelet filters $\widehat{\Psi^j}$, $j_0 \leq j \leq j_1$, and (iii) a high pass filter given by sampling $1 - \widehat{\Psi^{j_1}}$. The last filter is not a function in $L^2(\mathbb{R}^2)$, but it is well defined as function in $l^2(\mathbb{Z}_n^2)$.

This construction has several advantages over the standard dyadic wavelet transform. It allows us to compute arbitrary discrete scales, i.e. we can adapt the scales $a_{j_0}, \ldots, a_{j_1}$ to fit the size of the feature we are interested in. We have further the flexibility to choose a wavelet, which is adapted to a specific problem.

The computational load of this approach is dominated by the Fourier transform. To compute J scales we have to compute $J+1$ Fourier transforms. The same is true for reconstruction. The typical optimizations are possible. By using real wavelets, for example, we can compute two scales with only one complex Fourier transform.

In Fig. 8 we give an example of a 1-D signal decomposition using integrated wavelets.

The signal f is Gaussian. As a wavelet we use the integrated Gaussian. Figure 8(a) shows the classical dyadic wavelet decomposition, i.e. with scales in powers of 2. Scales are plotted on logarithmic axis. In Fig. 8(b) we use an irregular discretization of scales adapted to the expected maximum correlation of the wavelet with our signal. We can reconstruct the original signal from both, the dyadic and the irregular decomposition. In both cases we used the same number of 11 scales. With the adapted discretization we can determine the scale and value of peak correlation with higher accuracy.
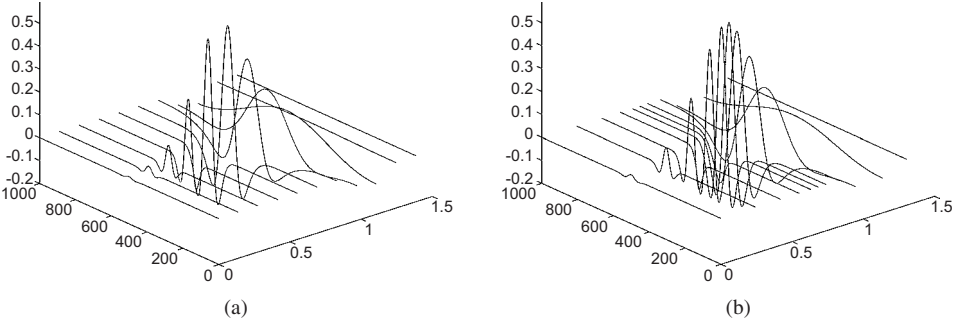
Fig. 8. Example of the integrated wavelet transform. The figure shows the wavelet transform of a Gaussian using the integrated Gaussian wavelet. (a) Dyadic wavelet decomposition using scales in powers of 2. (b) Irregular wavelet decomposition adapted to the expected maximum correlation. In both cases 11 scales are computed. The signal consists of 1000 samples. Scales are plotted on a logarithmic axis.

In the case of dyadic scales, specialized algorithms for the affine group exist. An interesting pyramidal algorithm using convolution in real space is given in M. A. Muschietti and B. Torrésani.[32] The exact integrated filters are replaced by approximate filters that fulfill a two-scale equation. With this replacement the authors give up the tight frame condition and eventually the reconstruction property. Moreover, the algorithm is slow due do computation in time domain. P. Vandergheynst and J.-F. Gobbers[40] extended this algorithm to the 2-D similitude group and improved its performance by operating in the Fourier domain. The authors state that their algorithm is twice as fast as the standard implementation of the dyadic wavelet transform, but again they do not provide a reconstruction formula.

With integrated wavelets we can adapt the discretization of scales dynamically. In applications there often is prior knowledge about the range of scales that are the most interesting ones. Specifically, we can adapt the scales to fit the size of microcalcifications which is well known. We have further the flexibility to choose the wavelet adapted to the problem of enhancing microcalcifications.

4.4. *Extensions to general settings*

Let us note that the ideas presented here are valid in a broader context. In the general setting of wavelet transforms over semidirect product groups $G = \mathbb{R}^m \rtimes H$, we replace dilation and rotation by a general dilation group H. Because the operation of H need not be on all of $\mathbb{R}^m$, we have to replace $L^2(\mathbb{R}^m)$ with the generalized Hardy space $\mathcal{H}_V := \left\{ f \in L^2(\mathbb{R}^m)\,;\, \hat{f}(\omega) = 0 \text{ for a. e. } \omega \in \widehat{\mathbb{R}^m} \backslash V \right\}$ defined over an open H-orbit $V := \gamma H$ for some $\gamma \in \widehat{\mathbb{R}^m}$.

One can show, compare e.g. H. Führ,[18] that in this setting the admissibility condition is the same as above. Let J be a discrete countable index set. Let $(H_j)_{j \in J}$ be a partition of H into measurable compact subsets. Then the integrated wavelet $\Psi^j \in \mathcal{H}_V$ with respect to a wavelet ψ and $(H_j)_{j \in J}$ is defined in Fourier

domain by

$$\left|\widehat{\Psi^j}(\omega)\right|^2 := \frac{1}{c_\psi} \int_{H_j} \left|\hat{\psi}(\omega h)\right|^2 d\mu_H(h), \quad j \in J, \ \omega \in V. \tag{9}$$

Because H_j is compact, the integrated wavelet is again admissible. If the wavelet is band-limited, its integrated wavelet is also band-limited. The results from the special case presented above apply also for the general setting. Thus one can construct tight frames and wavelet frames for general dilation groups with arbitrary choice of discrete scales.[21]

The aim is to apply wavelet transforms with other group operations to image processing problems. Due to lack of a fast implementation, this has been difficult. However, integrated wavelets and the straightforward implementation derived in the previous section open up new possibilities, and allows to exploit these wavelet transforms.

In this section we have introduced integrated wavelets as a refinement of wavelet transform techniques tailored for difficult multiscale image analysis problems. Integrated wavelets generate adaptable discrete versions of the continuous wavelet transform (CWT). They provide a reconstruction formula for arbitrary choice of discrete scales and orientation, and they do not restrict us to a limited set of wavelets. An efficient algorithm based on the FFT provides a straightforward implementation. With the integrated wavelet transform we do not have to overcome difficulties resulting from tensor product algorithms but can use appropriate group operations such as rotation to model our problems.

Especially in medical applications the focus is on extracting features for subtle pathologies, while the penalty on artifacts introduced by an image manipulation is particularly high. Here integrated wavelets provide an excellent means to bring together the precision and the modelling power of the CWT and the practical requirement for a fast discrete implementation.

5. Enhancement of Microcalcifications

We apply the methods introduced in the last section to real world images. Performance of integrated wavelets with adapted wavelet and adapted discrete range of scales is compared with the performance of hitherto used dyadic scales and an edge detection wavelet. Further, we improve the enhancement operation on wavelet coefficients by introducing a feature which is based on local regularity.

For the following examples we use the image shown in Fig. 9(a). Figure 9(b) shows an enlarged part of the image that contains a microcalcification cluster. To compare enhancement results we compute the local contrast of microcalcifications with respect to a neighborhood. In Fig. 9(c) both microcalcifications and the neighborhood that is used are marked. The local contrast C of microcalcifications o_f in an image f with respect to the neighborhood h_f is defined by

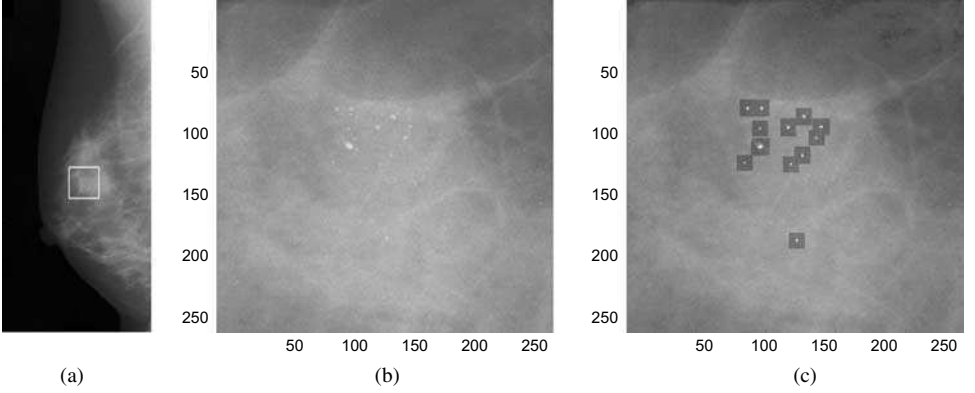$$C(f) := \frac{E(o_f) - E(h_f)}{E(f)}. \tag{10}$$

Fig. 9. Microcalcifications. (a) Digitized film image. (b) Magnified part with microcalcification cluster. (c) Microcalcification markers and neighborhood used for computation of local contrast.

We set h_f to 15×15 pixels for each calcification. We compute the local contrast on a normalized image which has the same mean and variance as the original image. The normalization reduces the dependency of the local contrast from the magnitude of the lowpass filter and allows better comparison of contrast values.

5.1. *Choosing an adapted wavelet*

We start with a model for microcalcifications. Microcalcifications MC vary in size and shape.[24] An accepted simplification is to think of them as ellipsoids of diameters between $0.05\,\text{mm}$ and $1\,\text{mm}$. Projection of such an ellipsoid MC onto the image plane results in a half-ellipsoid MC_{Proj}. Here we assume parallel X-rays. The detector measures the energy image $\exp(MC_{proj})$. Another aspect is the diameter of the focal spot which is about the same size as small microcalcifications. This can be modeled by a convolution with the characteristic function χ_U of the focal spot, i.e. $MC_{image} = \exp(MC_{proj}) * \chi_U$. We approximate MC_{image} by the Gaussian

$$\phi(x) := \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}|x|^2}, \quad x \in \mathbb{R}^2.$$

Figure 10 shows that the Gaussian is a good approximation of the modeled function. Figure 11 shows an empirical verification on cross sections of microcalcifications taken from a mammogram.

We now have a model for microcalcifications. To construct the matched filter ψ we need a stochastic model N for the background, i.e. the breast tissue. Let us follow the arguments of R. Strickland and H. Hahn.[38] They propose a texture model $N(t) = m(t) + R(t)$, $t \in \mathbb{R}^2$, with a nonstationary mean $m(t)$ and a stationary residuum R. Modelling R as ergodic random process with autocovariance $\gamma_R(x_1, x_2) := \sigma_R^2 e^{-\alpha\sqrt{x_1^2 + x_2^2}}$, $(x_1, x_2) \in \mathbb{R}^2$, yields the spectral density function $P_{RR}(\omega) = 2\sigma_R^2 \frac{\alpha}{\alpha^2 + \|\omega\|^2}$, $\omega \in \widehat{\mathbb{R}}^2$. Using the Gaussian as model for microcalcifications
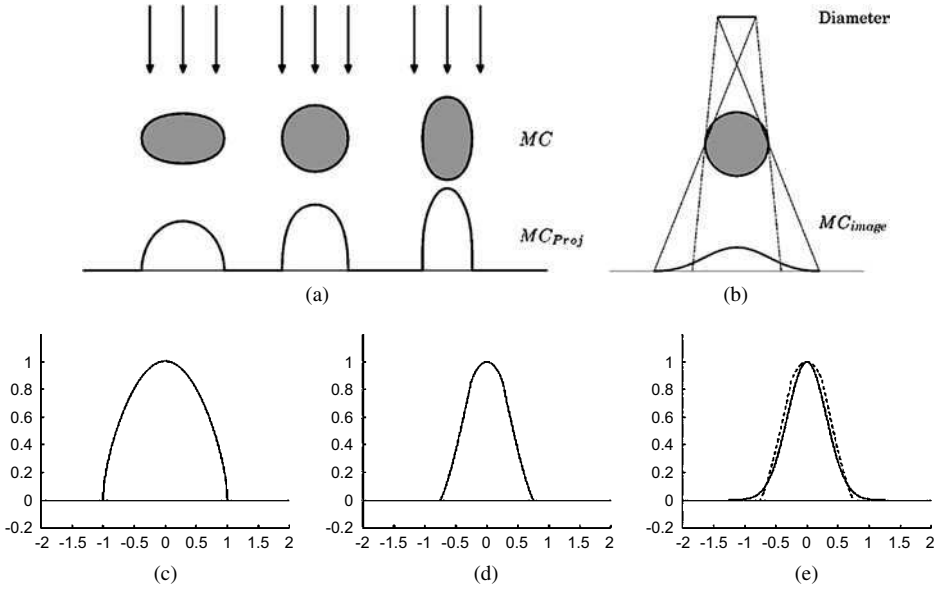
Fig. 10. Model for microcalcifications. (a) Projection of ellipsoid. (b) Unsharp radiation source. (c) Image $\exp(MC_{Proj})$ of microcalcifications. (d) Convolution with focal spot. (e) Approximation by Gaussian.
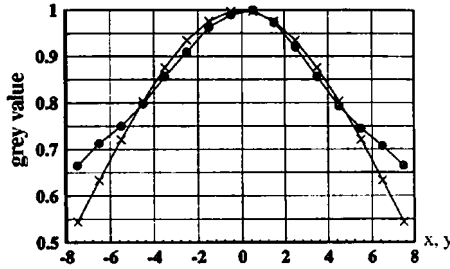
Fig. 11. Cross section of microcalcifications. Average of profiles of sample microcalcifications ($\bullet$) and Gaussian fit ($+$). The image is taken from R. Strickland and H. Hahn.[38]

we obtain a matched filter ψ_m by $\hat{\psi}_m(\omega) := \frac{\hat{\bar{\phi}}(\omega)}{P_{RR}(\omega)} = \frac{1}{2\alpha\sigma_R^2}e^{-\|\omega\|^2/2}(\alpha^2 + \|\omega\|^2)$, $\omega \in \widehat{\mathbb{R}^2}$. R. Strickland and H. Hahn[38] propose a simplification $\alpha \to 0$ based on results on the Nijmegen mammography database. This yields an approximation $P_{RR}(\omega) = \|\omega\|^{-2}$. Thus the matched filter is given by

$$\hat{\psi}(\omega) = \sigma_R^{-2}\|\omega\|^2 e^{-\|\omega\|^2/2}. \tag{11}$$

There is an analogy with a neural network based approach found in D. Rosen *et al.*[36] The template generated by training of the network looks similar to this filter.

The filter (11) is also an admissible function. We construct the integrated wavelet filters by applying formula (2). They are given in Fourier domain by

$$\left|\widehat{\Psi^{j,l}}(\omega)\right|^2 = \frac{1}{2}\left(\left(a_j^2\|\omega\|^2 + 1\right)e^{-a_j^2\|\omega\|^2}\right.$$
$$\left. - \left(a_{j+1}^2\|\omega\|^2 + 1\right)e^{-a_{j+1}^2\|\omega\|^2}\right)\eta_l(\arg(\omega)), \tag{12}$$

$\omega \in \widehat{\mathbb{R}^2}, j \in J, l \in \{0, \pi/2\}$, with highpass filter

$$\left|\widehat{\Psi^0}(\omega)\right|^2 = \frac{1}{2}\left(1 - \left(a_1^2\|\omega\|^2 + 1\right)e^{-a_1^2\|\omega\|^2}\right) \tag{13}$$

and lowpass filter

$$\left|\widehat{\Psi^{|J|}}(\omega)\right|^2 = \frac{1}{2}\left(a_{|J|}^2\|\omega\|^2 + 1\right)e^{-a_{|J|}^2\|\omega\|^2}. \tag{14}$$

A matched filter by itself does generally not allow a reconstruction of the original signal. In our approach we use the idea of matched filters to construct wavelets. By this we obtain a wavelet transform, thus an invertible transform. Further, the rich knowledge of the structure of wavelet coefficients can be exploited.

Some authors tried to detect microcalcifications with wavelets that are typically used for edge detection, for example spline-wavelets.[25,34,41] Figure 12 demonstrates that our choice of a matched filter wavelet approach yields a significantly more specific enhancement of microcalcifications than the edge wavelet. In Fig. 12(b) we observe that edges and line-like structures are ignored by the Gaussian wavelet. The enhancement operator yields a contrast improvement of 5.62. Contrast improvement with the edge wavelet is 6.04 (Fig. 12(a)). The superior local contrast value of the edge wavelet results from Gibbs-like phenomena that produce dark trenches
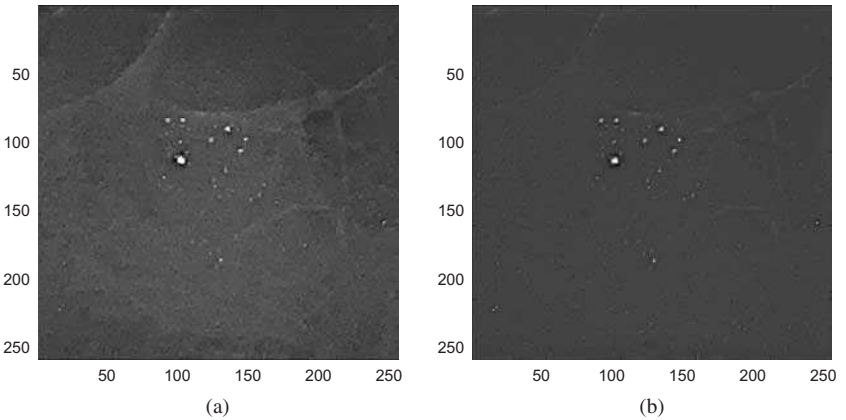

Fig. 12. Choosing the wavelet. Contrast improvement is 6.04 for the first order spline wavelet (a), and 5.62 for the Gaussian wavelet (b). Application of the Gaussian wavelet exhibits a significantly more specific enhancement than with the spline wavelet. The gain factor is set such that differences between the methods are clearly visible.

around the calcification. Especially when applied to high resolution mammograms, the adapted wavelet performs significantly better than edge wavelets. In that case small edges originate mostly from film noise and defects while calcifications have a smooth outline. Both wavelets yield comparably contrast enhancement on micro-calcifications but the choice of a matched wavelet leads to an enhancement more specific to microcalcifications.

5.2. *Enhancement operator*

Our aim is enhancement of microcalcifications without introducing artifacts. As operation on wavelet-coefficients we use a modified version of the spot enhance-ment introduced by A. Laine *et al.*[27] Our refined version also exploits the wavelet coefficient decay over scales for better separation of microcalcifications from line like structures.

The enhancement operates on the integrated wavelet coefficient which is a func-tion on the grid $\Gamma = \mathbb{Z}_n^2 \times J \times L$ with scales $J = \{0, 1, \ldots, j_{\max}\}$ and two orientations $L = \{x, y\}$. For each scale a gain factor K_j is computed. A threshold T_j and a local measures M_j of the isotropy determine whether this gain is applied to a certain point.

Let us define the enhancement operation $E : l^2(\Gamma) \to l^2(\Gamma)$ on the wavelet coefficient by

$$E(F)(b, j, l) := \begin{cases} (1 + S \cdot K_j)M_j(b)F(b, j, l), & \text{if } M_j(b) > T_j \\ F(b, j, l), & \text{if } M_j(b) \leq T_j, \end{cases}$$

with $F \in l^2(\Gamma)$, $(b, j, l) \in \Gamma$, $0 < j < j_{\max}$ and $l \in L$.

The factor $S > 0$ determines the overall strength of the enhancement. We exploit information within wavelet scales as well as the wavelet coefficient decay over scales.

Information within scales

The thresholds T_j were determined empirically by A. Laine *et al.*[27] The threshold T_j depends on the variance of the wavelet coefficient within the scale j.

$$T_j := 0.5\sqrt{\mathrm{var}\left((F(b, j, l))_{b \in \mathbb{Z}_n^2, l \in L}\right)}, \quad j \in J.$$

Essentially the threshold ignores small wavelet-coefficients that do not correlate with the matched filter.

The weight K_j is inverse proportional to the variance:

$$K_j := \frac{T_{\max}}{T_j} - 1, \quad j \in J,$$

with $T_{\max} := \max_{0 \leq j \leq j_{\max}}\{T_j\}$. The weight factors $K_j > 1$ preserve local scale maxima.

We apply the two-directional integrated Gaussian wavelet filters that we have derived in Sec. 5.1. Let

$$\mathcal{M}_j(b) = \sqrt{|F(b, j, l_x)|^2 + |F(b, j, l_y)|^2}, \tag{15}$$

if $\sqrt{|F(b, j, l_x)|^2 + |F(b, j, l_y)|^2}$ has a local maximum in b in direction $\arg(F(b, j, l_x), F(b, j, l_y))$ of the gradient, and $\mathcal{M}_j(b) = 0$ otherwise. The definition of $\mathcal{M}_j$ is based on the assumption that every spot generates a local scale maximum on a certain scale. Thus microcalcifications are contained in the set of local maxima.

Exploiting wavelet coefficient decay

Local regularity is a feature that can be used to distinguish between microcalcifications and film artifacts such as film defects and scratches on the film. Such film artifacts produce sharp edges while calcifications have a smoother outline.

The local regularity of a signal can be measured from the decay order of the magnitude of the corresponding wavelet coefficient over scales, compare S. Mallat and S. Zhong.[30] Figure 13 (top) shows the continuous wavelet analysis of a signal (bottom) which consists of intersections of four microcalcifications and four film artifacts. The corresponding maxima-lines are displayed in the center. Solid lines indicate microcalcifications, dashed lines indicate film artifacts. The gradient of the magnitude over scales in Fig. 13 (right) is proportional to the decay order of the wavelet coefficient.

Both, microcalcifications and film artifacts, generate long maxima lines with large magnitudes. However, they differ in the rate of decay. Maxima lines of film

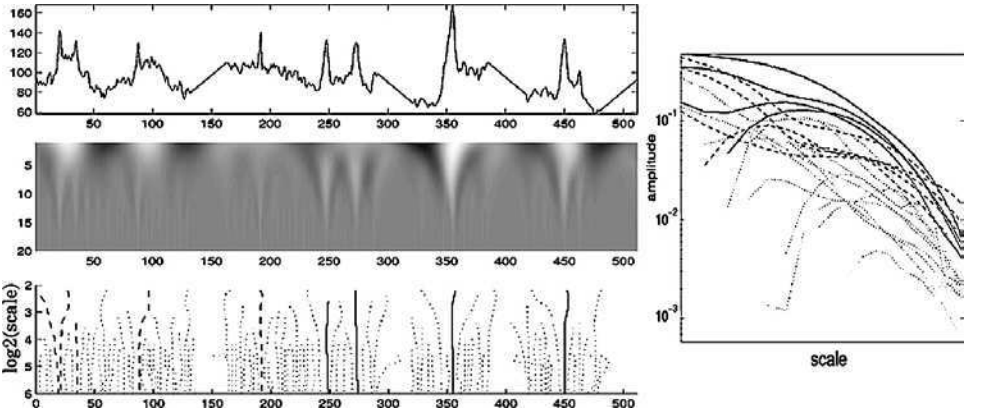Fig. 13.   Wavelet analysis of intersections. The signal f (top) is a composition of intersections of three microcalcifications and three film artifacts. From the continuous wavelet coefficient $\mathcal{WT}_\psi f$ (center) we extract maxima lines (bottom). Lines corresponding to film artifacts are dashed, lines corresponding to microcalcifications solid. Selected amplitudes of maxima lines are plotted on logarithmic axes (right).
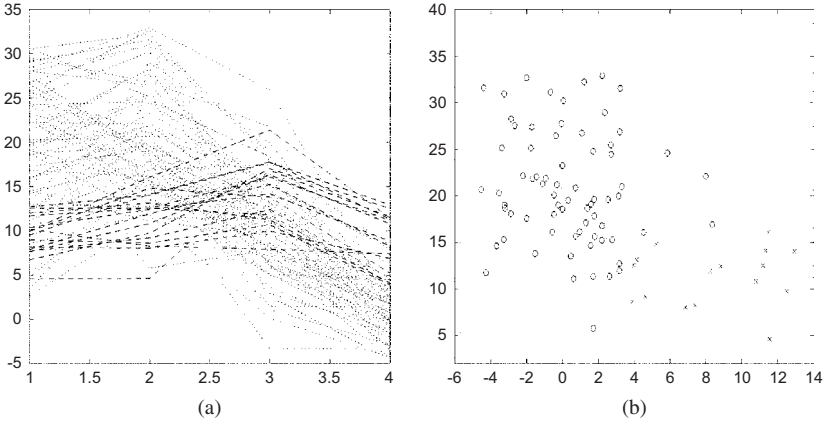
Fig. 14.   Decay of wavelet coefficients. (a) Amplitude of integrated wavelet coefficients of microcalcifications (dashed) versus film artifacts (dotted) over four scales. (b) Values from wavelet scales 2 and 4 plotted on the horizontal and vertical axes. Microcalcifications are marked by circles, film artifacts by stars.

artifacts have lower decay order than those of microcalcifications. This corresponds to the known lower regularity of film artifacts.

We translate this observation to discrete wavelet scales. Figure 14(a) shows integrated wavelet coefficients over four scales. Values corresponding to microcalcifications are dotted, film artifacts are dashed.

We see that we can still distinguish microcalcifications from film artifacts by the gradient of amplitudes. Figure 14(b) shows that by comparing coefficients of only two (properly chosen) scales, microcalcifications can be distinguished from film artifacts.

T. Netsch[33] computed another two-dimensional feature from scale-space. He uses the extremum over scales together with the value at the extremum. With this feature he detects microcalcifications properly. However, he points out that he cannot distinguish between film artifacts and microcalcifications.

To exploit these observations we define the enhancement weight M_j.

$$M_j(b) := \begin{cases} \mathcal{M}_j(b), & 2 < \mathcal{S}(WT_\psi(f)(b,.)) < 3, \\ 0 & \text{otherwise.} \end{cases} \tag{16}$$

Here $\mathcal{S}\big((j_1, \ldots, j_{\max})\big) := \big(\sum_{k=1}^{j_{\max}} k j_k\big)/\big(\sum_{k=1}^{j_{\max}} j_k\big)$ denotes the focal point over scales. The values 2 and 3 depend on the numbering of scales. They are chosen to comprise the expected sizes of microcalcifications.

The threshold $M_j(b) > T_j$ is an explicitly constructed classifier. By composing the features with a logic AND-operator we implicitly construct a partition of the feature space into rectangular cubes. One can do better by training of a classifier on these features. Here we only demonstrate the capability of our features to separate classes.

5.3. *Effect of discretization*

We introduced integrated wavelets because of their flexibility in choosing scales adaptively. Figures 15 and 16 show a comparison of enhancement on dyadic scales and on an adapted range of scales. In the first case we compute five filters on a dyadic range. The scales are $0.4, 0.8, 1.6$ and 3.2 millimeters. The parameters a_j are given by multiplying these values with the pixel resolution of the image. The pixel resolution of the images used is 20 pixel per mm. The actual filters are then given by formulas (12), (13) and (14).
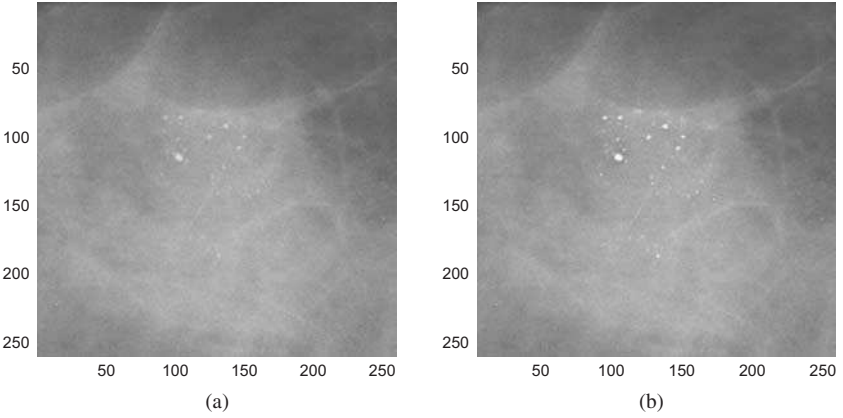

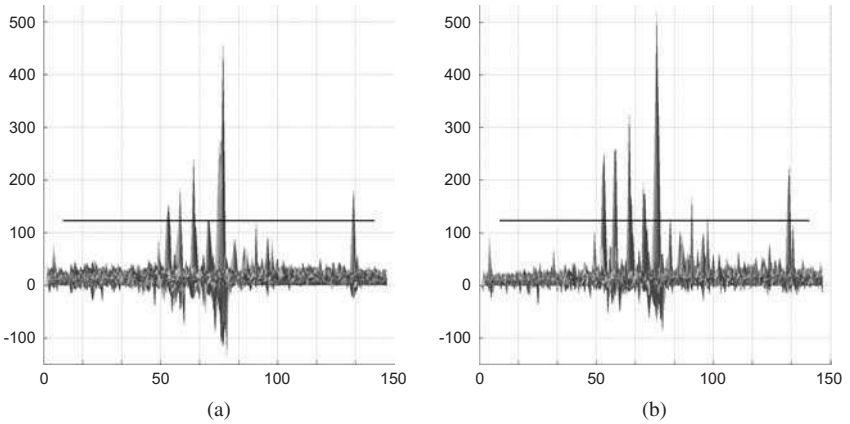Fig. 15.   Microcalcification Enhancement. Original (a) and result of enhancement (b).

Fig. 16.   Adapted discretization of scales. Magnitude images of enhancement using integrated Gaussian wavelets. (c) Dyadic discretization. (d) We added two scales in the range relevant for microcalcifications. Contrast improvement is (c) 5.62 and (d) 6.21.

For the adapted discretization we compute seven filters on a more dense range of scales in order to match better with the expected size of calcifications. The adapted scales are $0.55, 0.8, 1.25, 1.6, 2.1$ and 3.2 millimeter.

The magnitude images 16(a) and (b) are normalized to have the same mean and variance as the original image. There is an arbitrary threshold drawn into both images. We observe that, when using the integrated wavelet transform, more microcalcifications appear above the threshold. This is exactly what you can expect from an adapted algorithm. Contrast improvement in this example is 5.62 for dyadic scales. Adding more scales in the dyadic range did not improve this result: 0.2 millimeter is almost the pixel spacing of the Nijmengen images, thus the corresponding filter is already sampled at the Nyquist rate, while 6.2 millimeter is by far to large to respond to microcalcifications.

For the adapted discretization, contrast improvement is 6.21. On average, adapting the discretization of scales increases the contrast value about 5 to 10%. This improvement relies only on the adapted discretization. It would have been possible neither using classical dyadic wavelet frames nor using wavelet frames with interpolating scales with seven filters only.

5.4. *Robustness*

As discussed before, a major challenge in digital mammography is that algorithms are expected to work regardless of the imaging modality. In design of our algorithm we have carefully chosen robust features, such as scale-adapted integrated wavelets, that do not depend on pixel spacing or assumptions on image characteristics. We solved the problem of enhancement for digitized film, because there film-artifacts and noise are worst.

Our algorithm for enhancement of microcalcifications also has been evaluated on direct digital images at the Charité university hospital in Berlin by F. Diekmann *et al.*[11] Direct digital mammograms show less noise and contain practically no artifacts. Due to other phenomena, mainly due to reduced pixel resolution, microcalcifications appear less sharp than in film digitized at high resolution. Figure 17 shows examples of microcalcification clusters in direct digital mammograms. We have applied the same algorithm as described above to these direct digital images. Scales were computed for the different image resolution but with the same scheme as before. Results of the enhancement algorithm are shown in Fig. 18. One can see that visibility of the calcification cluster in each image segment is improved significantly. Overall the system performance on direct digital images shows the same performance as on film. Due to less noise and practically no artifacts from the imaging process the enhancement algorithm is even more specific to microcalcifications. Thus one can actually increase the gain factor which results in a higher overall contrast gain for direct digital images.

In this section we have presented in detail a method for enhancement of microcalcifications in digital mammograms. It improves several components of known
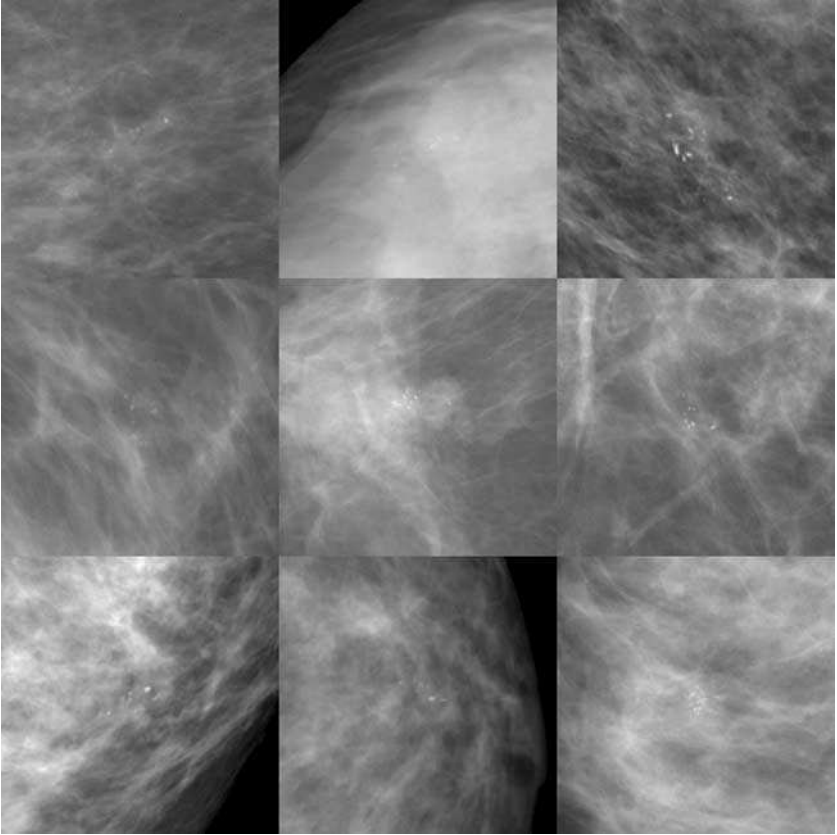
Fig. 17.   Examples for microcalcification clusters in direct digital mammograms.

wavelet-based enhancement algorithms. Due to exploiting the flexibility of the newly developed integrated wavelets we obtain contrast improvements of about 5 to 10% in comparison to hitherto used wavelet transforms. Integrated wavelets allow us to pick discrete scales where the image features are located, thus matching the filterbank with the size of the interesting structure. The exact choice of scales was done manually, following the model of microcalcifications. The algorithm adapts itself automatically to changing image resolution by converting scales given in millimeter to pixel resolution. We can compute the filterbank for decomposition and reconstruction directly, without applying slowly converging iterative schemes. Thus for arbitrary choice of scales, the computational load always depends only on the number of scales. By choosing an adapted wavelet we showed that enhancement becomes more specific to microcalcifications. The problem of film artifacts was targeted by using knowledge of the decay order of wavelet coefficient.
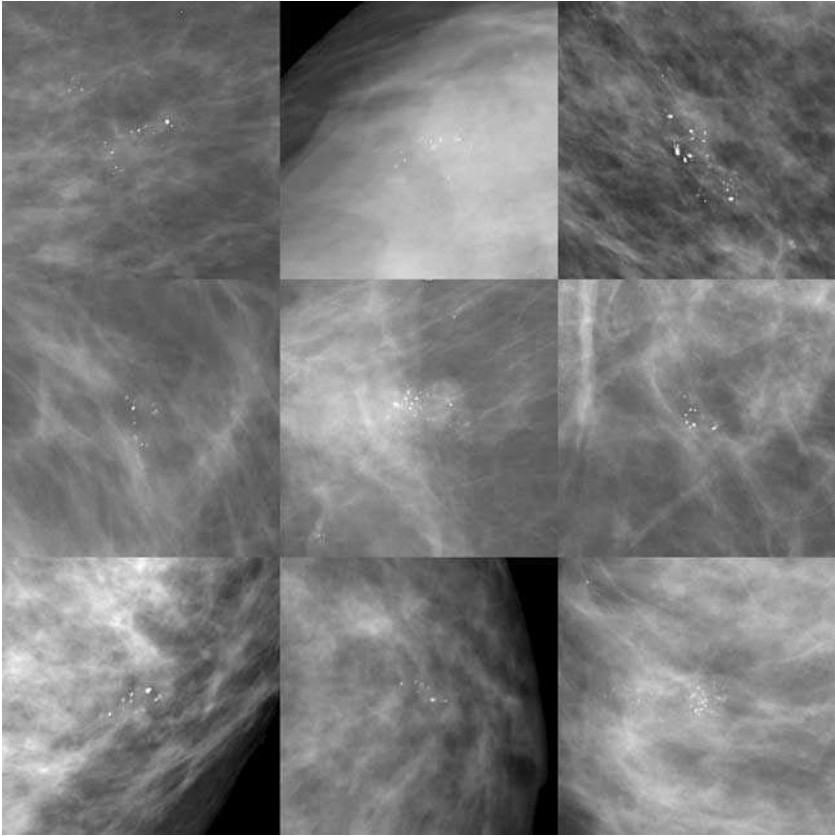
Fig. 18.   Enhancement of microcalcification clusters in direct digital mammograms.

6. Further Applications

6.1. *Exploiting directional information for detection of microcalcifications*

The ideas above can also be applied to the detection of microcalcifications. A digital mammogram (at $5\,\mathrm{LP/mm}$) contains about $2048\cdot2048 = 4.194.304$ pixels. Less than about 20 to 400 pixels correspond to calcifications — if there are any calcifications present at all. On the other hand, training of a classifier works best when all classes are of about the same size. This motivates a first pre-classification step. You can apply a threshold to the enhanced image after removal of the lowpass component. The threshold is to be chosen in a way that all calcifications appear above the threshold.

The enhancement algorithm is based on the assumption that microcalcifications are almost circular and on the use of an appropriate isotropic wavelet to
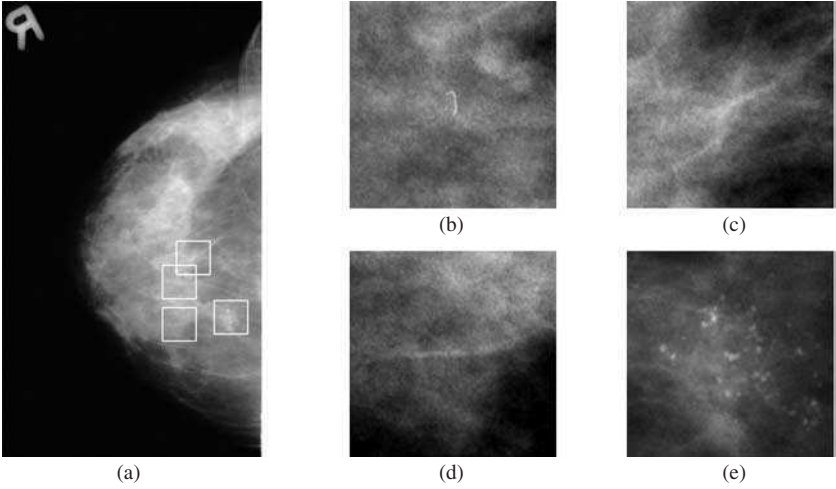
Fig. 19. Typical structures in a mammogram. (a) image c05c, Nijmegen mammography database with (b) film artifact, (c) tissue crossing, (d) blood vessel and (e) microcalcification.

detect spot-like structures. Overall this approach yields good results. An in depth analysis shows that there are difficulties to distinguish microcalcifications from similar structures such as film-artifacts, crossings of structured tissue or small line-like structures that locally fall apart into a series of small spots. On a small scale they all are spot-like, but they exhibit very different directional behavior on other scales. We apply the scale-angle representation of the continuous wavelet transform to measure this distinction. It was first proposed by J.-P. Antoine *et al.*[5,6] For fixed $b \in \mathbb{R}^2$ the scale-angle representation is defined as section $b \times \mathbb{R}_+^* \times SO(2)$ through the wavelet coefficient $WT_\psi f$. From this representation one can deduce dominant directed structures and their corresponding scale.

Figure 19 shows typical structures in a mammogram. Figure 20 shows corresponding scale-angle representations. Figure 20(a) corresponds to a film artifact. We observe the directed structure on small scales. Figure 20(b) shows a crossing of tissue forming a spot-like center. We observe four directed structures corresponding to the tissue. Figure 20(c) corresponds to a directed blood vessel. Figure 20(d) is calculated at the center of a microcalcification. As predicted from our model, the scale-angle representation exhibits a strong isotropic structure.

A major disadvantage of using the scale-angle representation directly for this type of problem is the amount of data to be computed from the continuous wavelet coefficient. Further, by sampling scales and directions we might miss important peaks. Thus instead of sampling the CWT, we apply integrated wavelets for the computation of directional information. First we introduce a measure

$$\mathcal{R}(b, \rho) := \int_{[a_j, a_{j+1}]} \mathcal{W}\mathcal{T}_\psi f(b, a, \rho) \, \frac{da}{a}$$
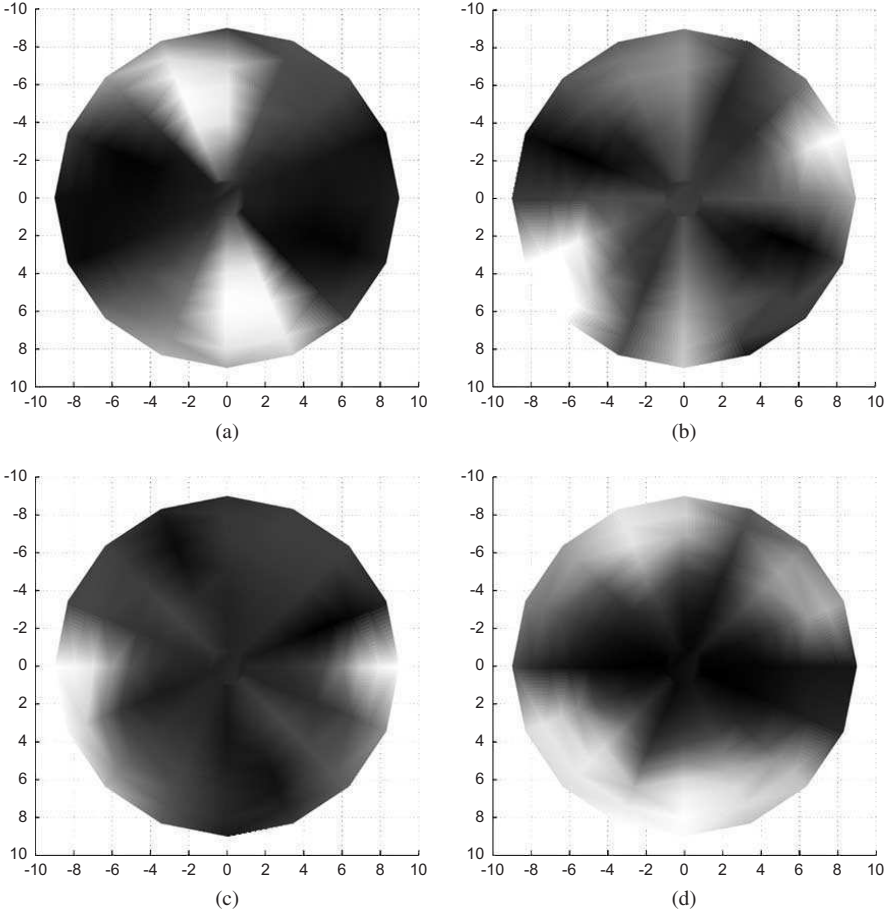
Fig. 20. Scale-angle representations. The representations correspond to the center of the (a) film artifact, (b) crossing of tissue, (c) blood vessel and (d) microcalcification from Fig. 19. Brightness corresponds to the absolute value of the wavelet coefficient. The inner circle corresponds to the smallest scale.

for directional information at a certain range of scales $[a_j, a_{j+1}]$. For a given point b it measures the magnitude of directional information in direction ρ on the range of scales $[a_j, a_{j+1}] \subset \mathbb{R}_+^*$. For ρ we choose 16 values, i.e. we partition $SO(2)$ by $K_l = [l/16, (l+1)/16]$. With this setting we can compute $\mathcal{R}$ from the integrated wavelet coefficient immediately:

$$\mathcal{R}(b, l) = \mathcal{WT}_\psi^I f(b, j, l).$$

A single integrated wavelet accumulates all scale information in the chosen range of scales and the given sector of directions. Further, from the tight frame property we know that independently of our choice of discrete grid, all directional information is contained in the integrated wavelet coefficients. With a classical sampling approach

we would have had to compute values of the CWT at several scales and then compute the sum over all scales. This would imply the risk to miss information with to sparse a grid. Thus by using integrated wavelets we save a large amount of computational load while at the same time obtaining more accurate results.

Microcalcifications typically are isotropic, while line-like structures exhibit strong periodicity. Irregular structures generate small autocorrelation values. Thus we compute the autocorrelation function $R_{\mathcal{R}}$ for the directional measure $\mathcal{R}$ with respect to the direction ρ:

$$R_{\mathcal{R}}(b, \rho) := \frac{1}{\|\mathcal{R}(b, .)\|_2} \int_{SO(2)} \mathcal{R}(b, r)\mathcal{R}(b, r + \rho) \, dr.$$

Figure 21 shows the autocorrelation function $R_{\mathcal{R}}$ for the points b selected in Fig. 20(a). Large values of $R_{\mathcal{R}}$ correspond to periodicity with a period corresponding to the angle they appear. In Figs. 21(a) and 21(c) we observe strong periodicity at angles 0 and π, corresponding to the line like structure of the film artifact and the blood vessel. Figure 21(b) shows several small extrema corresponding to the tissue crossing. Figure 21(d) is almost constant. This corresponds to the approximately isotropic structure of microcalcifications.

As we observe from the example, the directional information provided by the wavelet transform can help to distinguish similar structures by analyzing their anisotropy across several scales and directions. Especially the separation of
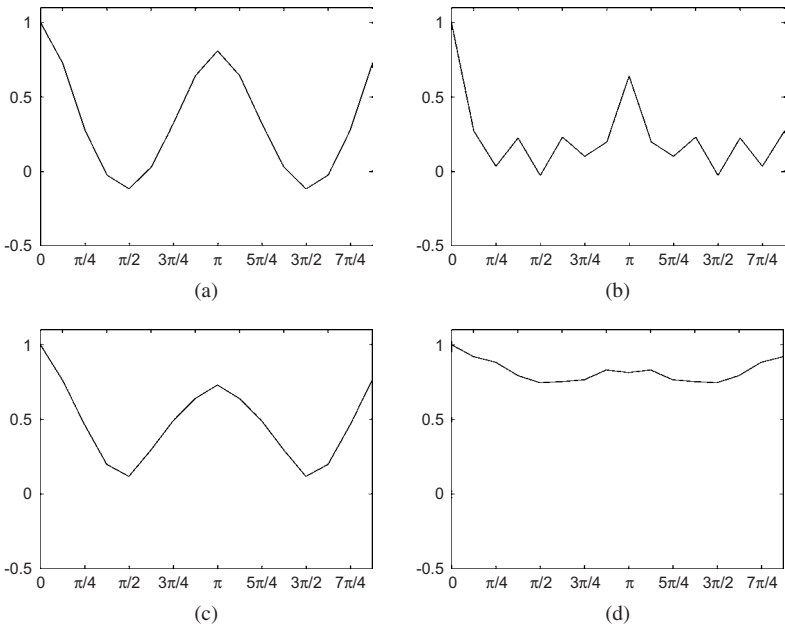

Fig. 21.   Autocorrelation function for directional measure. Autocorrelation w.r.t. ρ of the directional measure $\mathcal{R}$ corresponding to the scale-angle representations from Figs. 20(a)–(d).

microcalcifications from tissue crossings that generate spot-like structures is not possible with the typical approach with just two filters, that estimates the local gradient only.

6.2. *Denoising of mammograms*

Let us close with an other application of the general method. The problems of enhancement and denoising share the same concept. in both cases one intends to discriminate between signal and background. For enhancement "signal" is just microcalcifications. For denoising "background" is the noise. Based on such a discrimination the aim is to modify the image without introducing artifacts.

Artifacts are a major problem of classical denoising algorithms that are based on wavelet thresholding. Coefficients in wavelet domain correspond to wavelet functions in the signal domain. Thus modification of wavelet coefficients implies adding copies of the wavelet function to the signal. In standard algorithms these copies are visible as artifacts. To eliminate this effect, we propose the application of a different wavelet reconstruction scheme, based on a variation of integrated wavelets. The method generates less correlated results than DWT or dyadic wavelet transform based denoising schemes. The quality is measured by the autocorrelation of the reconstructed signal.

We find the results better suited for further processing by algorithms that rely on statistical features. An example from mammographic image analysis is the problem of mass detection in digital mammograms. One typically starts with a directional feature extraction method based on local image statistics. Removing the noise can improve detection of lines that indicate the presence of a spiculated lesion. On the other hand the preprocessing may not introduce line like artifacts and shall not change the statistical distribution of directional information. Thus our emphasis goes beyond smoothing visible artifacts or optimizing the overall SNR. Our aim is to minimize the correlation that is introduced by the denoising procedure.

The original wavelet thresholding method by D. Donoho and I. Johnstone[12] is based on computation in a wavelet basis. While smooth wavelets produce results that may look good to the eye, a subsequent operator can be significantly disturbed by these artifacts. Improved algorithms have been developed that try to reduce this effect by averaging over a redundant frame decomposition. This technique is often called cycle spinning. Most recent research is devoted to optimization of the threshold operator, cf. F. Shi and I. W. Selesnick[37] or P. L. Dragotti and M. Vetterli.[13] An alternative approach for application to piecewise smooth signals with sharp discontinuities comes from S. Durand and J. Froment.[14] They apply a total variation minimization approach, but present no rigorous quality analysis of the algorithm.

Morlet reconstruction

Our approach is based on the Morlet reconstruction. The continuous Morlet reconstruction is well known. Let us first derive the Morlet integrated wavelet transform,

extending the original ideas from M. A. Muschietti and B. Torrésani[32] into a general framework. All results also hold in the setting of wavelet transforms over semidirect product groups $G = \mathbb{R}^m \rtimes H$, where H denotes a general dilation group H. We apply the notation introduced in Sec. 4.

The image of the CWT is a reproducing kernel Hilbert space. The reproducing kernel allows to choose different functions for decomposition and reconstruction. Let χ be a reconstruction wavelet, i.e.

$$0 \neq \int_H \overline{\hat{\psi}(\omega_0 h)} \hat{\chi}(\omega_0 h) \, d\mu_H(h) := c_{\psi,\chi},$$

then there is a generalized reconstruction formula

$$f(x) = \frac{1}{c_{\psi,\chi}} \int_{\mathbb{R}^2} \int_H \mathcal{WT}_\psi f(g) \mathcal{U}_g \chi(x) \, d\mu(g) \tag{17}$$

in L^2-norm.

There is an extremal case. Formally, one can replace the reconstruction wavelet in (17) by the Dirac distribution which leads to the *Morlet reconstruction*

$$f(x) = \frac{1}{k_\psi} \int_H \mathcal{WT}_\psi f\big((x,h)\big) \big|\det(h)\big|^{-1/2} \, d\mu_H(h). \tag{18}$$

For this formula to hold, there is again an admissibility condition. We call $\psi \in L^2(\mathbb{R}^2)$ *Morlet-admissible*, iff

$$\int_H |\hat{\psi}(\omega_0 h)| d\mu_H(h) := K_\psi < \infty, \quad \text{and}$$
$$0 < \int_H \overline{\hat{\psi}(\omega_0 h)} d\mu_H(h) := k_\psi \tag{19}$$

for $\omega_0 \in \widehat{\mathbb{R}^2} \setminus \{0\}$. For Morlet admissible ψ the reconstruction (18) is valid for all $f \in L^2(\mathbb{R}^2)$ in the weak topology. If $\psi \in L^2(\mathbb{R}^2) \cap L^1(\mathbb{R}^2)$ the reconstruction holds in L^2-norm for all $f \in L^2(\mathbb{R}^2)$ with $\hat{f} \in L^1(\widehat{\mathbb{R}^2})$.

The Morlet reconstruction needs no integration over $\mathbb{R}^2$. Thus the Morlet reconstruction is by orders faster than the classical reconstruction. Further, the reconstruction induces no smoothing of wavelet coefficients. This is the key property to construct an artifact free denoising algorithm.

Admissibility does not imply Morlet admissibility or vice versa.[20] Nevertheless this is true for most functions: If $\hat{\psi}$ is bounded in a neighborhood of 0 then Morlet admissible implies admissible. In the other direction, if $\hat{\psi}(0) = 0$ and $\hat{\psi}$ is pointwise Lipschitz α in 0 for some $\alpha > 0$, then admissible implies Morlet admissible.

Morlet integrated wavelet transform

From the Morlet reconstruction we can derive the Morlet integrated wavelet transform in a way similar to the definition of integrated wavelets from Sec. 4. The Morlet reconstruction forces us into a l^1-setting. This was presented as the linear scheme by M. A. Muschietti and B. Torrésani.[32]

Let $(a_j)_{j\in\mathbb{Z}}$ be a strict monotonic decreasing sequence in $\mathbb{R}_+^*$ with $\lim_{j\to-\infty} a_j = \infty$ and $\lim_{j\to\infty} a_j = 0$. Further let $(K_l)_{l\in L}$ be a partition of $SO(2)$ with finite index set L. For Morlet admissible ψ we define the *Morlet integrated wavelet* $\Psi^{j,l} \in L^2(\mathbb{R}^2)$ in Fourier domain by

$$\widehat{\Psi^{j,l}}(\omega) := \frac{1}{k_\psi} \int_{a_{j+1}}^{a_j} \int_{K_l} \widehat{\mathcal{D}_a R_\rho \psi}(\omega) \frac{d\rho\, da}{a}. \tag{20}$$

Morlet admissibility implies a partition of unity in Fourier domain given by

$$\sum_{j\in\mathbb{Z}} \sum_{l\in L} \widehat{\Psi^{j,l}}(\omega) = 1 \quad \text{for a.e. } \omega \in \widehat{\mathbb{R}^2}. \tag{21}$$

For a given family $(\Psi^{j,l})_{j,l}$ of Morlet integrated wavelets we call

$$MWT_\psi : L^2(\mathbb{R}^2) \to L^2(l^1(\mathbb{R}^2 \times \mathbb{Z} \times L))$$
$$f \mapsto \langle f, T_b \Psi^{j,l}\rangle, \quad b \in \mathbb{R}^2, \; (j,l) \in \mathbb{Z} \times L.$$

the *Morlet integrated wavelet transformation (MWT)*.

Condition (19) implies the inequality

$$1 \le \sum_{j\in\mathbb{Z}} |\widehat{\Psi^j}(\omega)| \le \frac{K_\psi}{k_\psi}, \qquad \omega \in \widehat{\mathbb{R}^2}.$$

If we multiply the inequalities with $|\hat{f}(\omega)|$ and take the L^2-norm, then application of Parseval's formula yields the norm inequality

$$1 \cdot \|f\|^2 \le \left\| \sum_{j\in\mathbb{Z}} \sum_{l\in L} |\langle f, T_x \Psi^{j,l}\rangle| \right\|^2 \le \frac{K_\psi}{k_\psi} \cdot \|f\|^2. \tag{22}$$

This inequality is similar to the one obtained in the case of frames,[21] but here we have one component in l^1-norm.

If ϕ, defined by $\hat{\phi}(\omega) := \int_1^\infty \widehat{D_u \psi}(\omega) \frac{du}{u}$, $\omega \in \widehat{\mathbb{R}^2}$, is in $L^1(\mathbb{R}^2) \cap L^2(\mathbb{R}^2)$, then we call ϕ a *scaling function*. We then obtain from (21) the reconstruction

$$f(x) = \langle f, T_x D_{j_0} \phi\rangle + \sum_{j=j_0}^{\infty} \langle f, T_x \Psi^j\rangle, \quad x \in \mathbb{R}^2, \tag{23}$$

in the weak topology on $L^2(\mathbb{R}^m)$. Interpreting the sum as limit, the identity also holds in L^2-norm.[20,32]

For example the Gaussian wavelet $\psi(x) := (2\pi)^{-\frac{1}{2}}(1 - x^2) \exp(-\|x\|^2)$, $x \in \mathbb{R}^2$, has the associated Morlet integrated wavelet

$$\Psi^j(x) = (2\pi)^{-\frac{1}{2}} \big(a_j^{-\frac{1}{2}} \exp(-\|x\|^2/2a_j^2) - a_{j+1}^{-\frac{1}{2}} \exp(-\|x\|^2/2a_{j+1}^2) \big).$$

The scaling function is $\phi(x) := (2\pi)^{-1/2} e^{-\|x\|^2/2}$.

An implementation for discrete signals can be derived from formula (23) by sampling the Morlet integrated wavelet $\Psi^{j,l}$ in the Fourier domain. Let $m \in \mathbb{N}$, and let J be a finite set of scales. Take $f \in l^2(\mathbb{Z}_m^2)$ and a Morlet admissible wavelet

$\psi \in L^2(\mathbb{R}^2) \cap L^1(\mathbb{R}^2)$. Let F denote the discrete Fourier transform. Then the decomposition is given by

$$MWT_\psi^d : l^2(\mathbb{Z}_m^2) \to l^2(l^1(\mathbb{Z}_m^2 \times J)),$$

$$f \mapsto F^{-1}(Ff \cdot \widehat{\overline{\Psi^j}})(b).$$

There is a perfect reconstruction formula

$$f(x) = \sum_{j \in J} MWT_\psi^d f(x, j), \quad x \in \mathbb{Z}_m^2.$$

The Morlet integrated wavelet transform is flexible with respect to choice of discrete scales.[22] Further it is shift-invariant. For operators in wavelet domain we have a real-time reconstruction of the signal.

Application to image denoising

We adapt the well known results by D. Donoho and I. Johnstone[12] on nonparametric regression in wavelet bases to the Morlet integrated wavelet transform.

The noisy data is a realization of the stochastic process $X := f + N$ with signal f and additive noise N. Let $S := Z_m \times J$ and $f \in l^2(l^1(S))$. Denote $s = (x, j) \in S$. Further, let $w_s := w_{x,j} := T_x \Psi^j$. Let the noise $N := (N_s)_{s \in S}$ be stationary Gaussian with variance σ^2. Then the coefficients $\langle N, w_s \rangle$ are Gaussian random variables with variance $\sigma^2 \|w_s\|_2^2$. We investigate the error of the nonlinear coefficient selection estimator

$$\tilde{X} := \sum_{s \in S} \langle X, T_x \Psi^j \rangle \Theta(x, j), \quad s = (x, j),$$

with ideal coefficient selection $\Theta : S \mapsto \{0, 1\}$. Let $M := \{s \in S ; |\langle f, \omega_{x,j} \rangle| \geq \sigma \|\omega_{x,j}\|_2\}$. Then the pointwise error is

$$\langle f, w_s \rangle - \langle X, w_s \rangle \Theta(s) = \begin{cases} -\langle N, \omega_s \rangle; & s \in M, \\ \langle f, \omega_s \rangle; & s \notin M. \end{cases}$$

The optimal L^2-error ε of the estimator $\tilde{X}$ is

$$\varepsilon := E\left(\left\|f - \tilde{X}\right\|^2\right)$$

$$\overset{(22)}{\leq} E\left(\left\|\sum_{j \in J} |\langle f, w_{x,j} \rangle - \langle X, w_{x,j} \rangle \Theta(x, j)|^1\right\|_2^2\right)$$

$$= \sum_{x \in \mathbb{Z}_m} E\left(\left(\sum_{j \in J} \min\{|\langle f, w_{x,j} \rangle|, |\langle N, w_{x,j} \rangle|\}\right)^2\right)$$

$$\leq \sum_{x \in \mathbb{Z}_m} \left(\left(\sum_{j \in M^C} |\langle f, w_{x,j} \rangle|\right)^2 + \sum_{j \in M} \sigma^2 \|w_{x,j}\|^2\right).$$

Thus the optimal error is comparable to the error in the case of orthonormal bases.

The same reasoning as above applies to attenuation estimators. For Morlet-reconstruction we set

$$\tilde{Y} := \sum_{s \in S} \Theta_a\big(\langle X, w_s\rangle\big)$$

with hard-threshold or soft-threshold operator Θ_a. Then for threshold $T_s := \sigma\|\omega_s\|\sqrt{2\ln|S|}$ and $|S| > 3$ the average approximation error satisfies

$$\varepsilon = E\big(\|f - \tilde{Y}\|^2\big) \le c_J\big(2\ln|S| + 1\big)$$
$$\cdot \left(\sigma^2 + \sum_{x \in \mathbb{Z}_m}\left(\left(\sum_{j \in M^C}|\langle f, w_{x,j}\rangle|\right)^2 + \sum_{j \in M}\sigma^2\|w_{x,j}\|^2\right)\right).$$

Results on non-stationary denoising hold in the same way.

We have computed the signal to noise ratio (SNR) for several test-images. The results confirm the above theoretical estimation: The SNR improvement between the DWT and MWT is just 5 to 10% apart. Sometimes the MWT is better, more often the DWT has larger SNR.

To measure the quality of the denoising algorithm with respect to artifacts we apply the algorithm to pure noise. We use the soft threshold operator because it is a continuous operator in wavelet domain and thus does not introduce discontinuities. The noise is stationary Gaussian. We apply the threshold $3\sigma\|\Psi^j\|_2$. The low-pass filter is set to zero because it would dominate the autocorrelation result with its low frequency waves.

In Fig. 23 we compare results for the DWT, the dyadic wavelet transform and the Morlet integrated transform. Artifacts in the DWT and dyadic WT are obvious. The DWT, as a basis transform, shows a well centered autocorrelation. Main flaws come from the tensor-product construction. The dyadic WT exhibits far ranging

(a) Mammogram   (b) Dyadic WT vs. Morlet WT
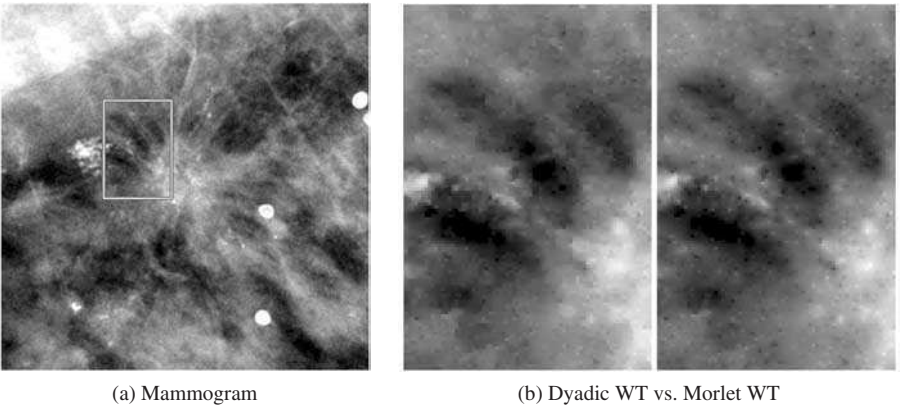
Fig. 22.   Denoising applied to digital mammogram. The dyadic WT introduces artificial directional lines. With the Morlet WT there just remain spikes from large realizations of the noise.

(a) Gaussian random noise

(b) DWT

(c)

(d) Dyadic WT

(e)
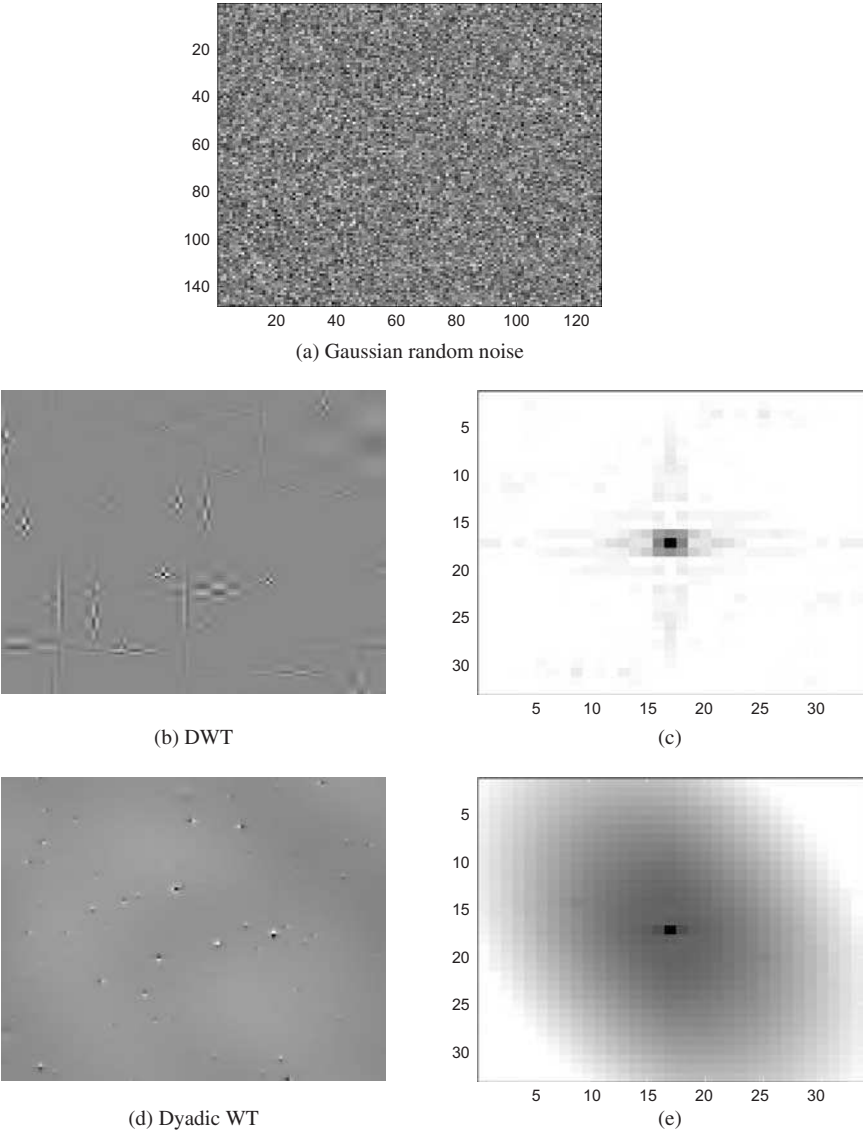
Fig. 23.   Example of denoising using different wavelet transforms. (a) Pure Gaussian noise. Application of the DWT exhibits correlation along axes (b) which is expressed in checkerboard artifacts (c). The Dyadic WT is correlated over scales (d), the wavelet clearly visible at varying scales (e). The Morlet-reconstruction is uncorrelated (f). No artificial structure is added to the image (g).

correlation that comes from the averaging over coefficients. In contrast the MWT is almost uncorrelated by construction.

In Fig. 22 we compare results for a mammogram with a spiculated mass. While the dyadic WT introduces artificial directional lines, with the Morlet-WT there just remain spikes from large realizations of the noise. The statistical feature that we

(f) Morlet WT   (g)
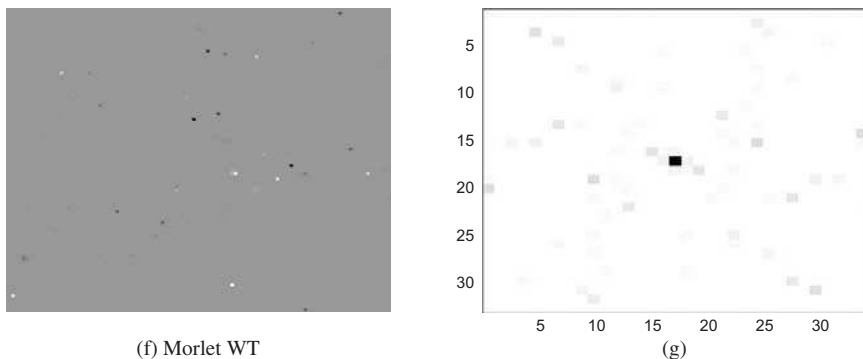
Fig. 23.   (*Continued*)

use to determine the presence of a spiculated mass is not disturbed and performs well on the Morlet denoised image.

7. Conclusion

We have presented a method for enhancement of microcalcifications in digital mammograms. It is based on application of the integrated wavelet transform. This transform provides a reconstruction formula for arbitrary choice of discrete scale and orientation. In exploiting this flexibility we obtain contrast improvements of about 5 to 10% in comparison to hitherto used wavelet transforms. The algorithms have been positively evaluated on direct digital images at the Charité University Hospital in Berlin, Germany.[11] The enhancement algorithms have been implemented in a digital mammography review workstation developed by Image diagnost, Munich, Germany.

References

1. A. Aldroubi and M. Unser, *Wavelets in Medicine and Biology* (CRC Press, New York, 1996).

2. S. T. Ali, J.-P. Antoine and J.-P. Gazeau, *Coherent States, Wavelets and Their Generalizations* (Springer, New York, 2000).

3. J.-P. Antoine, The continuous wavelet transform in image processing, preprint, 1998.

4. J.-P. Antoine, P. Carette, R. Murenzi and B. Piette, Image analysis with two-dimensional continuous wavelet transform, *Signal Processing* **31** (1993) 241–272.

5. J.-P. Antoine and R. Murenzi, Two-dimensional directional wavelets and the scale-angle representation, *Signal Processing* **52** (1996) 259–281.

6. J.-P. Antoine, R. Murenzi and P. Vandergheynst, Directional wavelets revisited: Cauchy wavelets and symmetry detecion in patterns, *Applied and Computational Harmonic Analysis* **6** (1999) 314–345.

7. U. Bick, Computerassistierte Diagnose in der Screeningmammographie, *Radiologie* **36** (1996) 72–80.

8. H. D. Cheng, X. Cai, X. Chen, L. Hu and X. Lou, Computer-aided detection and classification of microcalcifications in mammograms: A survey, *Pattern Recognition* **36** (2003) 2967–2991.

9. L. P. Clarke, M. Kallergi, W. Qian, H. D. Li, R. A. Clark and M. L. Silbiger, Tree-structured non-linear filter and wavelet transform for microcalcification segementation in digital mammography, *Cancer Letters* **77** (1994) 173–181.

10. F. Diekmann, S. Diekmann, M. Bollow, K. G. Hermann, K. Richter, P. Heinlein, W. Schneider and B. Hamm, Evaluation of a wavelet-based computer-assisted detection system for identifying microcalcifications in digital full-field mammography, *Acta Radiologica* (2004).

11. F. Diekmann, P. Heinlein, J. Drexl, S. Grebe, A. Gössler, W. Schneider and B. Hamm, *Visualization of microcacifications by Full Field Digital Mammography using a wavelet algorithm*, pages 526–530. In: Computer Assisted Radiology and Surgery CARS 2001 (G. Elsevier Science B.V., 2001).

12. D. Donoho and I. Johnstone, Ideal spatial adaptation by wavelet shrinkage, *Biometrika* **81** (1994) 425–455.

13. P. L. Dragotti and M. Vetterli, Wavelet transform footprints: Catching singularities for compression and denoising, in *Int. Conf. Image Processing* (September 2000).

14. S. Durand and J. Froment, Artifact free signal denoising with wavelets, in *Proc. ICASSP*, Vol. 6 (2001).

15. M. Duval-Destin, M. A. Muschietti and B. Torrésani, Continuous wavelet decompositions, multiresolution, and contrast analysis, *SIAM J. Math. Anal.* **24**(3) (1993) 739–755.

16. K. Ewen, *Moderne Bildgebung: Physik, Gerätetechnik, Bildbearbeitung und -kommunikation, Strahlenschutz, Qualitätskontrolle* (Thieme Verlag, Stuttgart, 1998).

17. U. Fischer, K.-P. Hermann and R. Vosshenrich, Bildgebende Mammadiagnostik, *Radiologie up2date*, 1, 2001.

18. H. Führ, *Zur Konstruktion von Wavelettransformationen in höheren Dimensionen.* dissertation, Munich University of Technology, Munich (1997).

19. M. Le Gal, J.-C. Durand, M. Laurent and D. Pellier, Conduite a tenir devant une mammographie revelatrice de microcalcifications groupées sans tumeur palpable, *La Nouvelle Presse Medicale* **5** (1973) 1621–1627.

20. P. Heinlein, *Wavelet-Methoden zur Analyse mammographischer Bilddaten.* dissertation, Munich University of Technology, Munich (2001).

21. P. Heinlein, Discretizing continuous wavelet transforms using integrated wavelets, *Applied and Computational Harmonic Analysis* **14**(3) (2003) 238–256.

22. P. Heinlein, J. Drexl and W. Schneider, Integrated wavelets for enhancement of microcalcifications in digital mammography, *IEEE Transactions on Medical Imaging* **22**(3) (2003) 402–413.

23. S. Heywang-Köbrunner *et al.*, *Bildgebende Mammadiagnostik: Untersuchungstechnik, Befundmuster und Differentialdiagnostik in Mammographie, Sonographie und Kernspintomographie* (Thieme Verlag, Stuttgart, 1996).

24. R. P. Highnam and M. Brady, *Mammographic Image Analysis* (Kluwer Academic Publishers, Doodrecht, 1999).

25. J. K. Kim, M. Park, K. S. Song and H. W. Park, Adaptive mammographic image enhancement using first derivative and local statistics, *IEEE Transactions on Medical Imaging* **16**(5) (1997) 405–502.

26. A. Laine, J. Fan and W. Yang, Wavelets for contrast enhancement of digital mammography, *IEEE Engineering in Medicine and Biology* **14**(5) (1995) 536–550.

27. A. F. Laine, S. Schuler, J. Fan and W. Huda, Mammographic feature enhancement by multiscale analysis, *IEEE Transactions on Medical Imaging* **13**(4) (1994) 725–740.

28. S. Mallat, *A Wavelet Tour of Signal Processing* (Academic Press, 1998).

29. S. Mallat and W. L. Hwang, Singularity detection and processing with wavelets, *IEEE Transactions on Information Theory* **38** (1992) 617–643.

30. S. Mallat and S. Zhong, Characterization of signals from multiscale edges, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **14**(7) (1992) 710–732.

31. R. Murenzi, *Ondelettes multidimensionelles et applications a l'analyse d'images*, PhD thesis, Université Catholique de Louvain, Louvain-La-Neuve, Belgien (1990).

32. M. A. Muschietti and B. Torrésani, Pyramidal algorithms for Littlewood-Paley decompositions, *SIAM J. Math. Anal.* **26**(4) (1995) 925–943.

33. T. Netsch, *Automated Detection of Clustered Microcalcifications in Digital Mammograms.* PhD thesis, Universität Bremen (1998).

34. W. Qian, L. P. Clarke, B. Zheng, M. Kallergi and R. Clark, Computer assisted diagnosis for digital mammography, *IEEE Engineering in Medicine and Biology* **14**(5) (1995) 561–569.

35. W. Quian, L. P. Clarke, M. Kellergi and R. A. Clark, Tree-structured nonlinear filters in digital mammography, *IEEE Transactions on Medical Imaging* **13**(1) (1994) 25–36.

36. D. Rosen, B. Martin, M. Monheit, G. Wolff and M. Stanton, *A Bayesian Neural Network to Detect Microcalcifications in Digitized Mammograms*, in *Digital Mammography*, eds. K. Doi and M. L. Giger (G. Elsevier Science B.V. 1996), pp. 277–282.

37. F. Shi and I. W. Selesnick, Class of heavy-tailed multivariate non-Gaussian probability models for wavelet coefficients, in *Proc. SPIE: Wavelets X*, Vol. 5207 (August 2003).

38. Robin N. Strickland and Hee Il Hahn, Wavelet transforms for detecting microcalcifications in mammogramms, *IEEE Transactions on Medical Imaging* **15**(2) (1996) 218–229.

39. A. Teolis, *Computational Signal Processing with Wavelets* (Birkhäuser, Boston, 1998).

40. P. Vandergheynst and J.-F. Gobbers, Directional dyadic wavelet transforms: Design and algorithms, *IEEE Transactions on Image Processing* **11** (2002) 363–372.

41. T. C. Wang and N. B. Karayiannis, Detection of microcalcifications in digital mammograms using wavelets, *IEEE Transactions on Medical Imaging* **17**(4) (1998) 498–509.

42. H. Yoshida, W. Zhang, W. Cai, K. Doi, R. M. Nishikawa and M. L. Giger, Optimizing wavelet transform based on supervised learing for detection of microcalcifications in digital mammograms, *ICIP* (1996), pp. 152–155.

This page intentionally left blank

CHAPTER 2

TECHNIQUES IN THE DETECTION
OF MICROCALCIFICATION CLUSTERS
IN DIGITAL MAMMOGRAMS

ISSAM EL NAQA

Department of Radiation Oncology, School of Medicine
Washington University, 4921 Parkview Place
St. Louis, MO 63110, USA

YONGYI YANG

Department of Electrical and Computer Engineering
Illinois Institute of Technology, 3301 S. Dearborn Street
Chicago, IL 60616, USA

In this chapter we provide a review of some state-of-the-art techniques in detection of microcalcification (MC) clusters in digital mammograms. Clustered MCs can be an important early indicator of breast cancer, a leading cause of death in women. Computer aided diagnosis (CAD) systems have been developed in recent years for providing a "second reader" opinion in breast cancer diagnosis. In this review we focus on those most recently published computer techniques for MC detection. Based on the underlying image-processing methods used, we divide these techniques into four categories: (1) basic image enhancement, (2) multiscale decomposition, (3) stochastic modeling, and (4) machine learning. For each category we begin with an overview of the different existing methods, and then provide an in-depth description of one or two representative methods. We also present some numerical results to illustrate these representative methods. In the end we furnish some evaluation results to demonstrate the performance of these methods.

Keywords: Microcalcification; digital mammograms; computer-aided detection; image processing.

1. Introduction

Breast cancer is presently a leading cause of death among women in developed countries. According to the American Cancer Society (ACS), an estimated 215,990 new cases of invasive breast cancer are expected to occur among women in the U.S. during 2004, and approximately 41,110 women are anticipated to die from breast cancer in the same year. Moreover, breast cancer is second only to lung cancer in mortality rates.[1] Research has shown that the key to successfully treating breast cancer is early detection. The five-year survival rate for the disease is nearly five times as great when it is diagnosed in its early stage (96 percent) as opposed to its later stages (20 percent).[1]

Mammography currently provides the most effective strategy for early detection of breast cancer. The sensitivity of mammography is approximately 90%.[2] However, complicated by the presence of a wide variety of structures, accurate reading of mammograms is very difficult. Some breast cancers produce changes in the mammogram that are subtle and difficult to recognize. It has been reported that 10–30% of lesions are misinterpreted during routine screening of mammograms. One study has found that those interpretation errors (e.g. the radiologist sees the cancer, but thinks it is benign) account for 54% of the missed cancers.[3] Furthermore, it is very difficult to mammographically distinguish benign lesions from malignant ones. Between 2 and 10 women are biopsied for every cancer detected.[4,5] This low specificity results in a relatively large interobserver variability that can lead to malignant lesions not being biopsied and benign lesions being biopsied.[6] Unnecessary biopsy is often cited as one of the "risks" of screening mammography. Surgical, needle-core, and fine-needle aspiration biopsies are expensive, invasive, and traumatic for the patient.

There has been a great deal of research in the last 15 years focused on developing computer aided diagnosis (CAD) tools[7] for detection and diagnosis of breast cancer. This intensive research has resulted in several FDA approved commercial systems which aim to play the role of a virtual "second reader" by highlighting suspicious areas for further review by the radiologist.[8,9]

The computerized secondary review is expected to be a valuable tool in breast cancer detection. In fact, studies have shown that 12–19.5% of breast cancers missed would be detected with the use of such CAD systems.[10,11] Such detection is a crucial precursor for accurate diagnosis of cancer (i.e. the differentiation between malignant and benign cases).[12]

In this chapter, we provide an overview of computer techniques reported in recent years in the literature for detection of microcalcifications (MCs) in mammograms. Clustered microcalcifications (MCC) can be an important indicator of breast cancer. They appear in 30–50% of mammographically diagnosed cases.[13] MCs are tiny calcium deposits that appear as small bright spots in mammograms. As an example, Fig. 1 shows a craniocaudal-view mammogram containing a MC cluster, where individual MCs are indicated by circles.

Individual MCs are sometimes difficult to detect because of the surrounding breast tissue, their variation in shape (from granular to rod shapes), orientation, brightness and diameter size (typically, 0.05–1 mm).[13] Round or oval shapes are probably of benign origin, while malignant lesions usually have irregular boundaries. Due to this difficulty, they could be easily overlooked by radiologists.

The rest of the chapter is organized as follows: In Sec. 2, we review some of the basics in mammography and computer-aided diagnosis (CAD). In Sec. 3, we provide a review of various techniques for detection of MC clusters. This is followed in Sec. 4 by a discussion on a set of evaluation results using several representative methods. Finally, we draw conclusions in Sec. 5.
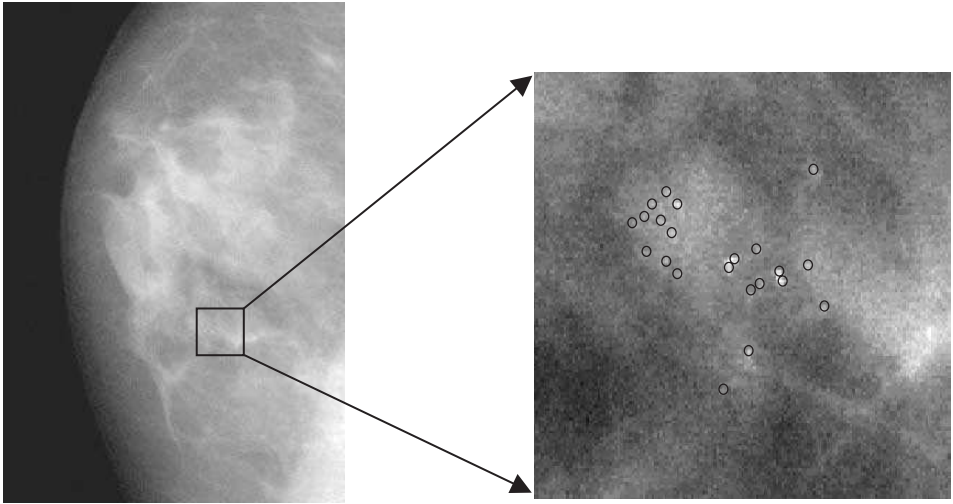
Fig. 1. *Left*: A craniocaudal view mammogram; *Right*: expanded view showing clustered microcalcifications (MCs). MCs are small granule-like deposits of calcium, and appear as bright spots in a mammogram.

2. Basics of Mammography and Computer Aided Diagnosis

2.1. *Mammography*

Mammography is a medical imaging technique that is based on x-ray examination of the breast. The resulting x-ray images are called mammograms (e.g. Fig. 1). Typically, the breast is imaged from two different angles: mediolateral oblique view (MLO) and cranio-caudal view (CC). This procedure has been proven to be an effective tool in reducing mortality due to breast cancer.

As in any other x-ray examination, the quality of the mammograms depends on the attenuation of the different tissues. Generally, a clear separation between normal and cancerous tissues is not perfect. In older women, such a separation is more noticeable due to the diminishing of glandular tissue.

2.2. *Breast abnormalities*

There are two types of lesions that could be observed in mammograms;[13] the first type is localized density of tissue called mass, and the second one appears as bright spots of calcium deposits that are called microcalcifications (MCs). Clusters of MCs (MCC) can be an important early indicator of breast cancer. The likelihood of malignancy is considered to be high if the MCs show a large degree of polymorphism and branching. Most such malignant MCs arise in ductal carcinoma.

MCs could be found within the ducts, alongside ducts, in the lobular acini, in vascular structures, in the interlobular stroma, in fat, or in skin. Depending on

their etiology and location they could be punctuate, branching, linear, spherical, fine, coarse, cylindrical, smooth, jagged, regular in size or shape, or heterogeneous.

The characteristics of MCs have been extensively studied in the literature over the past 30 years, these efforts had led to the establishment of the Breast Imaging Reporting and Data System (BI-RADS) lexicon by the American College of Radiology (ACR) in 1993.[14] Interested readers are referred to the ACR publications[1] for more details.

2.3. *Computer-aided detection of microcalcification clusters*

We show in Fig. 2 a block diagram to illustrate the building blocks in a typical computer-aided framework for MCC detection. In the block diagram, a mammogram image is first digitized, either from film or from a digital detector, with a resolution that could range from 30–$200\,\mu\mathrm{m}$ and 10–16 bits per pixel. A region of interest (ROI) of the image could be selected for processing. In most techniques the image is enhanced prior to detection for the purposes of improving contrast and reducing noise by using techniques such as histogram equalization or digital filtering. This is followed by a detection algorithm for identifying the MCs from the image. The detected individual MCs are then grouped into clusters by a clustering algorithm.

The performance by a computer-aided detection algorithm is typically evaluated by means of free-response receiver operating characteristics (FROC) curves.[15] An FROC curve is an extension of the classical receiver operating characteristics (ROC)[16] curve. It is a plot of the correct detection rate (i.e. true-positive fraction) achieved by a classifier versus the average number of false positives per image varied over the continuum of the decision threshold. FROC is considered as an extension of ROC in that it also takes into account the correct localization of the detected clusters. An FROC curve provides a comprehensive summary of the trade-off between detection sensitivity and specificity.

In the literature, there also exist several other strategies for the evaluation of detected clustered MCs, and therefore caution should be taken when comparing the published results of different algorithms.[17]

It should be mentioned in passing that there exists a large body of work on methods for classification or differentiation of benign from malignant MCs for the purpose of computer-aided diagnosis;[12] in addition, there also exists work on computer-aided detection and diagnosis of masses.[18] However, this chapter focuses only on detection of MCs.
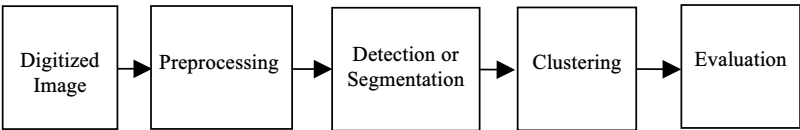

Fig. 2.   Generic block diagram for MCC detection.

3. Techniques for Detection of Microcalcification Clusters (MCCs)

Computer-assisted mammography has been the subject of extensive research over the past two decades and several review articles have been written on this subject.[19–23] In this chapter we focus on methods for detection of MCs, particularly those reported most recently in the literature. Due to the large volume of existing work, it is not possible to describe each method in detail. Instead, we adopt the following approach: based on the underlying image-processing methods used, we organize the detection techniques into four major categories. Within each category, we first provide a summary of existing methods in the literature, and then give an in-depth description of one or two representative methods. Furthermore, we will also present some numerical results to illustrate these representative methods whenever it is useful.

Based on our initial review of some 200 recent journal articles related to the topic of CAD systems for MCC detection, we selected among them 39 articles for our review in this chapter. Our selection was based on the following criteria: the MCC detection methods in these articles all involved image processing techniques, and were reported with results from clinical mammogram data. We categorize these detection methods into the following four categories: (1) basic image-enhancement, (2) multiscale decomposition, (3) stochastic modeling, and (4) machine learning. A summary of selected techniques in these categories is given in Table 1.

As pointed out above, such categorization of the different techniques is purely based on the underlying image-processing methods employed, and is by no means intended to be exclusive. Some of the techniques can be easily found to belong to more than one category. For example, a neural network approach may use wavelet-based features as input. However, we find it appropriate for the presentation of the material. Finally, and most importantly, this review is mainly intended for pedagogical purposes, and is not intended to be comprehensive (as several rather comprehensive reviews already exist).

3.1. *Basic image enhancement techniques*

Methods in this category are motivated by the fact that MCs tend to be brighter than their surroundings. The basic idea is to first employ image enhancement methods to improve the contrast of MCs, and then apply thresholding to separate them out. Specifically, Karssemeijer[24] applied adaptive histogram equalization for MC detection. Morrow *et al.*[25] applied a region based enhancement method, wherein

Table 1.   Summary of selected MCC detection techniques.

Category	Number of publications	Percentage
Image enhancement	22	56%
Multiscale decomposition	4	10%
Stochastic model	8	21%
Machine learning	5	13%

adaptive region growing was applied to identify MC regions, and the image contrast was computed for each region with respect to the background. Chan *et al.*[26] used unsharped masking, in which a mammogram image was first lowpass filtered and then subtracted from its original; this process was equivalent to highpass filtering, where high frequency features such as MCs and edges were emphasized and slowly varying background tissues were suppressed. Dengler *et al.*[27] developed a method based on the difference of Gaussians (DoG), and applied morphological operators to extract the shape features of the detected MCs. Bankman *et al.*[28] applied "region growing" in conjunction with active contours, where the seed points were selected as the local maxima found by an edge detection operator.

An exemplary method in this category is the filtering approach developed by Nishikawa *et al.*[29] The method is based on a difference image technique followed by morphological post-processing. It applies two linear filters, whose coefficients are given in Fig. 3. The first filter (i.e. matched filter) enhances the MCs while the second filter (i.e. box-rim filter) suppresses them. The "difference image" is obtained by subtracting the suppressed image from the enhanced one. In practice, the two filters are combined into a single linear filter. This is followed by morphological erosion such that very small signals (less than 3 pixels in area) that are likely caused by random noise are eliminated.

To illustrate these steps, we show in Fig. 4 some results obtained with this technique. Figure 4(a) is a sample ROI image, where MC locations are identified by circles; Fig. 4(b) is the difference image obtained with filtering, where MCs are visibly enhanced and background tissues are suppressed; Fig. 4(c) is the thresholded result, where only the brightest 2% of the pixels in Fig. 4(b) are shown; Fig. 4(d) is the result after morphological erosion.

In summary, the advantage of image enhancement techniques lies in their simplicity and computational efficiency. However, they are often highly susceptible to

0.75	0.75	0.75
0.75	1	0.75
0.75	0.75	0.75

(a)

0.125	0.125	0.125	0.125	0.125	0.125	0.125	0.125	0.125
0.125	0	0	0	0	0	0	0	0.125
0.125	0	0	0	0	0	0	0	0.125
0.125	0	0	0	0	0	0	0	0.125
0.125	0	0	0	0	0	0	0	0.125
0.125	0	0	0	0	0	0	0	0.125
0.125	0	0	0	0	0	0	0	0.125
0.125	0	0	0	0	0	0	0	0.125
0.125	0.125	0.125	0.125	0.125	0.125	0.125	0.125	0.125

(b)

Fig. 3.   The two filters used in the image difference technique: (a) the matched filter, and (b) the box-rim filter.
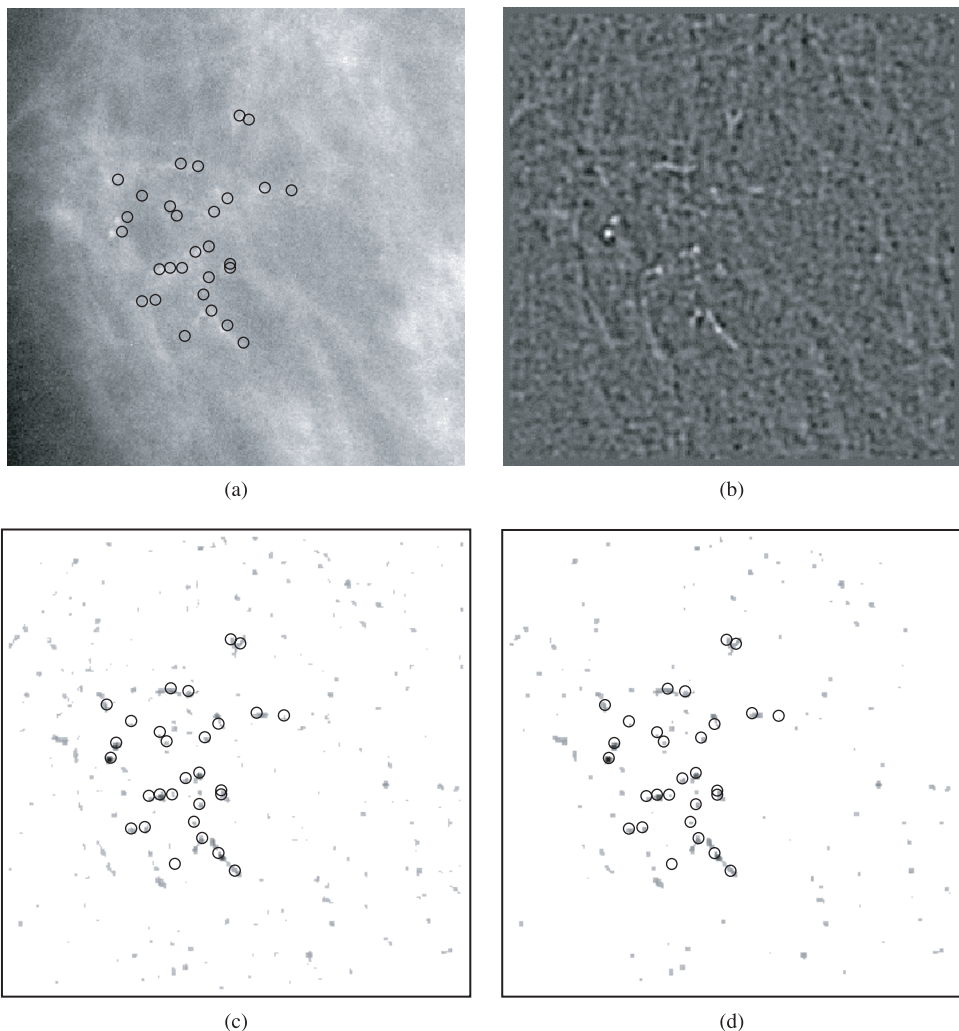
Fig. 4. Illustration of the image difference method: (a) a sample ROI with 29 MCs, (b) difference filtering output, (c) output after thresholding (brightest 2% of the pixels), and (d) result after morphological erosion.

errors, because under-enhanced regions will increase the false negative rate, (i.e. missing detection of true MCs), while over-enhancement will increase the false positive rate, (i.e. normal tissues being detected as MCs[23]).

3.2. *Multiscale decomposition methods*

Methods in this category are based on the fact that the bright MC spots are different in frequency content from their surrounding background. In particular, Nitsch and Peitgen[30] applied multiscale analysis using the Laplacian-of-Gaussian filter.

By varying the width of the filter, the mammogram image was transformed into different scales. The response of the filter was thresholded to detect MCs.

Perhaps the most popular techniques in this category are those based on the use of wavelet transforms. Yoshida *et al.*[31] used a decimated wavelet transform and the difference image technique for the detection of MCs; the wavelet transform was applied as an enhancement technique to improve the visibility of MCs by reconstructing mammograms from selected subbands, namely, subbands 2 and 3. On the other hand, Strickland and Hahn[32,33] used undecimated biorthognal wavelet transforms and optimal subband weighting for detecting and segmenting clustered MCs. The undecimated wavelet transform has the advantage of being translation invariant.

Below we provide a detailed explanation of the undecimated wavelet transform approach by Strickland and Hahn. The wavelet transform uses biorthognal basis of length four. It functions as a bank of multi-scale matched filters under the assumption that the MCs can be represented on average by a 2-D circular Gaussian shape of which the width varies along the different scales. Four-octave decomposition is used along with two voices inserted between octaves 2-3 and 3-4 to account for the coarse approximation of the Gaussian shape model. Optimal weighting is applied to the detail subbands using a Fisher discriminant. At each pixel location $\mathbf{x}$ a feature vector is formed by:

$$v(\mathbf{x}) = [v_1(\mathbf{x}), v_2(\mathbf{x}), \ldots, v_{2N+1}(\mathbf{x})]^T, \tag{1}$$

where N is the number of octaves (total of 6: 4 scales plus 2 voices), and the components are given by:

$$\begin{aligned} v_{2j-1}(\mathbf{x}) &= d^j_{LH}(\mathbf{x}) + d^j_{HL}(\mathbf{x}) \\ v_{2j}(\mathbf{x}) &= d^j_{HH}(\mathbf{x}), \quad 1 \le j \le N, \\ v_{2N+1}(\mathbf{x}) &= s^N_{LL}(\mathbf{x}) \end{aligned} \tag{2}$$

where $d^j_{LH}(\mathbf{x}) + d^j_{HL}(\mathbf{x})$ is the summation of the vertical and the horizontal details at octave j, $d^j_{HH}(\mathbf{x})$ is the diagonal detail, and $s^N_{LL}(\mathbf{x})$ is the approximation at the last octave. The linear discriminant function is given by:

$$f(\mathbf{x}) = \mathbf{w}^T v(\mathbf{x}), \tag{3}$$

where the parameters $\mathbf{w}^T$ are obtained during the training phase to maximize the separation between the two classes.

In Fig. 5 we show some results to illustrate these steps. Figure 5(a) is a sample ROI, where MCs are marked by circles; Fig. 5(b) is the output of the Fisher discriminant. In Fig. 6 we show the undecimated subbands corresponding to the 3rd octave, where the MCs are visibly enhanced.

In summary, the multiscale decomposition methods tend to be more robust than simple image-enhancement techniques. However, it relies on the assumed models for the MC (e.g. approximating it by the Gaussian shape), which may not be always accurate due to the wide variation in the morphology of MCs.
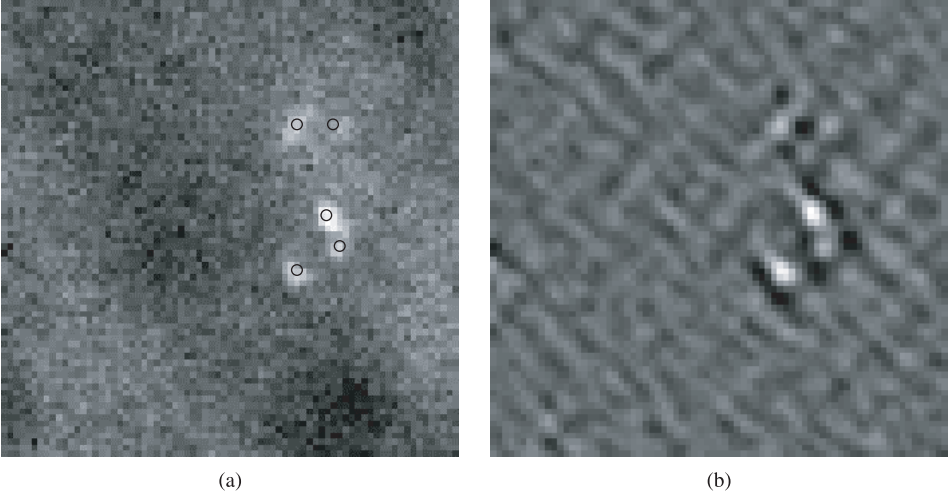
Fig. 5.   Illustration of wavelet based detection: (a) a sample ROI with 5 MCs, and (b) the final Fisher discriminant output.

3.3. *Stochastic model methods*

Methods in this category aim to directly exploit the statistical difference between MCs and their surroundings. Namely, MCs are known to exhibit strong local spatial correlations, which could be used to separate them from the surrounding tissues. Gurcan *et al.*[34] proposed an approach based on higher order statistics, wherein it was conjectured that areas with no MCs would have a Gaussian like distribution and areas with MCs would have non-zero skewness (3rd moment) and kurtosis (4th moment), of which both are zero for a Gaussian distribution. In this approach, a mammogram image was first processed using a subband filterbank decomposition. The bandpass subimage was divided spatially into nonoverlapping blocks. The higher order statistics were then computed for each block in the bandpass subimage to detect the MCs.

Karssemeijer and Veldkamp[35,36] developed a method based on Bayesian estimation using a Markov random field (MRF) to segment MCs. In this approach, an MRF is defined for modeling the pixels in a mammogram image. The MRF model is specified by the conditional probability distribution of the label at each pixel given its gray level and the labels of its neighboring pixels. With a Gaussian distribution assumed for the fluctuation of gray levels due to noise, this conditional probability is modeled as

$$\log P[x_i = k | y_i, \mathbf{n}_i] \propto -\alpha(k) + \beta g(k) - \frac{(y_i - \mu_l)^2}{2\sigma_l^2}, \tag{4}$$

where x_i is the label of pixel i ($k = 0$ for background and $k = 1$ for MC, respectively), y_i is the gray level of pixel i, $\mathbf{n}_i$ denotes the labels of its neighboring pixels, $\alpha(k)$ is an offset value, $g(k)$ is the number of neighbors with label k, β is a parameter

Approximation $s^3_{LL}(\mathbf{x})$ Horizontal detail $d^3_{LH}(\mathbf{x})$

Vertical detail $d^3_{HL}(\mathbf{x})$ Diagonal detail $d^3_{HH}(\mathbf{x})$
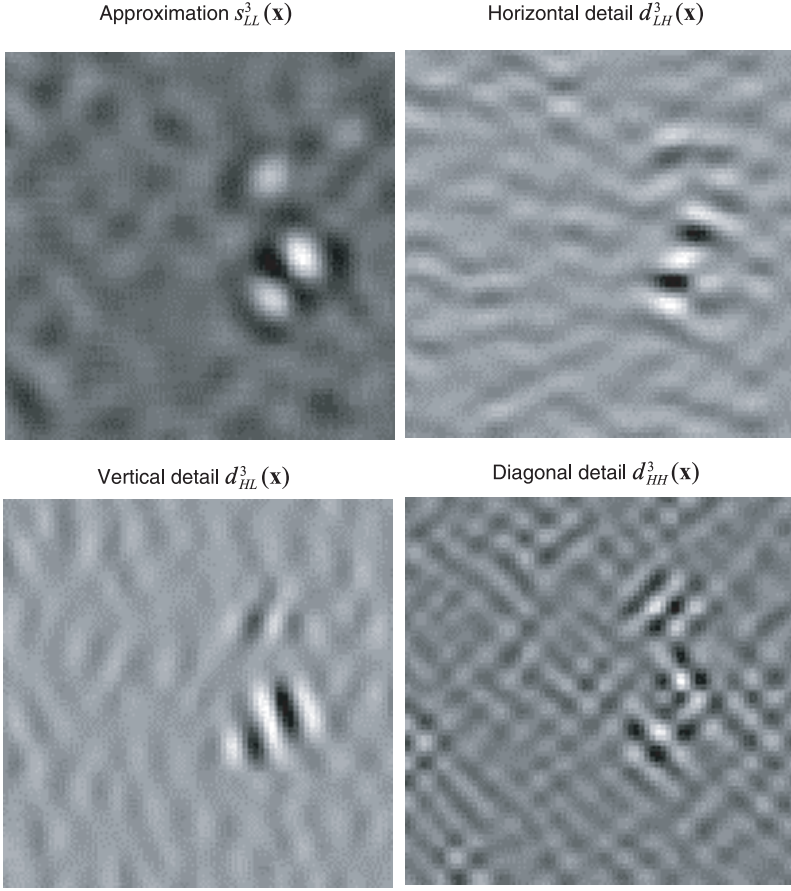
Fig. 6.   The 3rd octave subbands of the image in Fig. 5(a).

used to model the *a priori* knowledge that neighboring pixels likely have the same labels, μ_l and σ_l are the respective mean and standard deviation for each of the two classes. The neighborhood used for $\mathbf{n}_i$ consists of the four nearest neighbors of pixel i. The pixel label x_i is estimated by maximizing the conditional probability in (4), that is,

$$\tilde{x}_i = \arg\max_k \{\log P[x_i = k|y_i, \mathbf{n}_i]\}. \tag{5}$$

The maximization in (5) is performed using an iterative procedure known as estimation by iterated conditional modes (ICM). It is noted that this maximization requires an initial estimate of the labeling. This can be obtained by initially setting $g(k)$ to zero. Interested readers are referred to Karssemeijer and Veldkamp[35,36] for more details.

3.4. *Machine learning methods*

Methods in this category aim to learn dependencies from data.[37] In the context of MC detection, it is a binary classification process, where the goal is to determine whether an MC is present at a pixel location or not. There are two machine learning models frequently used in the literature: one is based on neural networks (NN), and the other is based on support vector machines (SVMs).

3.4.1. *Neural networks (NN)*

Neural networks represent a powerful technique that has been applied successfully to model many human decision-making tasks. In particular, because of their adaptive and learning capabilities, neural networks have been applied to solve a variety of pattern recognition problems.[38] Feedforward neural networks have been widely used for MC detection and classification.[39–41] As an example, we describe below a two-stage network approach by Yu and Guan,[42] where wavelet components, gray-level statistics, and shape features were used to train the network, which was formed by a cascade of two separate 3-layer feedforward neural networks.

The first stage of the network by Yu and Guan was used to segment out the potential MC pixels in a mammogram image. It had 6 neurons in the hidden layer. At each pixel location $\mathbf{x}$, the input feature vector was formed by

$$v(\mathbf{x}) = [v_1(\mathbf{x}), v_2(\mathbf{x}), v_3(\mathbf{x}), v_4(\mathbf{x})]^T, \tag{6}$$

where $v_1(\mathbf{x}), v_2(\mathbf{x})$ were wavelet features constructed from the second and the third subbands, respectively, $v_3(\mathbf{x})$ was the median contrast, and $v_4(\mathbf{x})$ was the normalized gray level value within a window of $M \times M$ centered at $\mathbf{x}$. The wavelet transform was based on Daubechies' orthogonal wavelet of length four.

After the first neural network was trained, its output was thresholded to generate a binary map of potential MCs. The threshold value was selected empirically during the training phase such that the majority of MCs appeared in the binary map. A set of 15 features based on shape and texture, chosen by a sequential backward selection procedure, was used to train the second neural network to detect the final individual MCs. The second neural network had 8 neurons in the hidden layer. The input features were: (1) area of the object, (2) compactness of the object (perimeter^2/area), (3) standard deviation of the gray-level values in the object, (4) elongation (ratio of the major to the minor axis of the best fitted ellipse), (5) foreground-background difference, (6) the 4th shape moment, (7) the 5th invariant moment. Features 8–15 are second-order histogram features, which include: (8) correlation, (9) variance, (10) contrast, (11) entropy, (12) sum entropy, (13) inverse difference moment, (14) difference ratio, and (15) angular second moment.

In Fig. 7 we show some results to illustrate this two-stage network approach. Figure 7(a) shows a sample ROI with MCs marked by circles, Fig. 7(b) shows the output by the first neural network, and Fig. 7(c) shows the output by the second neural network.
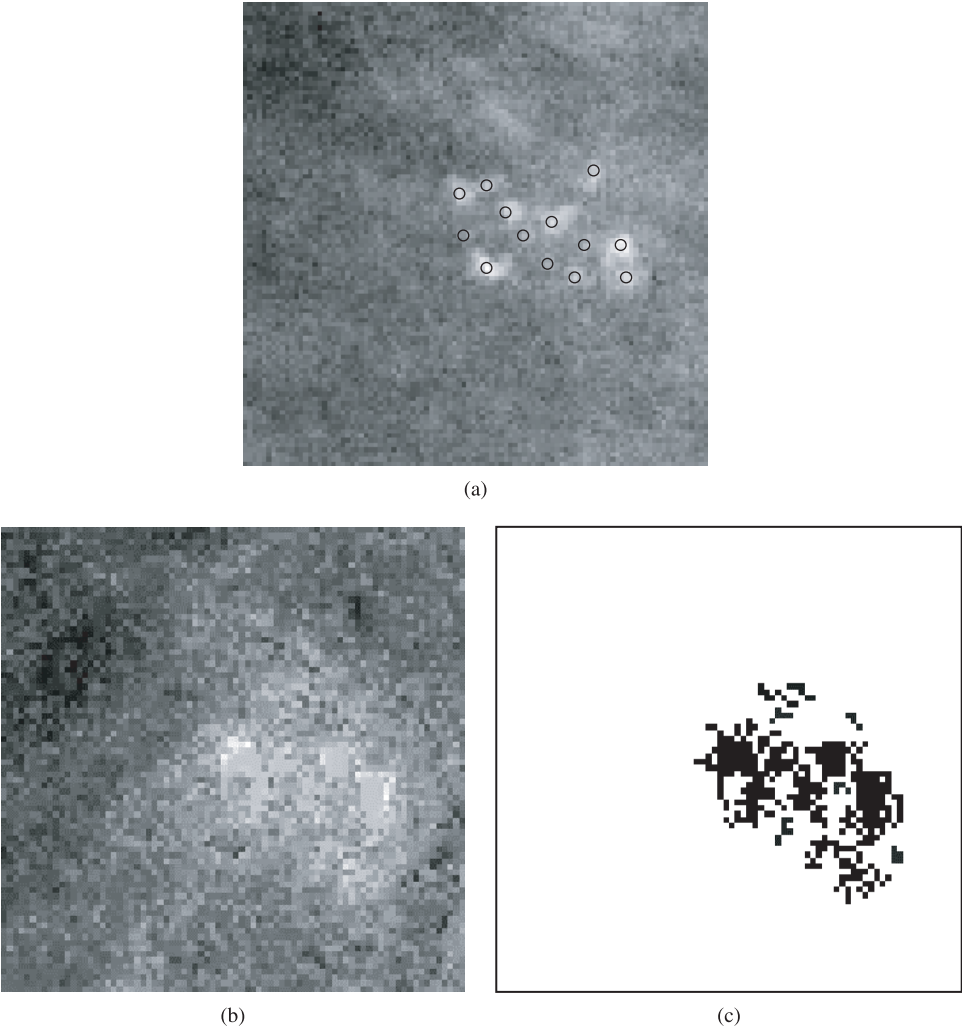
Fig. 7.   Illustration of the two-stage neural network approach: (a) a sample ROI with 13 MCs, (b) output of the first network, and (c) output of the second network.

3.4.2. *Support vector machine (SVM)*

SVM is a universal constructive learning procedure based on statistical learning theory.[43] It can be applied both in the context of classification (e.g. detection of MCs) and in the context of regression. In classification, the main idea of this technique is to separate the classes with a hyperplane that maximizes the margin between them (Fig. 8). It is an implementation of the structural risk minimization principle, which aims to minimize the bound on the generalization error of a model rather than minimizing the mean square error over the data set. The discriminant function in SVM is characterized by a subset of the training samples known as

Fig. 8.   Support vector machine classification with a hyperplane that maximizes the separating margin between the two classes (indicated by data points marked by 'X's and 'O's). Support vectors are elements of the training set that lie either within the margin or on the boundary hyperplanes of the two classes.

support vectors $\mathbf{s}_i$:

$$f(\mathbf{x}) = \sum_{i=1}^{ls} \alpha_i y_i K(\mathbf{s}_i, \mathbf{x}) + \alpha_o, \tag{7}$$

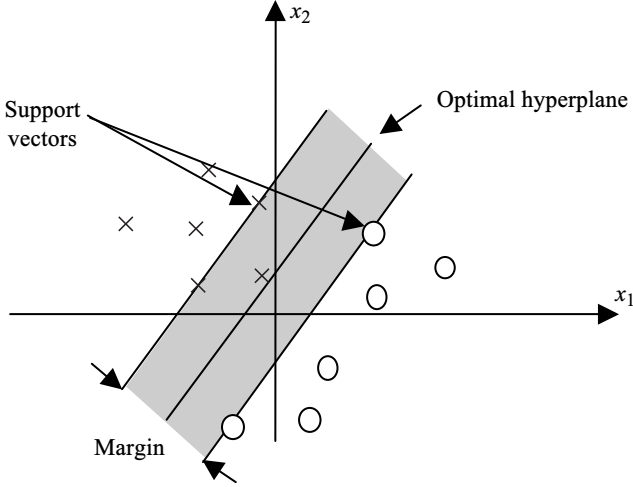where l_s is the number of support vectors, y_i are the training decisions, $K(\cdot, \cdot)$ is a kernel function satisfying Mercer's conditions,[43] and α_i are coefficients determined by quadratic programming.

Bazzani *et al.*[44] proposed a method for MC detection based on multiresolution filtering analysis and statistical testing, in which an SVM classifier was used to reduce the false detection rate. In our own work,[45] we developed an SVM approach for detection of MCs in mammogram images, wherein an SVM classifier was trained through supervised learning to determine whether an MC is present or not at any location under consideration in a mammogram. Below we describe this approach in detail.

In our approach, we define the input pattern $\mathbf{x}$ to the SVM classifier to be a small $M \times M$ pixel window centered at the location of interest. This is based on the fact that individual MCs are well localized in a mammogram; therefore, to detect whether an MC is present at a given location in a mammogram, we can simply examine the image content surrounding that location.

The window size is chosen to be large enough to contain most MCs, but small enough to avoid potential interference from neighboring MCs. A small window size is also desirable for computational efficiency. In our study, the mammograms used were digitized at a resolution of 0.1 mm/pixel, and we chose $M = 9$.
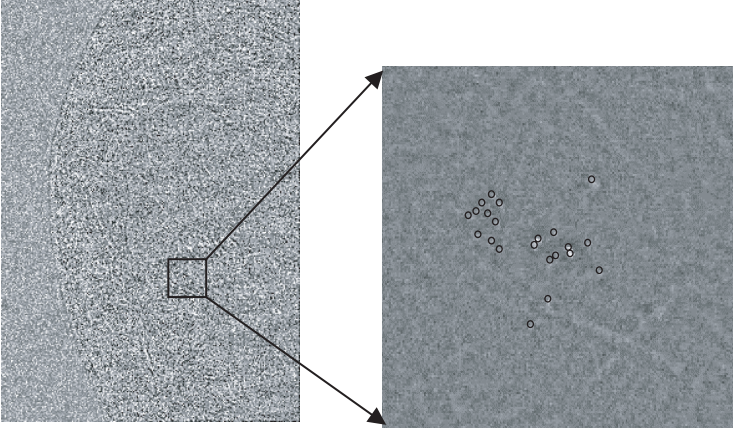
Fig. 9. The mammogram in Fig. 1 after preprocessing with a highpass filter. The inhomogeneity of the background has been notably mitigated.

It is important to note, as can be seen in Fig. 1, that the background in a mammogram is not homogeneous. To ensure that the input patterns to the SVM classifier are somewhat uniform throughout the mammogram, we preprocessed the mammogram by applying a sharp high-pass filter to it. This filter was designed as a linear-phase FIR filter with 3-dB cutoff frequency $\omega_c = 0.125$ and length 41. As an example, we show in Fig. 9 the result after filtering the mammogram in Fig. 1 with this filter. It appears to be effective in mitigating the impact of inhomogeneity of the background.

For the determination the parameters in the SVM function in (7), a set of training mammograms was used. For each MC location in the training set mammogram, a window of $M \times M$ image pixels centered at its center of mass was extracted; the vector formed by this window of pixels, denoted by $\mathbf{x}_i$, was then treated as an input pattern for the "MC present" class ($y_i = +1$). "MC absent" samples were collected ($y_i = -1$) similarly, except that their locations were selected randomly from the set of all "MC absent" locations in the training mammograms. In this procedure no window in the training set was allowed to overlap with any other training window. The reason for using only a random subset of "MC absent" examples was that there were too many "MC absent" examples to be used at once practically. The SVM parameters were estimated using a multi-fold cross validation procedure, in which it was concluded that a radial basis function performed best.

We show in Fig. 10 some examples of the support vectors obtained for both MC and non-MC image windows. For comparison, some randomly selected examples from the training set are also shown. Note that the support vectors appear to be the difficult-to-classify, "borderline" cases; i.e. the MC support vectors are MCs that could be easily mistaken for background regions, and the non-MC support vectors are background regions that appear like MCs.

"MC present" class templates   "MC present" class support vectors

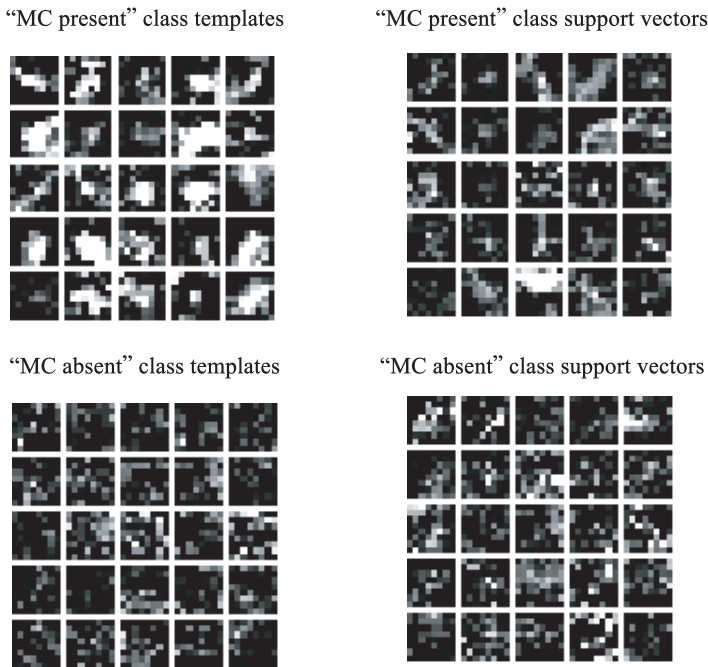"MC absent" class templates   "MC absent" class support vectors

Fig. 10. Examples of 9×9 image windows and support vectors. Image windows with and without microcalcifications (MCs) are shown at top-left and bottom-left, respectively. Support vectors representing the "MC present" and "MC absent" classes of image windows are shown at top-right and bottom-right, respectively. Note that the support vectors (SVs) represent the borderline examples from each class that are difficult to categorize ("MC present" SVs could be mistaken for "MC absent" image regions; "MC absent" SVs might be mistaken for MCs). The support vectors shown are for the case of a SVM with Gaussian kernel.

A difficult problem in training a classifier for MC detection is that there are a very large number of image locations where no MC is present, so that the training set for the "MC absent" class can be impractically large. Thus, there arises an issue of how to select the training examples so that they well represent the class of "MC absent" locations. To solve this problem, we proposed a technique called successive enhancement-learning (SEL) to select the training examples. SEL selects iteratively the "most representative" MC-absent examples from all the available training images while keeping the total number of training examples small. The steps of SEL are illustrated in Table 2.

We show in Fig. 11 some results to illustrate the SVM approach. Figure 11(a) shows a sample ROI with MCs marked by circles; Fig. 11(b) shows the SVM output. Note that the image in Fig. 11(a) is a particularly difficult case because of the apparent similarity between the MCs and the background. Nevertheless, the SVM output successfully highlights the presence of MCs, reducing false alarm detection.

Table 2.   Successive enhancement learning algorithm.

---

1. Extract an initial set of training examples from the available training
    images(e.g., through random selection).
    Let $Z = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \ldots, (\mathbf{x}_l, y_l)\}$ denote this resulting set of training
    examples.
2. Train the SVM classifier

$$f(\mathbf{x}) = \sum_{i=1}^{ls} \alpha_i y_i K(\mathbf{s}_i, \mathbf{x}) + \alpha_o \quad \text{with } Z.$$

3. Apply the resulting classifier $f(\mathbf{x})$ to all the mammogram regions (except
    those in Z) in the available training images and record the   "MC absent"
    locations that have been misclassified as "MC present".
4. Gather N new input examples from the misclassified "MC absent" locations;
    update the set Z by replacing N "MC absent" examples that have been
    classified correctly by $f(\mathbf{x})$ with the newly collected "MC absent" examples.
5. Retrain the SVM classifier with the updated set Z.
6. Repeat Steps 3–5 until convergence is achieved.


3.5. *Other methods*

In addition to the methods described in the four categories above, there also exist
other interesting approaches. Li *et al.*[46] developed a technique, in which the image
background was modeled by fractals, which resulted in enhanced MC detection in
the difference between the original and the modeled image.

Cheng *et al.*[47] proposed a fuzzy logic approach, wherein a mammogram image
was transformed into a fuzzified image according to the maximum fuzzy entropy
principle. This was based on the fact that mammograms have some degree of fuzzi-
ness such as indistinct borders, ill-defined shapes, and different densities. Geomet-
rical statistics were then used to measure the nonuniformity of different regions. To
further reduce the FP fate, a curve detector was employed to remove those line-like
or curve-like irrelevant breast structures.

4. Evaluation Studies and Discussions

In the literature the different MC detection methods were often developed using
different mammogram data sets, and were evaluated using different performance
criteria. Thus, it becomes extremely difficult to compare these methods directly
based on their respectively reported results. Here we present a set of evaluation
results of the different methods, which was obtained in the development of our
SVM approach.[45]

In our study we used a data set provided by the Department of Radiology at
the University of Chicago. This data set consisted of 76 clinical mammograms, all
containing multiple MCs. These mammograms were of dimension 1000×700 pixels,
with a spatial resolution of $0.1\,\text{mm/pixel}$ and 10-bit grayscale. Collectively, there
are a total of $1120\,\text{MCs}$ in these mammograms, which were identified by a group of
experienced mammographers.
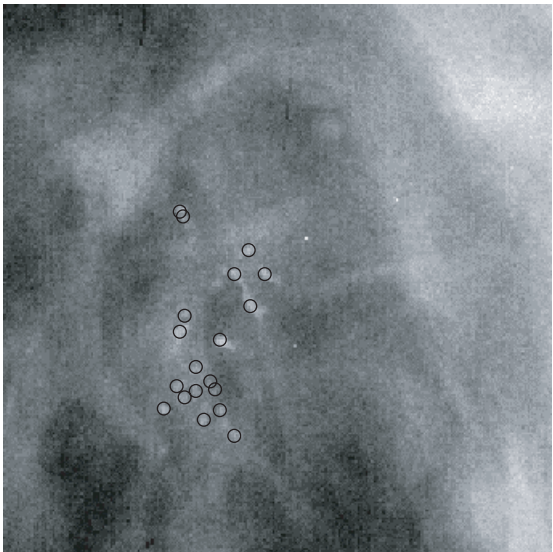
(a)
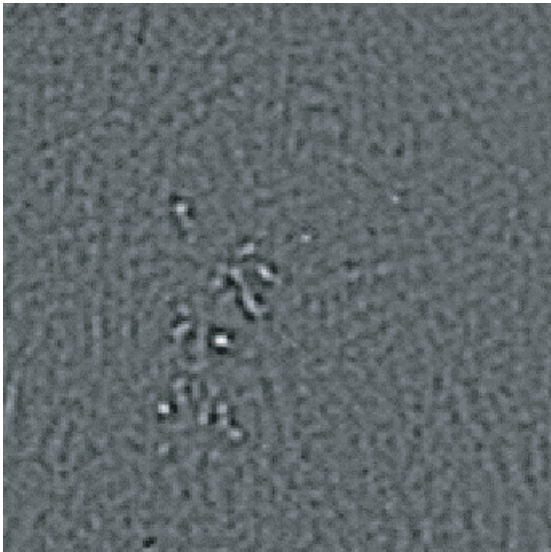
(b)

Fig. 11.   Illustration of the SVM approach: (a) a sample ROI with 19 MCs, and (b) SVM classifier output.

The data set was divided randomly into two separate sets, each consisting of 38 images. One set was used exclusively for training purposes (required in the case of SVM, wavelet, and neural network-based methods) and the other set was used exclusively for testing of all methods.

The evaluation was performed using FROC curves generated in accordance with a criterion recommended by Kallergi *et al.*[17] for identifying MC clusters. Specifically, a group of objects classified as MCs is considered to be a true positive (TP) cluster only if: (1) the objects are connected with nearest-neighbor distances less than 0.2 cm, and (2) at least three true MCs are detected by the algorithm within an area of 1 cm^2. Likewise, a group of objects classified as MCs is labeled as a false positive (FP) cluster if the objects satisfy the cluster requirement but contain no true MCs. Such a criterion has been reported to yield more realistic performance than several other alternatives.

Besides our SVM approach, the following four methods for MC detection were also considered: (1) the image difference technique (IDTF) by Nishikawa *et al.*,[29] (2) the difference of Gaussians (DoG) method by Dengler *et al.*,[27] (3) the wavelet-decomposition based (WD) method by Strickland and Hahn,[32,33] and (4) the two-stage multi-layer neural network (TMNN) method by Yu and Guan.[42] These methods were chosen to represent the different categories of methods that we reviewed above, with the exception of the stochastic model methods. However, in Yu and Guan[42] a comparison of TMNN with the stochastic method by Karssenmeijer[35] was provided using the Nijmegen database, where superior performance by TMNN was demonstrated.

The performance of these different methods is summarized in Fig. 12 with FROC curves. As can be seen, the SVM classifier offered the best detection result. In addition, the proposed SEL scheme further improved the correct detection rate. In particular, it achieved a sensitivity of approximately 85% when the false detection rate was at an average of one FP cluster per image.

We note that the FROC results in Fig. 12 for WD and IDTF are very similar to those described in the original reports of these methods. For the DoG method (for which no FROC information was given in its original report), the detection rate is close to that of the IDTF when the FP rate is around two FP clusters per image. This is not surprising because both methods operate under a similar principle (the detection kernels in both cases behave like a band-pass filter).

In addition, the FROC results indicate that the TMNN method outperforms the other three methods we compared (WD, IDTF, and DoG) when the FP rate is above one FP cluster per image. The numerical FROC results we obtained for the TMNN are somewhat different from those in its original report. There are several possible explanations: (1) the mammogram set used was different, (2) the detection criterion for MC clusters used in performance evaluation was different; and (3) in the original work[42] the MC clusters used for training were also included in testing.

Finally, we demonstrate in Fig. 13 that even the method of defining MC clusters has an influence on the FROC curves, making it difficult to compare reported results in the literature when they were derived using various criteria. The results in Fig. 13 were obtained when the nearest-neighbor-distance threshold for MC cluster detection was increased from 0.2 cm to 0.5 cm. In particular, the sensitivity of the SVM approach increased to as high as 94% while the FP rate remains at one FP
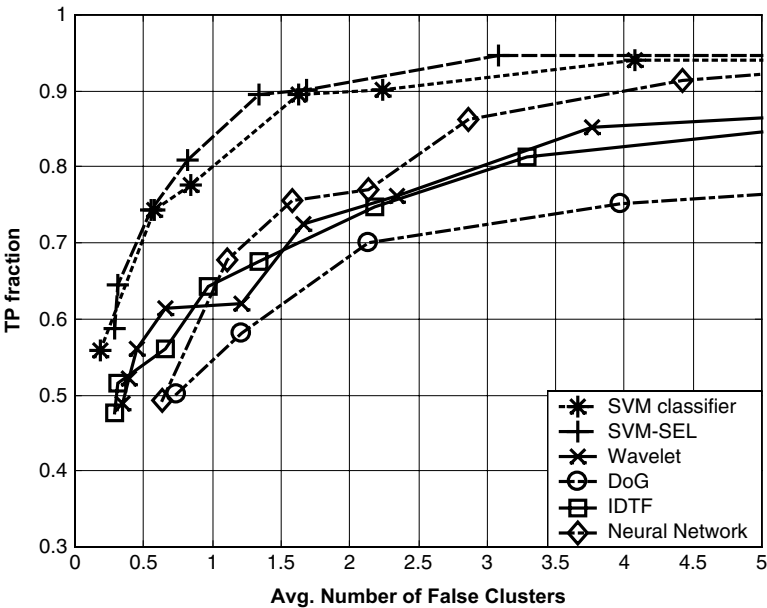
Fig. 12. FROC comparison of the different methods tested. A higher FROC curve indicates better performance. The best performance was obtained by a successive learning SVM classifier, which achieves 85% detection rate at a cost of one FP cluster per image. The nearest neighbor distance threshold used for cluster detection is 0.2 cm.
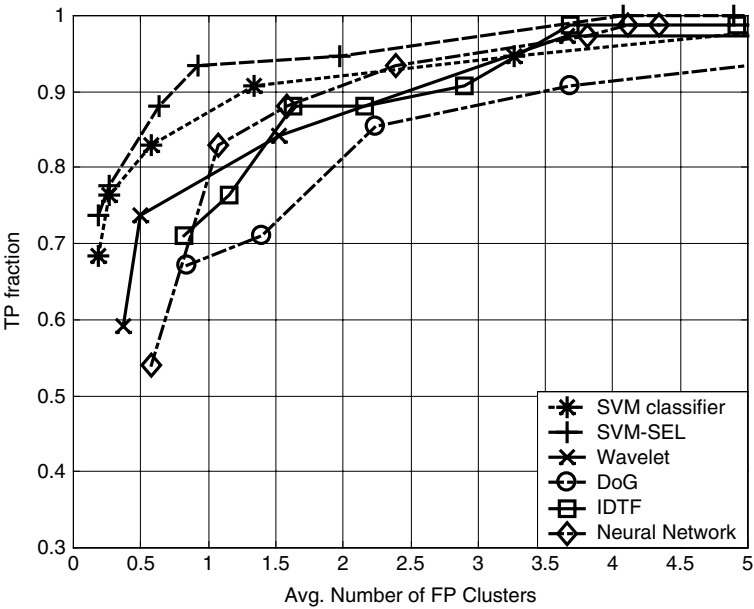
Fig. 13. FROC comparison of the different methods tested. The nearest neighbor distance threshold used for cluster detection is 0.5 cm.

cluster per image. Note that, while different criteria may affect the numerical FROC results, the relative ordering of performance of the methods is preserved.

5. Conclusions

Clustered microcalcifications are important features for early detection of breast cancer. In this chapter, we provided an overview of a number of state-of-the-art computerized methods for detection of MCs in mammograms, ranging from methods based on simple image enhancement, multiscale decomposition, or stochastic modeling, to those based on machine learning. While those methods based on image enhancement are simple and can offer reasonably good performance, more sophisticated methods based on machine learning tend to be more accurate at the expense of increased computational complexity. Despite recent advances in this field, computerized MC detection is still far from being perfect.

Acknowledgments

References

1. Breast Cancer Facts and Figures 2003–2004, *American Cancer Society* (2004).
2. A. I. Mushlin, R. W. Kouides and D. E. Shapiro, Estimating the accuracy of screening mammography: A meta-analysis, *Am. J. Prev. Med.* **14** (1998) 143–53.
3. R. E. Bird, T. W. Wallace and B. C. Yankaskas, Analysis of cancers missed at screening mammography, *Radiology* **184** (1992) 613–617.
4. E. A. Sickles, Mammographic features of 300 consecutive non-palpable breast cancers, *American Journal of Radiology* **146** (1986) 661–663.
5. D. B. Kopans, The positive predictive value of mammography, *American Journal of Radiology* **158** (1992) 521–526.
6. J. G. Elmore, C. K. Wells, C. H. Lee, D. H. Howard and A. R. Feinstein, Variability in radiologists' interpretations of mammograms, *N. Eng. J. Med.* **331** (1994) 1493–1499.
7. M. L. Giger, Computer-aided diagnosis in Radiology, *Acad. Radiol.* **9**(3) (2002).
8. P. E. Shile and J. A. Guingrich, Detecting "missed" breast cancers: A comparison of CAD systems, *Supplement to Applied Radiology*, October (2002).
9. M. T. Nelson and M. C. Lechner, Comparison of three commercially available FDA approved computer-aided detection (CAD) systems, *Proceedings RSNA 2002*, Chicago, IL, *Radiology*, **225** (2002) 600.
10. P. Bandodkar, R. L. Birdwell and D. M. Ikeda, Computer aided detection (CAD) with screening mammography in an academic institution. Preliminary findings, *Proceedings RSNA 2002*, Chicago, IL, *Radiology* **225** (2002) 600.
11. T. W. Freer and M. J. Ulissey, Screening mammography with computer-aided detection: Prospective study of 12,860 patients in a community breast center, *Radiology* **220** (2001) 781–786.

12. M. Kallergi, Computer-aided diagnosis of mammographic microcalcificayion clusters, *Medical physics* **34**(2) (2004) 314–326.

13. D. B. Kopans, Breast Imaging, *J.B. Lippincott company* (1989).

14. American College of Radiology, Illustrated Breast Imaging Reporting and Data System, 3rd edn. (1998).

15. P. C. Bunch, J. F. Hamilton, G. K. Sanderson and A. H. Simmons, A free-response approach to the measurement and characterization of radiographic-observer performance, *J. Appl. Eng.* **4** (1978) 166–177.

16. R. G. Swensson, Unified measurement of observer performance in detecting and localizing target objects on images, *Med. Phys.* **23** (1996) 1709–1725.

17. M. Kallergi, G. M. Carney and J. Garviria, Evaluating the performance of detection algorithms in digital mammography, *Med. Phys.* **26** (1999) 267–275.

18. M. K. Markey, J. Y. Lo and C. E. Floyd, Jr., Differences between the computer-aided diagnosis of breast masses and that of calcifications, *Radiology* **223** (2002) 489–493.

19. N. Karssemeijer and J. Hendriks, Computer-assisted reading of mammograms, *European Radiology* **7** (1997) 743–748.

20. K. Doi, H. MacMahon, S. Katsuragawa, R. M. Nishikawa and Y. L. Jiang, Computer-aided diagnosis in radiology: Potential and pitfalls, *European Journal of Radiology* **31** (1999) 97–109.

21. M. L. Giger, Computer-aided diagnosis of breast lesions in medical images, *Computing in Science & Engineering* **2** (2000) 39–45.

22. L. Zhang, R. Sankar and W. Qian, Advances in microcalcification clusters detection in mammography, *Computers in Biology and Medicine* **32** (2002) 515–528.

23. H. D. Cheng, X. P. Cai, X. W. Chen, L. M. Hu and X. L. Lou, Computer-aided detection and classification of microcalcifications in mammograms: A survey, *Pattern Recognition* **36** (2003) 2967–2991.

24. N. Karssemeijer, Adaptive noise equalization and image analysis in mammography, *Information Processing in Medical Imaging*: 13th International Conference, IPMI '93, AZ, USA (1993) 472–486.

25. W. M. Morrow, R. B. Paranjape, R. M. Rangayyan and J. E. L. Desautels, Region-based contrast enhancement of mammograms, *IEEE Trans. Med. Imag.* **11**(3) (1992) 392–406.

26. H. P. Chan, C. J. Vyborny, H. MacMahon, C. E. Metz, K. Doi and E. A. Sickles, Digital mammography: ROC studies of the elects of pixel size and unsharp-mask filtering on the detection of subtle microcalcifications, *Invest. Radiol.* **22**(7) (1987) 581–589.

27. J. Dengler, S. Behrens and J. F. Desaga, Segmentation of microcalcifications in mammograms, *IEEE Trans. on Medical Imaging* **12**(4) (1993) 634–642.

28. I. N. Bankman, T. Nizialek, I. Simon, O. B. Gatewood, I. N. Weinberg and W. R. Brody, Segmentation algorithms for detecting microcalcifications in mammograms, *IEEE Trans. on Information Technology in Biomedicine* **1**(2) (1997) 141–149.

29. R. M. Nishikawa, M. L. Giger, K. Doi, C. J. Vyborny and R. A. Schmidt, Computer aided detection of clustered microcalcifications in digital mammograms, *Med. Biol. Eng. Compu.* **33** (1995) 174–178.

30. T. Netsch and H. O. Peitgen, Scale-space signatures for the detection of clustered microcalcifications in digital mammograms, *IEEE Trans. Med. Imag.* **18**(9) (1999) 774–786.

31. H. Yoshida, K. Doi and R. M. Nishikawa, Automated detection of clustered microcalcifications in digital mammograms using wavelet transform techniques, *Medical Imaging, SPIE* (1994) 868–886.

32. R. N. Strickland and H. L. Hahn, Wavelet transforms for detecting microcalcifications in mammograms, *IEEE Trans. on Medical Imaging* **15**(2) (1996) 218–229.
33. R. N. Strickland and H. L. Hahn, Wavelet transforms methods for object detection and recovery, *IEEE Trans. Image Processing* **6** (1997) 724–735.
34. M. N. Gurcan, Y. Yardimci, A. E. Cetin and R. Ansari, Detection of microcalcifications in mammograms using higher order statistics, *IEEE Signal Processing Letters* **4**(8) (1997) 213–216.
35. N. Karssemeijer, A stochastic model for automated detection of calcifications in digital mammograms, *12th International Conference IPMI*, Wye, UK, 1992, pp. 227–238. **8** (1997) 491–504.
36. W. J. H. Veldkamp and N. Karssemeijer, An improved method for detection of microcalcification clusters in digital mammograms, *The SPIE Conference on Image Processing* **3661** (1999) 512–522.
37. V. Cherkassky and F. Mulier, *Learning from Data* (John Wiley, 1998).
38. B. Ripley, *Pattern Recognition and Neural Networks* (Cambridge, 1996).
39. A. P. Dhawan, Y. Chitre and M. Moskowitz, Artificial neural network based classification of mammographic microcalcifications using image structure features, *Proc. SPIE* **1905** (1993) 820–831.
40. B. Verma and J. Zakos, A computer-aided diagnosis system for digital mammograms based on fuzzy-neural and feature extraction techniques, *IEEE Trans. Inform. Technol. Biomed.* **5**(1) (2001) 46–54.
41. Y. Jiang, R. M. Nishikawa, D. E. Wolverton, C. E. Metz, M. L Giger, R. A. Schmidt, C. J. Vyborny and K. Doi, Malignant and benign clustered microcalcifications: Automated feature analysis and classification, *Radiology* **198** (1996) 671–678.
42. S. Yu and L. Guan, A CAD system for the automatic detection of clustered microcalcifications in digitized mammogram films, *IEEE Trans. Med. Imag.* **19** (2000) 115–126.
43. V. Vapnik, *Statistical Learning Theory* (Wiley, New York, 1998).
44. A. Bazzani, A. Bevilacqua, D. Bollini, R. Brancaccio, R. Campanini, N. Lanconelli, A. Riccardi and D. Romani, An SVM classifier to separate false signals from microcalcifications in digital mammograms, *Phys. Med. Biol.* **46** (2001) 1651–1663.
45. I. El Naqa, Y. Yang, M. N. Wernick, N. P. Galatsanos and R. Nishikawa, A support vector machine for approach for detection of Microcalcifications, *IEEE Trans. Med. Imaging* **21**(12) (2002) 1552–1563.
46. H. Li, K. J. Liu and S. Lo, Fractal modeling and segmentation for the enhancement of microcalcifications in digital mammograms, *IEEE Trans. on Medical Imaging* **16**(6) (1997) 785–798.
47. H. D. Cheng, Y. M. Lui and R. I. Freimanis, A novel approach to microcalcification detection using fuzzy logic technique, *IEEE Trans. on Medical Imaging* **17**(3) (1998) 442–450.

CHAPTER 3

FUZZY REGION GROWING AND FUSION METHODS FOR THE SEGMENTATION OF MASSES IN MAMMOGRAMS

DENISE GULIATO

Laboratório de Computação Científica
Faculdade de Computação, Universidade Federal de Uberlândia
P.O. Box 593, Campus Santa Mônica, Uberlândia, Minas Gerais, Brazil
guliato@ufu.br

RANGARAJ M. RANGAYYAN

Department of Electrical and Computer Engineering and Department of Radiology
University of Calgary, 2500 University Drive N.W.
Calgary, Alberta T2N 1N4, Canada
ranga@enel.ucalgary.ca

The segmentation of regions of interest in medical images is, in general, a difficult problem. Methods proposed for a specific application are usually not applicable to images in other applications. In particular, defining criteria to determine precisely the boundaries of masses in mammograms is a difficult task. 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 less-dense tissues.

We present two segmentation methods that incorporate fuzzy concepts. The first method determines the boundary of a mass or tumor by region growing after a preprocessing step based on fuzzy sets to enhance the region of interest (ROI). Contours provided by the method have demonstrated good match with the contours drawn by a radiologist, as indicated by good agreement between the two sets of contours for 47 mammograms. The second segmentation method is a fuzzy region growing method that takes into account the uncertainty present around the boundaries of tumors. The difficult step of deciding upon a crisp boundary is obviated by the method. Measures of inhomogeneity computed from the pixels present in a suitably defined fuzzy ribbon have indicated potential use in classifying the masses and tumors as benign or malignant, with a sensitivity of 0.8 and a specificity of 0.9.

In view of the difficulty in the detection of masses and tumors in mammograms, we explore the combined use of multiple segmentation approaches. We present an abstract concept of information fusion based on a finite automaton and fuzzy sets to integrate and evaluate the results of multiple image segmentation procedures. We give examples on how the method can be applied to the problem of mammographic image segmentation, combining the results of region growing and closed-contour detection techniques. We also describe a measure of fuzziness to assess the agreement between a segmented region and a reference contour. Application of the fusion technique to breast tumor detection in mammograms indicates that the fusion results agree with the reference contours provided by a radiologist to a higher extent than the results of the individual methods.

Keywords: Fuzzy region growing; segmentation; fusion method; breast masses; mammograms.

1. Medical Image Segmentation

Although a radiologist typically begins an examination of a given medical image by scanning it in its entirety, diagnostic features of interest often manifest in local regions. It is not common for a disease to alter an image over its entire extent. In the context of a breast cancer screening program, the radiologist would scan the entire mammographic image(s) presented, and search for abnormal features that could be associated with disease; more attention would then be paid to specific regions of interest (ROIs). In a diagnostic application, the radiologist would concentrate on the region of suspected abnormality, and examine its characteristics to decide if the region exhibits signs related to a particular disease. In the situation of computer-aided diagnosis (CAD), an important role for image processing is to detect the ROIs for the specific application considered.[1,2] After the ROIs are detected, the subsequent steps could relate to objective characterization and classification of the regions. The following list provides a few examples of ROIs in medical images:

- Calcifications in mammograms.
- Lung nodules in X-ray images of the chest.
- Tumors and masses in mammograms.
- The white matter, the gray matter, the cerebro-spinal fluid, and the corpus callosum in magnetic resonance (MR) images of the head (brain).
- The liver and the kidneys in computed tomography (CT) images of the abdomen.

An image segmentation algorithm is a procedure that divides an image into its constituent parts or ROIs. Segmentation is an essential step before image analysis, pattern recognition, or diagnostic classification. Two major approaches to image segmentation are based on the detection or analysis of the following characteristics of images:

- Discontinuity — changes in gray level; edges.
- Similarity — homogeneity of the internal parts of an object.

Commonly used approaches to image segmentation are based upon edge detection, gray-level thresholding, region growing, and region splitting/merging. Depending upon the nature of the given image and the ROI, one could attempt to detect the edges of the ROI, or attempt to a grow region to span the ROI. In some cases, an ROI may include several disjoint components that may need to be matched and connected, or analyzed as a group. The results of methods as above could be labelled as "crisp" regions in the sense that a pixel either belongs to the region or does not belong to the region.

The abundance of reports on medical image segmentation available in the literature indicates that the problem is a difficult one: it is not possible to generalize the problem or any potential solution proposed for a specific application. More often than not, the ROI may not be clearly defined: even an expert radiologist may not be able to "segment" the ROI by cutting the part of interest from the image in

a definitive manner. However, in practice, a radiologist is not required to identify the precise boundaries of an ROI: it is adequate to identify the general area of the abnormality and give a diagnostic decision. This practical aspect suggests a path that is different from the traditional sequence of steps toward image analysis as described above: we may allow imprecision in the definition of the ROI — that is, allow it to be "fuzzy" — and arrive at the diagnostic decision using subsequent procedures that derive features by taking into account the lack of precise boundaries of the ROI. In this chapter, we shall explore methods based upon fuzzy sets and fuzzy fusion for the segmentation of masses and tumors in mammograms, and demonstrate how features may be derived from fuzzy regions to discriminate between benign masses and malignant tumors.

2. Breast Cancer and Mammography

Breast cancer is one of the leading causes of death from cancer among women. Early diagnosis can increase the chances of survival. Although mammography is being used for breast cancer screening,[3,4] analysis of masses and tumors on mammograms is, at times, difficult because developing signs of cancer may be minimal or masked by superimposed tissues, making their visual detection and analysis difficult. 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, 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, the growth of tumors into surrounding tissues, and distortion caused by developing densities.[5] Annotation of mammograms with objective measures may assist radiologists in diagnosis.[6]

Various segmentation algorithms have been presented in the literature to extract tumor regions from mammographic images (refer to recent papers by Rangayyan *et al.*[7,8] and Mudigonda *et al.*[9,10] for reviews on this topic). In general, the reported segmentation techniques attempt to define precisely an ROI, such as a tumor or a mass. However, it is difficult to define a criterion to obtain precise regions on mammograms. The problem is complicated by the fact that most malignant tumors possess fuzzy boundaries with slow and extended transition from a dense core region to the surrounding tissues. Very few works consider the uncertainty present around the ROI boundaries.[11,12]

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.*[13] proposed thresholding and multitolerance region growing methods for the detection of potential calcification regions and the extraction of their contours. Karssemeijer,[14] Laine *et al.*,[15] and Miller and Ramsey[16]

proposed methods for tumor detection based on scale-space analysis. Zhang *et al.*[17] proposed an automated detection method for the initial identification of spiculated lesions based on an analysis of mammographic texture patterns. Matsubara *et al.*[18] described an algorithm based on an adaptive thresholding technique for mass detection. Kupinski and Giger[19] 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. The difference among the methods lies in the utility function to determine the final lesion area. However, defining criteria to realize precisely the boundaries of masses in mammograms is difficult.

An alternative to address this problem is to represent tumor or mass regions by fuzzy sets.[20] The most popular algorithm that uses the fuzzy-set approach is the Fuzzy C-Means (FCM) algorithm.[12,21,22] The FCM 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[12] uses FCM 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[11] 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 using the same procedure. The authors presented results of application of the method to mammograms with four levels of segmentation.

In this chapter, we describe two segmentation methods that incorporate fuzzy concepts.[23–25] 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.[23] The method is simple and easy to implement, always produces closed contours, and has yielded good results even in the presence of high levels of noise. The second segmentation method is a fuzzy region growing method that takes into account the uncertainty present around the boundaries of tumors.[24] The method produces a fuzzy representation of the ROI, and preserves the uncertainty around the boundaries of tumors. We demonstrate the potential use of features derived from the results of segmentation in pattern classification of the regions as benign masses or malignant tumors.

Given the difficult nature of the problem of detection of masses and tumors in a mammogram, the question arises if the problem could benefit from the use of multiple approaches: How may we combine the results of several approaches — which may be considered to be complementary — so as to obtain a possibly better result?

In generic terms, image segmentation may be defined as a procedure that groups the pixels of an image according to one or more local properties.[26] 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 the discontinuity of local properties, and those based on the similarity of local properties.[27] 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 (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 results of segmentation procedures can offer good results, there are very few publications devoted to this subject.[28−34] 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 different image segmentation results.

Using the theory of fuzzy sets, it is possible to define several classes of fusion operators that generalize Boolean operators. We describe a general fusion operator, oriented by a finite automaton, to combine information from different sources.[35,36] In particular, we apply the idea to the problem of mammographic image segmentation for the detection of breast tumors, combining results obtained via contour detection and region growing. The final fuzzy set can classify pixels with more certainty, and preserve more information than either of the individual methods. The results are evaluated using a measure of agreement with reference to the contours of the tumors drawn independently by an expert radiologist specialized in mammography.

3. 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, pixels in and around the ROI will be displayed according to their degree of similarity with respect to the features of the mass ROI. We describe a method based on fuzzy set theory to achieve such a preprocessing step.

A fuzzy set can be defined by assigning to each possible element in the set X a value representing its grade of membership in the fuzzy set.[37,38] This grade corresponds to the degree with which the element is similar to or compatible with the concept represented by the fuzzy set. Let $\Gamma : X \rightarrow L$ be a membership function that maps X 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 $[0, 1]$; we will use this interval in the present work. Crisp sets can be seen as a particular case of fuzzy sets where $\Gamma : X \rightarrow \{0, 1\}$, i.e. the range includes only the discrete values 0 and 1.

The enhancement 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;
- is symmetric 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.

We may consider 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 Fig. 1, is given by the function $\Gamma(p) = \frac{1}{1+\beta|\mathbf{A}-\mathbf{B}|}$, where p is the pixel being processed, $\mathbf{A}$ is the feature vector of the mass (gray level in this work), $\mathbf{B}$ 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.

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

Fuzzy Membership Degree

Difference between the feature vector of the ROI
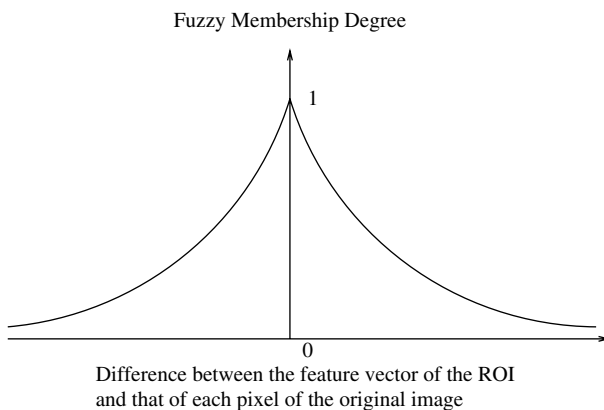and that of each pixel of the original image

Fig. 1. 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.

Fig. 2. Original version of a 700×700-pixel portion of a mammogram with a spiculated malignant tumor. Pixel size $= 62.5\,\mu m$. 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.
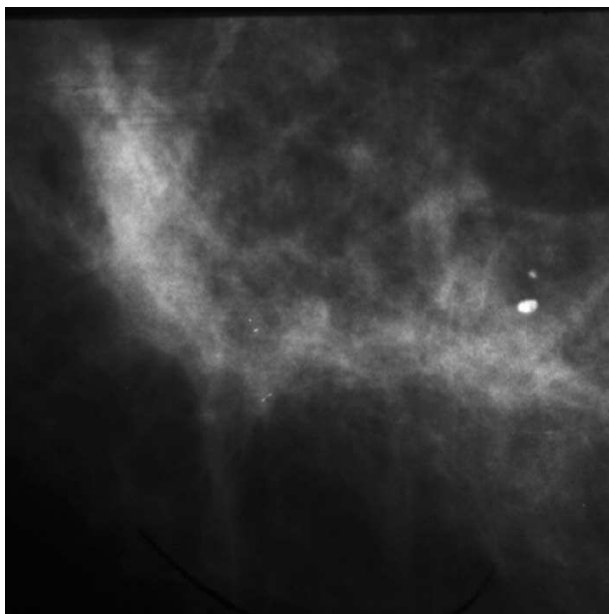
opposite case results in a low membership degree. The membership degree may be used as a scale factor to obtain gray levels and display the result as an image. The contrast of the ROI in the resulting image depends upon the parameter β.

Figures 2 and 3 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 Fig. 3 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 the 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 Sec. 4.1.

4. 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.[39] Statistical measures appear to be good tools for defining homogeneous regions. The success of image segmentation is directly associated with the choice of the measures and a suitable threshold. Various region growing algorithms based on statistical measures have
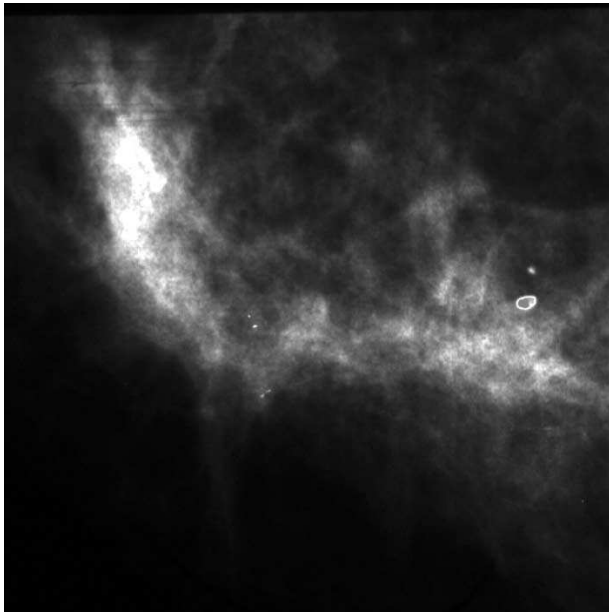
Fig. 3. Fuzzy-set-based ROI enhancement with $\beta = 0.007$ for the image in Fig. 2. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "Segmentation of breast tumors in mammograms using fuzzy sets", Journal of Electronic Imaging, 12(3): 369–378, 2003. © SPIE and IS&T.

been presented in the literature.[13,27,40,41] In particular, mean and standard deviation measures are often used as parameters to control region growing; however, these measures are highly influenced by extreme pixel values. As a consequence, the final shape of the region grown depends upon the strategy used to traverse the image. Furthermore, the algorithm could present unstable behavior, i.e. 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 scanning or 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 the following sections, we describe two image segmentation methods: the first based on classical region growing with the fuzzy-set preprocessed image, and the second based on fuzzy region growing using statistical measures in homogeneity criteria. The problems discussed so far are overcome by both methods.

4.1. *Detection of mass boundaries by classical region growing with fuzzy-set preprocessed images*

The boundary of the mass ROI may be obtained by performing region growing upon the fuzzy-set preprocessed image. The pixel values in the preprocessed image

represent 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 (mass).

The region growing process starts with the seed region. Four-connected neighboring pixels that are above the threshold are labelled as zero, the neighbors of the pixels labelled as zero are inspected, and the procedure continues. If the connected pixel is less than the threshold, it is labelled as one, indicating a contour pixel, and its neighborhood is not processed. The recursive procedure continues until all 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 has been evaluated with a number of synthetic test images as well as medical images such as CT and nuclear medicine images, and has produced good results even in the presence of high levels of noise. Figure 4 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

(a) (b)

(c) (d)
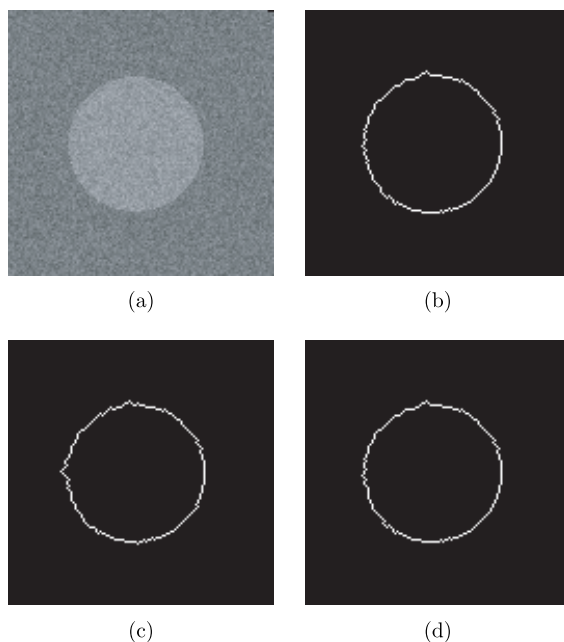
Fig. 4. Illustration of the effects of seed pixel and threshold selection on fuzzy-set preprocessing and region growing: (a) Original image (128×128 pixels) with additive Gaussian noise, with $\sigma = 12$ and signal-to-noise ratio (SNR) = 2.66; (b) Seed pixel (60, 60) and threshold = 0.82; (c) Seed 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.

Fig. 5. Contour extraction by fuzzy-set preprocessing and region growing for the mammogram in Fig. 2 with a spiculated malignant tumor; see also Fig. 3. The black line represents the boundary drawn by a r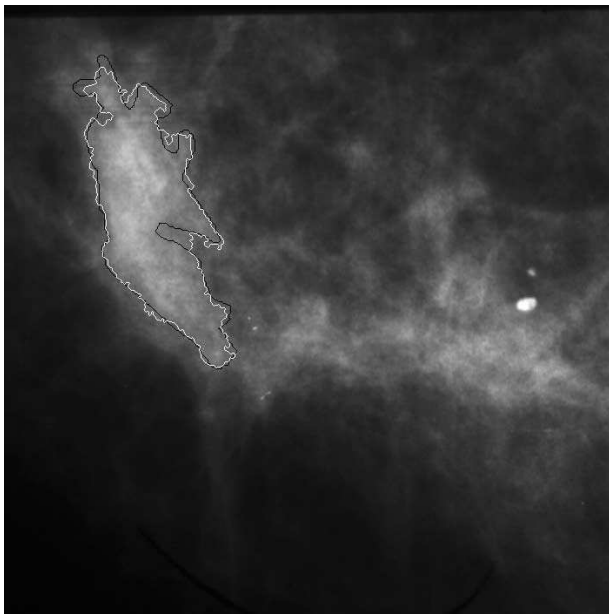adiologist (shown for comparison). The white line is the contour extracted by the region growing method. $\beta = 0.007$, threshold $= 0.63$. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "Segmentation of breast tumors in mammograms using fuzzy sets", Journal of Electronic Imaging, 12(3): 369–378, 2003. © SPIE and IS&T.

robustness of the method in the presence of noise. (Each result of segmentation shown in this chapter was obtained by using a single seed pixel.)

Figure 5 shows the contour extracted for the mammogram in Fig. 2 with a spiculated malignant tumor. Figure 6 shows a part of mammogram with a circumscribed benign mass. Figure 7 shows the result of enhancement of the mass ROI using fuzzy-set-based preprocessing, and Fig. 8 shows the corresponding contour obtained. The images with the results of segmentation are superimposed with the contours given by the region growing method in white; the contour in black is the boundary drawn independently by an experienced radiologist, shown for comparison.[7]

4.2. *Fuzzy region growing*

We now describe a fuzzy region growing algorithm to obtain mass regions in mammograms. 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 pixel or region that lies inside the ROI and spreads by adding to the region 8-connected pixels that have similar properties.
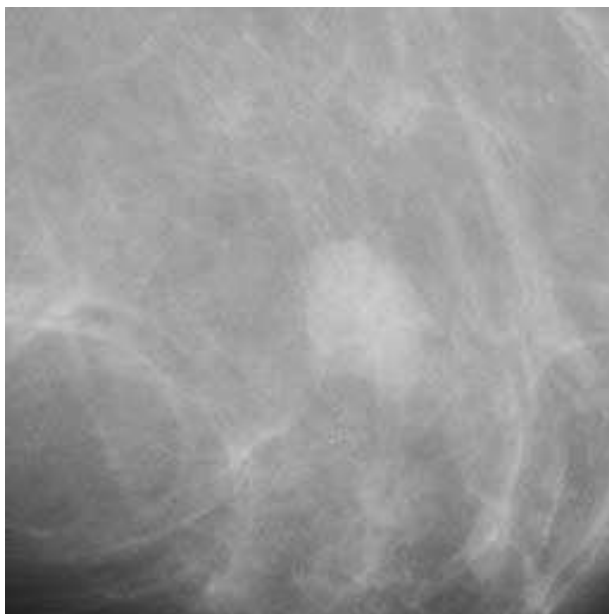
Fig. 6. Original version of a 1,024 × 1,024-pixel portion of a mammogram with a benign mass. Pixel size = 50 μm. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "Segmentation of breast tumors in mammograms using fuzzy sets", Journal of Electronic Imaging, 12(3): 369–378, 2003. © SPIE and IS&T.
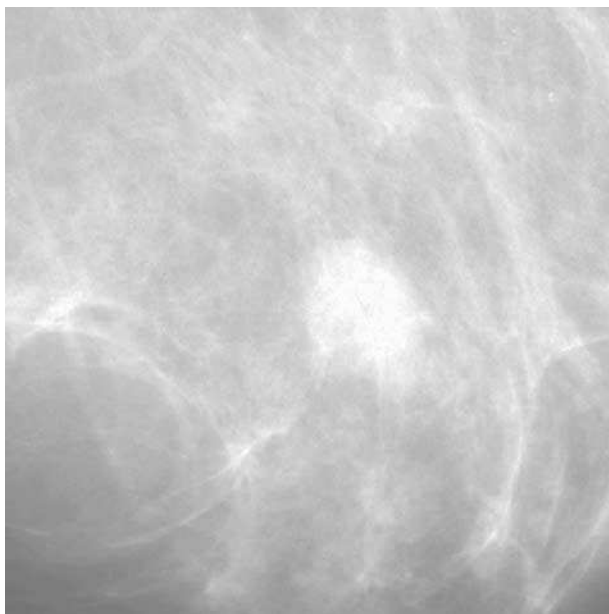
Fig. 7. Fuzzy-set-based ROI enhancement with $\beta = 0.007$ for the image in Fig. 6.

Fig. 8.   The white line shows the contour extracted by the closed-contour detection method for the mammogram in Fig. 6. The black line represents the boundary drawn by a radiologist (shown 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.

The homogeneity of the region is evaluated by calculating the mean (μ), standard deviation (σ), and the coeff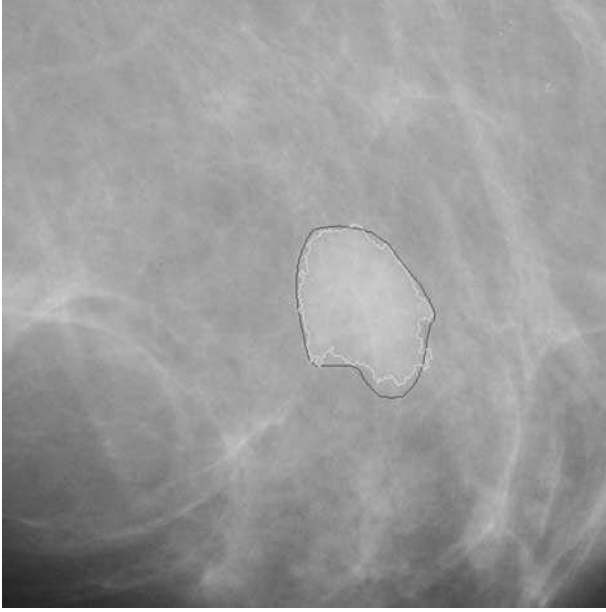icient of variation $CV = \sigma/\mu$. CV gives a measure of inhomogeneity of the region, and allows one to compare different regions.

Let $\Delta\mu_{\max}$, $\Delta CV_{\max}$, 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. $\Delta CV_{\max}$ 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. Let μ and σ be the mean and the standard deviation of the region already grown. The segmentation algorithm is executed in two steps:

(1) **Comparison 1**: $|I(p) - \mu| \leq \Delta\mu_{\max}$
   If this condition is not satisfied, then the pixel is labelled as *rejected*. If the condition is satisfied, p is temporarily added to the subregion and μ_{new} and σ_{new} are calculated.
(2) **Comparison 2**: $|\sigma/\mu - \sigma_{\text{new}}/\mu_{\text{new}}| \leq \Delta CV_{\max}$
   If the condition is satisfied, then p must definitely be added to the subregion labelled as *accepted*, and μ and σ must be updated, i.e. $\mu = \mu_{\text{new}}$ 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.

Comparison 2 given above analyzes the distortion (in terms of CV) that the pixel p can produce if added to the subregion. At the beginning of the procedure, 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 special procedure is executed for Comparison 2: $|\sigma/\mu - \sigma_{\text{new}}/\mu_{\text{new}}| \leq 2\Delta CV_{\text{max}}$. 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 procedure. 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.

We now define a fuzzy membership function that maps the pixel values of the region resulting from the procedure described above to the unit interval $[0, 1]$ 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 seed region must be 1;
- the membership degree of a pixel labelled 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 labelled 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 in the present work is illustrated in Fig. 9, where $a = |mean\_seed\_region + standard\_deviation\_seed\_region - \mu|$ and $b = \Delta\mu_{\text{max}}$; in this expression, μ is the mean of the region obtained by the preceding region growing procedure. 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 $\Gamma(p) = 0$
    else $\Gamma(p) = 1/\{1 + \beta[|I(p) - \mu|]\}$.

The method has been tested on several synthetic images with various levels of noise. Figure 10 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.

The fuzzy region provided by the method for the malignant tumor case shown in Fig. 2 is illustrated in Fig. 11. Figure 12 shows the fuzzy region obtained for the
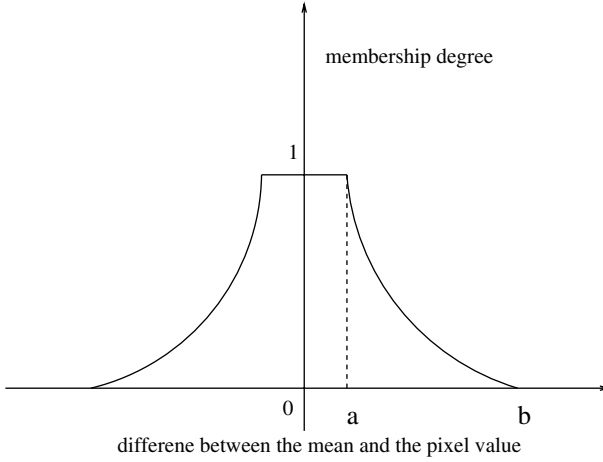
Fig. 9. Fuzzy membership function for region growing, where $a = |mean\_seed\_region + standard\_deviation\_seed\_region - \mu|$, and $b = \Delta\mu_{max}$. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "Segmentation of breast tumors in mammograms using fuzzy sets", Journal of Electronic Imaging, 12(3): 369–378, 2003. © SPIE and IS&T.

(a)   (b)

(c)   (d)
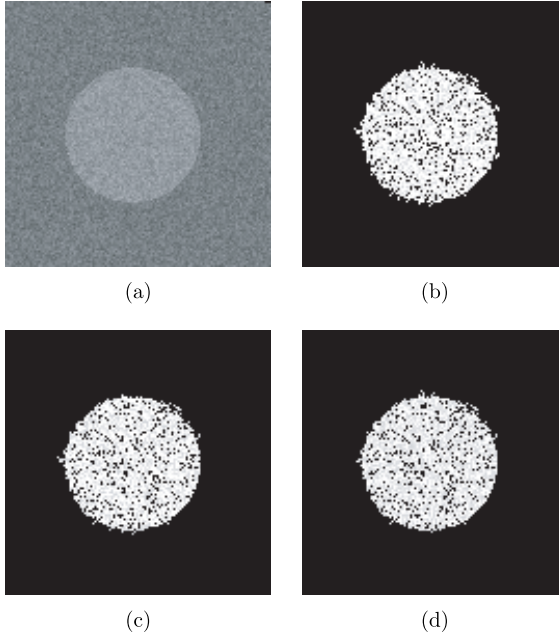
Fig. 10. Illustration of the effects of seed pixel selection on fuzzy region growing: (a) Original image (128×128 pixels) with Gaussian noise, with $\sigma = 12$ and SNR = 2.66; (b) Seed pixel (60, 60), $\Delta\mu_{max} = 18$, $\Delta CV_{max} = 0.007$, $\beta = 0.01$; (c) Seed pixel (68, 60), $\Delta\mu_{max} = 18$, $\Delta CV_{max} = 0.007$, $\beta = 0.01$; (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.

Fig. 11. Result of fuzzy region growing for the tumor displayed in Fig. 2, with $\Delta\mu_{\max} = 45$, $\Delta CV_{\max} = 0.01$, $\beta = 0.07$. 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.

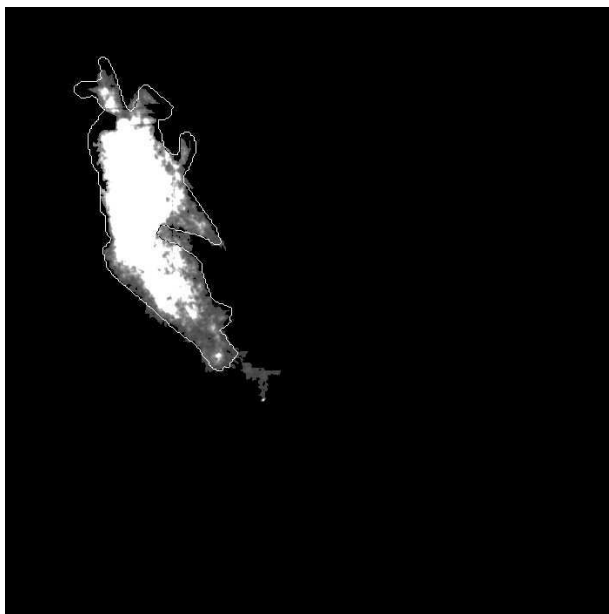benign mass case shown in Fig. 6. Good agreement between the fuzzy regions and the contours drawn by the radiologist (also shown in Figs. 11 and 12 for comparison) was obtained for most of the cases tested.

5. Classification of Masses Using the Results of Segmentation

An interactive graphical interface has been developed in Java, using an object-oriented architecture with controller classes, to integrate the programs that apply the segmentation methods, implemented in the C language, to mammographic images.[42] Some of the features of the interface are fast and easy upgradability, portability, and threads to support parallelism between tasks. The interface is ergonomic, easy to use, and includes online help. 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 or revise contours, and image format conversion.

The following subsections provide the results of classification obtained with 47 mammograms, including circumscribed/spiculated and benign/malignant
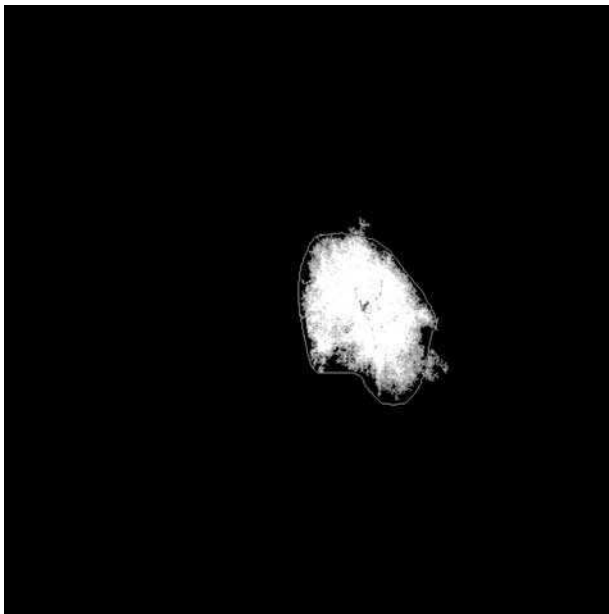
Fig. 12.   Result of fuzzy region growing for the benign mass displayed in Fig. 6. The contour drawn by the radiologist is superimposed for comparison. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "Segmentation of breast tumors in mammograms using fuzzy sets", Journal of Electronic Imaging, 12(3): 369–378, 2003. © SPIE and IS&T.

masses from the Mammographic Image Analysis Society (MIAS, UK) database (http://www.wiau.man.ac.uk/services/MIAS/MIAScom.html) and cases from the Foothills Hospital, Calgary. The dataset includes 22 benign masses and 25 malignant tumors. All diagnoses were proven by pathological examination of resected tissue, i.e. biopsy. The mammograms were digitized to a spatial resolution of $50\,\mu m$ in the case of the MIAS images and $62.5\,\mu m$ in the case of the Foothills Hospital images, and quantized to the range $0 - 255$ in gray scale. Contours of the masses were drawn by a radiologist on digital images independent of the present study, with no additional information made available.[7] The contours drawn by the radiologist were used for reference (as the "gold standard") to evaluate the results of segmentation.

5.1. *Results of classical region growing with fuzzy-set preprocessed images*

For the study described in this chapter, we manually selected the seed pixel and threshold value for each case; the threshold 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 the images in the study. In most of the cases analyzed, the boundaries obtained closely matched those drawn by the radiologist.
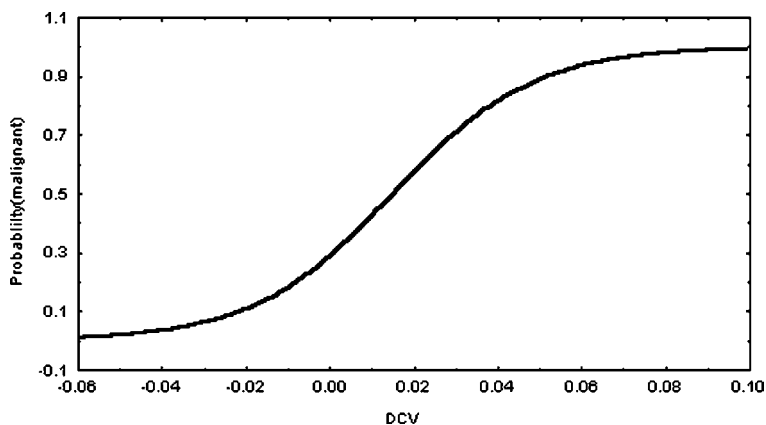
Fig. 13. The probability of malignancy derived from the parameter DCV. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "Segmentation of breast tumors in mammograms using fuzzy sets", Journal of Electronic Imaging, 12(3): 369–378, 2003. © SPIE and IS&T.

In order to derive a parameter for discriminating between benign masses and malignant tumors, the following procedure was applied.[42] 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 labelled as DCV was computed from the fuzzy-set preprocessed image, by taking the difference between the coefficient of variation 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; the result is illustrated in Fig. 13. 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.

The result of segmentation depends upon the choice of the seed to start region growing and the threshold. Automatic selection of the seed pixel or seed region is a difficult problem that requires further work. The threshold could be derived as a function of the statistics (mean, standard deviation, etc.) of the fuzzy-set preprocessed image; we are conducting further work to explore this possibility.

5.2. *Results of fuzzy region growing*

The fuzzy region growing method was applied to the 47 test images maintaining the same values of $\beta = 0.07$ and $\Delta CV_{max} = 0.01$, and by 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 the radiologist. The $\Delta \mu_{max}$ parameter ranged from 5 to 48 for the 47 masses and tumors analyzed.

In order to derive parameters for pattern classification, we 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 Fig. 14. Shape factors of mass contours and measures of edge sharpness and texture have been previously proposed for the purpose of classification.[7–10,43] However, important information is lost in analysis based on crisply defined contours: the uncertainty present in and/or around the ROI is not considered. In this work, we evaluate 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 tumor in Fig. 14 is darker than that of the benign mass in Fig. 15. 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, we computed the coefficient of variation CV_{fr} of the membership values of the pixels lying only within the fuzzy ribbon, and the ratio η of the number of pixels with membership degree less than 0.5 to the total number of pixels within the fuzzy ribbon. It is expected that the fuzzy ribbons of malignant tumors will possess higher CV_{fr} and η than those of benign masses.

Discrimination between benign masses and malignant tumors with the parameter η had no statistical significance. The probability of malignancy curve based
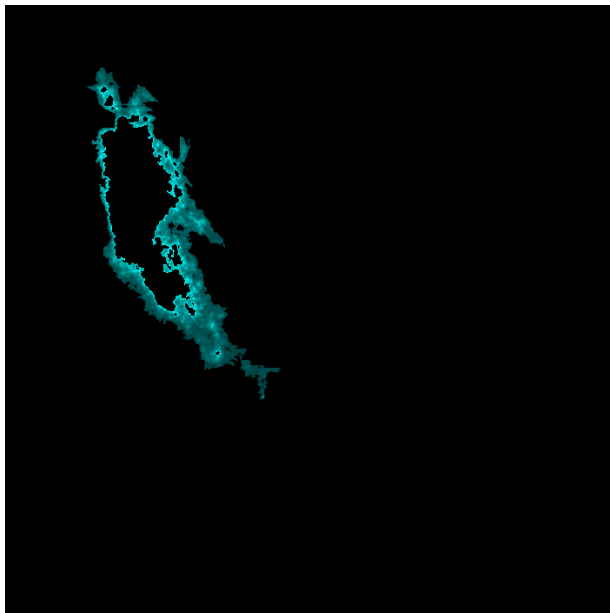
Fig. 14. The fuzzy ribbon of the fuzzy region shown in Fig. 11, for the malignant tumor shown in Fig. 2. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "Segmentation of breast tumors in mammograms using fuzzy sets", Journal of Electronic Imaging, 12(3): 369–378, 2003. © SPIE and IS&T.
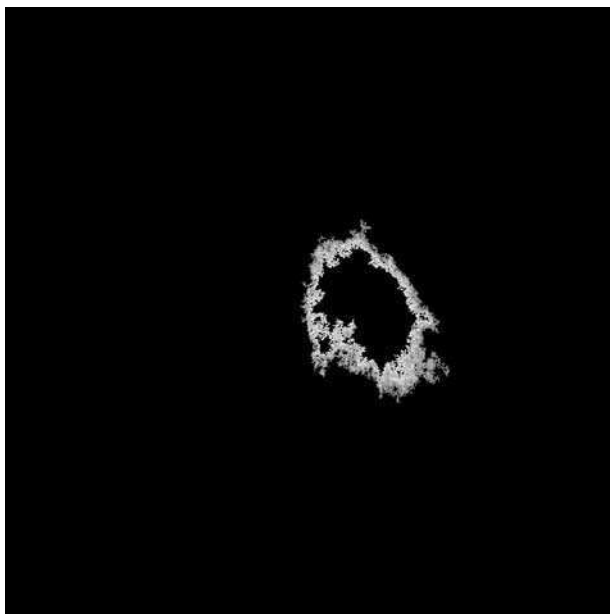
Fig. 15. The fuzzy ribbon of the fuzzy region shown in Fig. 12 for the benign mass in Fig. 6. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "Segmentation of breast tumors in mammograms using fuzzy sets", Journal of Electronic Imaging, 12(3): 369–378, 2003. © SPIE and IS&T.
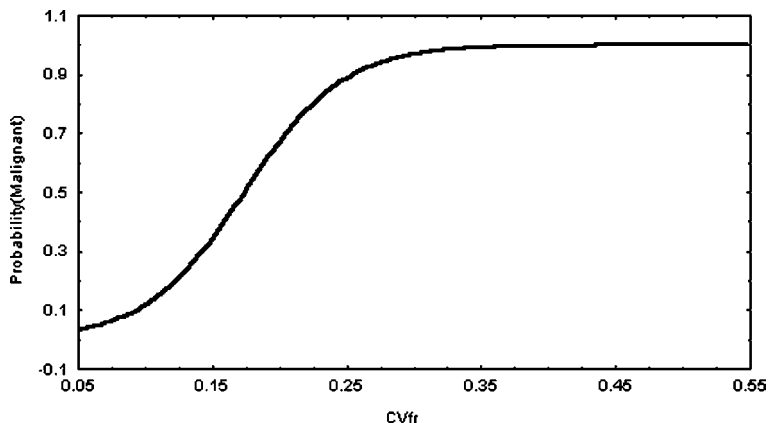
Fig. 16. The probability of malignancy derived from the parameter CV_{fr}. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "Segmentation of breast tumors in mammograms using fuzzy sets", Journal of Electronic Imaging, 12(3): 369–378, 2003. © SPIE and IS&T.

upon CV_{fr} is illustrated in Fig. 16. The cut point of 0.18 resulted in the correct classification of 20 out of the 25 malignant tumors and 20 out of the 22 benign masses processed, leading to a sensitivity of 0.8 and a specificity of 0.9.

The difficult step of deciding upon a crisp boundary has been obviated in the region growing method by using measures of inhomogeneity computed from the pixels present in a suitably defined fuzzy ribbon. The method is thus capable of assisting in not only identifying a mass ROI, but also in classifying the region as benign or malignant.

6. Elementary Concepts of Fusion Operators

A *fusion operator* over fuzzy sets is formally defined as a function $h : [0, 1]^n \to [0, 1]$, 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,[38,44] as follows:

- An operator is said to be *conjunctive* if $h(a_1, a_2, \ldots, a_n) \leq \min\{a_1, a_2, \ldots, a_n\}$, 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 $h(a_1, a_2, \ldots, a_n) \geq \max\{a_1, a_2, \ldots, a_n\}$. 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 $\min\{a_1, a_2, \ldots, a_n\} \leq h(a_1, a_2, \ldots, a_n) \leq \max\{a_1, a_2, \ldots, a_n\}$. Compromise operators are those that produce an intermediate measure between items of information obtained from several sources. They present *cautious* behavior.

Bloch[45] presents 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 (CICB) operators that maintain the same behavior independent of the input variables; Context-Independent Variable-Behavior (CIVB) operators, whose behavior varies according to the input variables; and Context-Dependent (CD) operators, whose behavior varies as in the previous case, taking into account also the agreement between and the reliability of the sources.

We describe a class of fusion operators that generalize CD operators, taking into consideration different degrees of confidence in the sources, specific knowledge, and spatial context while operating with conceptually distinct sources.

7. Considerations in the Fusion of the Results of Complementary Segmentation Techniques

Figure 17 illustrates a schematic 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 two results of segmentation are not the same: different segmentation algorithms may produce different results for the same ROI. A fusion operator designed to aggregate these entities should produce a third entity that takes into consideration the inputs and is better than either input on its own. To realize this, the fusion operator must be able to identify regions of certainty and uncertainty during its execution.

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 with high certainty, or is low when p does not belong to the region with high certainty; a value near 0.5 indicates a high level of uncertainty of membership. With respect to the fusion operator, four situations may be identified considering the position of p (see Fig. 18):

(1) p belongs to the intersection of S_r and S_c [i.e. $\Gamma_{S_r}(p)$ 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 [i.e. $\Gamma_{S_r}(p)$ 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.
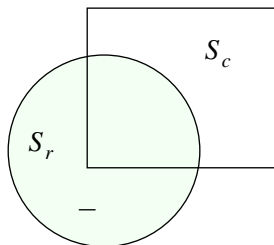
Fig. 17. Superimposition of the results of two complementary segmentation techniques. The circular region S_r represents the result of region growing. The square box S_c represents the result of contour detection. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, 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.
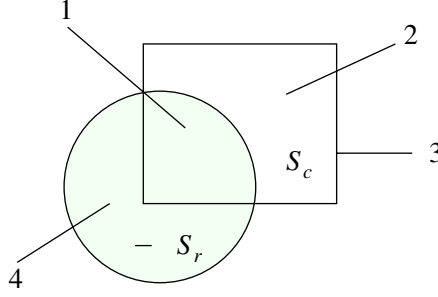
Fig. 18.   The four different situations treated by the fusion operator. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "Fuzzy fusion operators to combine results of complementary medical image segmentation techniques", Journal of Electronic Imaging, 12(3): 379–389, 2003. © SPIE and IS&T.

(3) p belongs to the contour line of S_c [i.e. $\Gamma_{S_c}(p)$ is high] and does not belong to S_r [i.e. $\Gamma_{S_r}(p)$ is low or zero]. As in Item 2 above, this is an uncertainty situation. However, although the inputs are different from those presented in Item 2, the result of the fusion operator is expected to represent uncertainty.

(4) p belongs to S_r [i.e. $\Gamma_{S_r}(p)$ is high] and is outside of S_c [i.e. $\Gamma_{S_c}(p)$ is zero]. Here again we have an uncertainty case. Observe that although the inputs are similar to those in Item 1 above [i.e. $\Gamma_{S_r}(p)$ 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 on these observations, we present 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 in Sec. 8) 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. Whereas parameters are included in the definitions of the operators to represent the reliability of the sources, it is assumed that the parameters will be determined using other methods.

8. Fusion Operators Based on Finite Automata

We define now a general notion of fusion operators based on finite automata that are inductively defined from a finite set of *basic 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[38]:

(1) $H = \{h_1, h_2, \ldots, h_k\}$ is a finite set of *basic fusion operators*, where h_i are functions that map $[0,1]^n \to [0,1]$, $n \geq 2$.
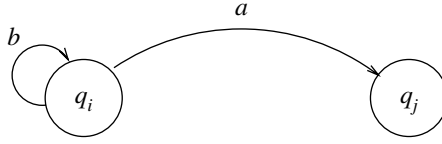
Fig. 19. Graphical representation of the transition function given by $\delta(q_i, a) = q_j$ and $\delta(q_i, b) = q_i$. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "Fuzzy fusion operators to combine results of complementary medical image segmentation techniques", Journal of Electronic Imaging, 12(3): 379–389, 2003. © SPIE and IS&T.

(2) $M = (Q, \Sigma, \delta, q_0, F)$ is a finite automaton, where:

    (a) Q is a finite set of states,

    (b) Σ is a finite input alphabet,

    (c) δ is a transition function that maps $Q \times \Sigma \to Q$, where $\times$ is the Cartesian product operator,

    (d) $q_0 \in Q$ is an initial state, and

    (e) $F \subset Q$ is the set of final states.

In our particular 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 as $\Sigma = \{0, 1, 2\}$ 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 iff 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. 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 Fig. 19.

9. Application of the Fusion Operator to Image Segmentation

We now present a specific application of the fusion operator to the problem of image segmentation. The fusion operator is designed to combine the results obtained from two segmentation techniques that explore complementary characteristics of the image: in our case, one is based on region growing (described in Sec. 4.2), and the other is based on closed-contour detection (described in Sec. 4.1).[23-25] The result of fusion obtained is a fuzzy set that represents the agreement or disagreement between the input sources.[35,36]

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 Fig. 17. The

process starts with a seed pixel selected by the user. The seed pixel must belong to the intersection $S_r \cap S_c$. We suppose 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 given by $OP = \langle H, M \rangle$, where $H = \{h_1, h_2, \ldots, h_6\}$ 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.

9.1. *Description of the basic fusion operators for image segmentation*

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 is the influence of the source. Methods need to be developed to derive automatically suitable values for these parameters.

The result of each basic fusion operator should give us information about the agreement among the sources being analyzed. The absence of conflict is represented by a membership degree equal to 1 or 0, i.e. 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 jth pixel of the segmented image S_i and $\Gamma_{S_i}(p_{ij})$ be the membership degree of the pixel p_{ij}, where $i \in \{r, c\}$, $j = 1, 2, \ldots, m$, and m is the total number of pixels in the image S_i. (Note that 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) $h_1 = \max\{\Gamma_{S_r}(p_{rj}) * C_r, C_c, 0.5\}$
   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 Fig. 20).

(2) if $\max(C_r, C_c) \leq 0.5$
   then $h_2 = 0.5$
   else if $(C_r \leq 0.5)$
       then $h_2 = \Gamma_{S_c}(p_{cj}) * C_c$
       else if $(C_c < 0.5)$
           then $h_2 = \Gamma_{S_r}(p_{rj}) * C_r$
           else $h_2 = 1/(C_r + C_c) * \{\Gamma_{S_r}(p_{rj}) * C_r + \Gamma_{S_c}(p_{cj}) * C_c\}$
   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 Fig. 20).
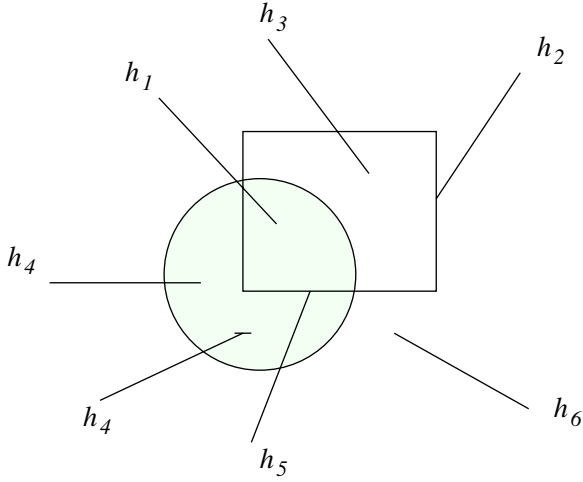
Fig. 20. The regions where the six basic fusion operators are applied are indicated by $\{h_1, h_2, h_3, h_4, h_5, h_6\}$. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "Fuzzy fusion operators to combine results of complementary medical image segmentation techniques", Journal of Electronic Imaging, 12(3): 379–389, 2003. © SPIE and IS&T.

(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 $h_3 = \Gamma_{S_r}(p_{rj}) * C_r$
           else $h_3 = 1/(C_r + C_c) * \{\Gamma_{S_r}(p_{rj}) * C_r + C_c\}$
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 Fig. 20).

(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 $h_4 = \Gamma_{S_r}(p_{rj}) * C_r$
           else $h_4 = 1/(C_r + C_c) * \{\Gamma_{S_r}(p_{rj}) * C_r + [1 - C_c^{0.5}]^2\}$
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 Fig. 20). Artifacts within the region growing result (as indicated schematically by the line segment inside the circle in Fig. 20) are rejected by this operator.

(5) $h_5 = \max\{\Gamma_{S_r}(p_{rj}) * C_r, \Gamma_{S_c}(p_{cj}) * C_c, 0.5\}$

This is a disjunctive operator that acts on the transition pixels lying in the intersection $S_r \cap S_c$ (see h_5 in Fig. 20).

(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 $h_6 = \Gamma_{S_r}(p_{rj}) * C_r$

  else $h_6 = \min\{\Gamma_{S_r}(p_{rj}) * C_r, [1 - C_c]\}$

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 Fig. 20).

9.2. *Description of the finite automaton for fuzzy fusion*

The finite automaton $M = (Q, \Sigma, \delta, q_0, F)$ in OP is defined by:

(1) $Q = \{a, b, c\}$ is a set of finite states, where

(a) state a indicates that the pixel being analyzed belongs to the interior of the contour,

(b) state b indicates that the pixel being analyzed belongs to the contour,

(c) state c indicates that the pixel being analyzed belongs to the exterior of the contour; see Fig. 21.
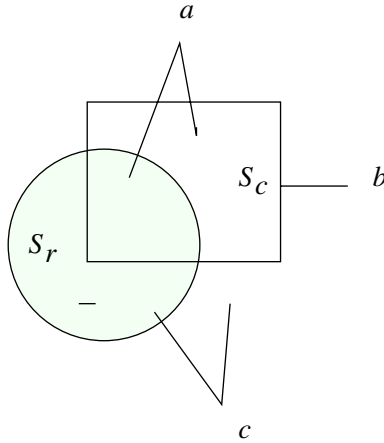

Fig. 21. The three states of the automaton. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "Fuzzy fusion operators to combine results of complementary medical image segmentation techniques", Journal of Electronic Imaging, 12(3): 379–389, 2003. © SPIE and IS&T.

(2) $\Sigma = \{I_1, I_2, I_3, I_4\}$ is a finite input alphabet.

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} is in the interval *high* if $\Gamma_{S_i}(p_{ij}) \geq 0.5$ for $j = 1, 2, \ldots, m$, and

p_{ij} is in the interval *low* if $\Gamma_{S_i}(p_{ij}) < 0.5$ for $j = 1, 2, \ldots, m$,

where p_{ij}, $i \in \{r, c\}$, and $j = 1, 2, \ldots, m$, identify the ith source and the jth 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 detection result. This input, as discussed in Sec. 7, represents a certainty or uncertainty situation depending on the spatial position of the pixel being analyzed; see Fig. 22.
- $\mu(high, high) = I_2$: the pixel being analyzed presents a high membership degree in the region-growing segmentation result and a high membership degree in the closed-contour detection result. This indicates an intersection case; see Fig. 22.
- $\mu(low, high) = I_3$: the pixel being analyzed presents a low membership degree in the region-growing segmentation result and a high membership degree in the closed-contour detection result. This indicates an uncertainty case; see Fig. 22.
- $\mu(low, low) = I_4$: the pixel being analyzed presents a low membership degree in the region-growing segmentation result and a low membership degree in the
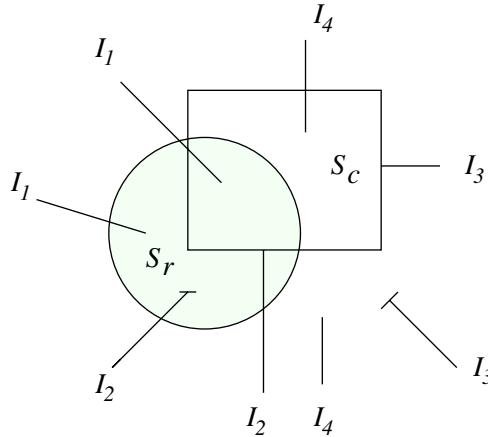
Fig. 22. The four possible input values $\{I_1, I_2, I_3, I_4\}$ for the fusion operator. The short line segments with the labels I_2 and I_3 represent artifacts in the segmentation result. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "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.
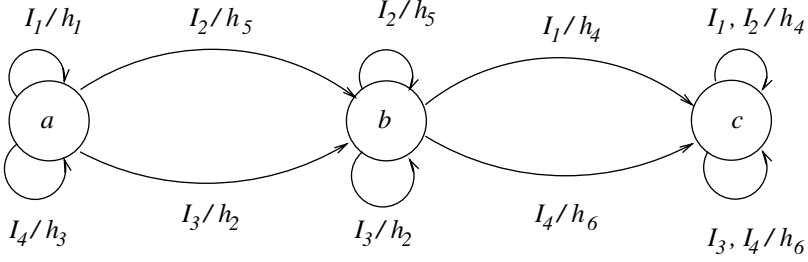
Fig. 23. Transition diagram that governs the actions of the fusion operator. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "Fuzzy fusion operators to combine results of complementary medical image segmentation techniques", Journal of Electronic Imaging, 12(3): 379–389, 2003. © SPIE and IS&T.

closed-contour detection 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 Fig. 22.

(3) The transition diagram δ of M is shown in Fig. 23.

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 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 operator h_4 or $h_6$ is 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 Fig. 23, 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.

(4) $q_0 \in Q$, with $q_0 = \{a\}$ as the initial state.

(5) $F = \{c\}$ is the set of final states, where $F \subseteq Q$. In the present case, F has only one element.
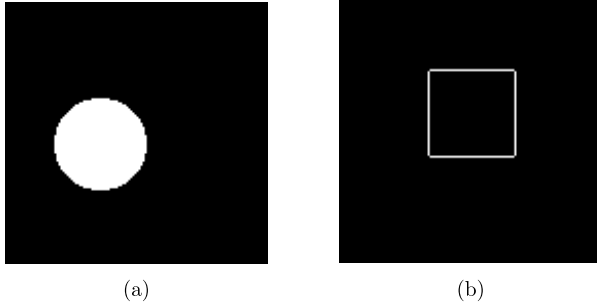
(a) (b)

Fig. 24. Illustration of the region and contour segmentation results to be aggregated: (a) Classical region segmentation result. (b) Classical contour segmentation result.

9.3. *Behavior of the basic fusion operators*

The fusion operator oriented by finite automaton described here can combine several results of segmentation, two at time. The result yielded by the fusion operator is a fuzzy set that identifies the certainty and uncertainty present in and among the inputs to the fusion process. It is expected that maximal certainty will be represented by a membership degree equal to 1 or 0 (i.e. the pixel being analyzed certainly belongs to, or does not belong to, the final segmentation result). When the individual segmentation results disagree with respect to a pixel belonging or not belonging to the final result, or when both the 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 intervals (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.

We present below illustrative studies on the behavior of the basic fusion operators described in Sec. 9.1, taking into consideration a limited set of entries. We consider both region and contour segmentation results represented by the classical Boolean set, i.e. the membership degree of the pixels inside of the region or belonging to the contour segmentation result is 1, and 0 otherwise, as shown in Fig. 24.

The starting condition of the fusion operator based on finite automaton is the state a and the entry 1. The starting pixel must lie in the intersection of S_r and S_c (see Figs. 20, 21, and 23). The basic fusion operator to start the procedure is $h_1 = \max\{\Gamma_{S_r}(p_{rj}) * C_r, C_c, 0.5\}$. Table 1 provides explanatory commentaries describing the behavior of h_1 for several values of the reliability parameters and inputs from the two sources. Table 2 presents the behavior of the fuzzy fusion operator including all of the basic operators h_1, h_2, h_3, h_4, h_5, and h_6, according to a given reliability. Figure 25 illustrates the fuzzy fusion results after the application of the fuzzy fusion operator to each situation presented in Table 2.

Table 1. Behavior of the basic fusion operator h_1. 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.

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
0.8	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
0.2	1.0	0.8	0.0	0.8	result depends on the source with the higher reliability

Table 2. Behavior of the basic fusion operators.

C_r	C_c	h_1	h_2	h_3	h_4	h_5	h_6
1.0	1.0	1.0	0.5	0.5	0.5	1.0	0.0
1.0	0.0	1.0	0.0	0.0	1.0	1.0	0.0
0.0	1.0	1.0	1.0	1.0	0.0	1.0	0.0
0.8	1.0	1.0	0.56	0.56	0.44	1.0	0.0
0.8	0.8	0.8	0.5	0.5	0.5	0.8	0.0
0.2	0.8	0.8	0.8	0.8	0.0	0.8	0.0

10. Application of Fuzzy Fusion to the Segmentation of Breast Masses

The result of fusion of the contour and the region obtained for the malignant tumor shown in Fig. 2, given in Figs. 5 and 11, respectively, is provided in Fig. 26. Figure 27 shows the result of fusion for the benign mass displayed in Fig. 6, with the corresponding contour and fuzzy region shown in Figs. 8 and 12, respectively. The results have been superimposed with the corresponding contours drawn by an expert radiologist for comparison. Contours of tumors were drawn by the radiologist on digital images, independent of the present study, with no additional information made available.[7]

Figure 28 shows a 700×700-pixel portion of another mammogram with a spiculated malignant tumor. Figures 29 and 30 show the detected contour and the fuzzy region that preserves the transition information around the boundary, respectively. Figure 31 shows the result of fusion of the ROIs represented in Figs. 29 and 30.

It should be observed that the fusion results reduce the uncertainty present in the interior of the regions, and also reduce the certainty of the boundaries. The features of the results of the individual segmentation procedures have contributed 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.
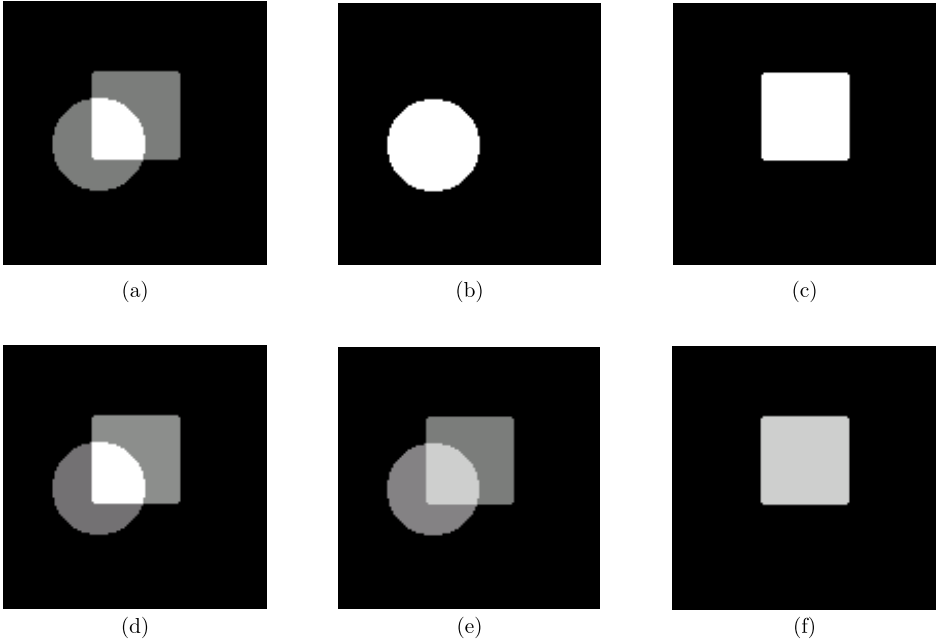
Fig. 25. Illustration of the effects of applying the fuzzy fusion operator, according to Table 2, to the schematic results of segmentation represented in Fig. 24: (a) $C_r = 1.0$, $C_c = 1.0$; (b) $C_r = 1.0$, $C_c = 0.0$; (c) $C_r = 0.0$, $C_c = 1.0$; (d) $C_r = 0.8$, $C_c = 1.0$; (e) $C_r = 0.8$, $C_c = 0.8$; (f) $C_r = 0.2$, $C_c = 0.8$.

11. Evaluation of the Results of Segmentation and Fusion Using a Measure of Fuzzyness

In order to evaluate the results of the fusion operator, we compare 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 their fusion. We aggregate the reference contour and a segmentation result using the fusion operator concept presented in Sec. 9. (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.) As discussed earlier, the fusion operator yields a fuzzy set which 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. Thus, if we quantify the uncertainty presented by the fusion result, we can evaluate the

Fig. 26.   Result of the fusion operator applied to the segmented regions in Figs. 5 and 11 related to the malignant tumor shown in Fig. 2, with $C_r = 1.0$, $C_c = 1.0$. 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.

degree of agreement among two different information sources. In order to quantify the uncertainty, we present a measure of fuzzyness.

11.1. *Measure of fuzzyness*

In general, a *measure of fuzzyness* is a function

$$f : \mathcal{F}(X) \to R^+, \tag{1}$$

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 fuzzyness of A. The function f must satisfy the following three requirements:

- $f(A) = 0$ iff A is a crisp set;
- $f(A)$ assumes its maximal value iff A is maximally fuzzy, i.e. all of the elements of A are equal to 0.5; and
- if set A is undoubtedly sharper (i.e. has more certainty) than set B, then $f(A) \leq f(B)$.

There are different ways of measuring fuzzyness that satisfy all of the three essential requirements.[38] We have chosen to measure fuzzyness in terms of the
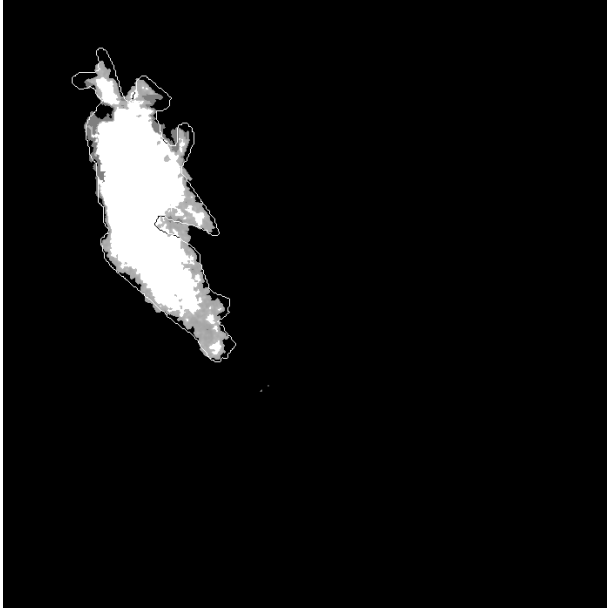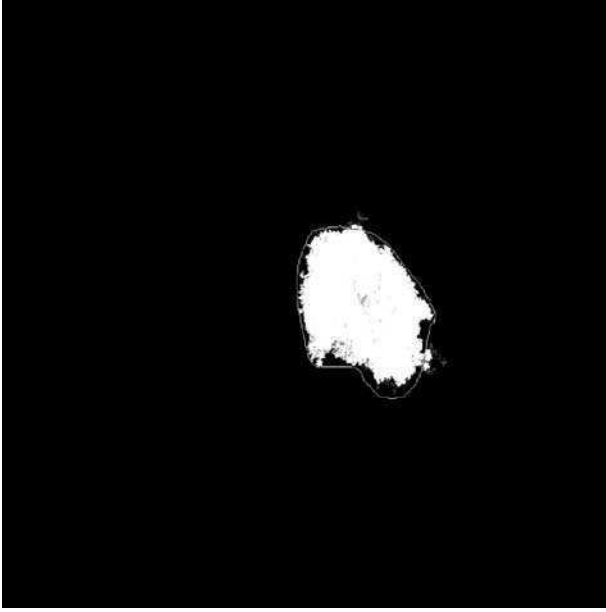
Fig. 27. Result of the fusion operator applied to the segmented regions in Figs. 8 and 12, for the benign mass shown in Fig. 6, with $C_r = 1.0$, $C_c = 1.0$. The contour drawn by the radiologist is superimposed for comparison. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, 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.

distinctions between a set and its complement. Indeed, it is the lack of distinction between a set and its complement that distinguishes a fuzzy set from a crisp set. The implementation of this concept depends on the definition of the fuzzy complement. We will employ the standard complement, where $\bar{A}(x) = 1 - A(x)$, for all $x \in X$. Choosing the Hamming distance, the local distinction between a given set A and its complement is measured by

$$|A(x) - \{1 - A(x)\}| = |2A(x) - 1|, \tag{2}$$

and the lack of local distinction is given by

$$1 - |2A(x) - 1|. \tag{3}$$

The measure of fuzzyness, $f(A)$, is then obtained by adding the local measurements:

$$f(A) = \sum_{x \in X} \{1 - |2A(x) - 1|\}. \tag{4}$$

The range of the function f is $[0, |X|]$: $f(A) = 0$ iff A is a crisp set; $f(A) = |X|$ when $A(x) = 0.5$ for all $x \in X$.
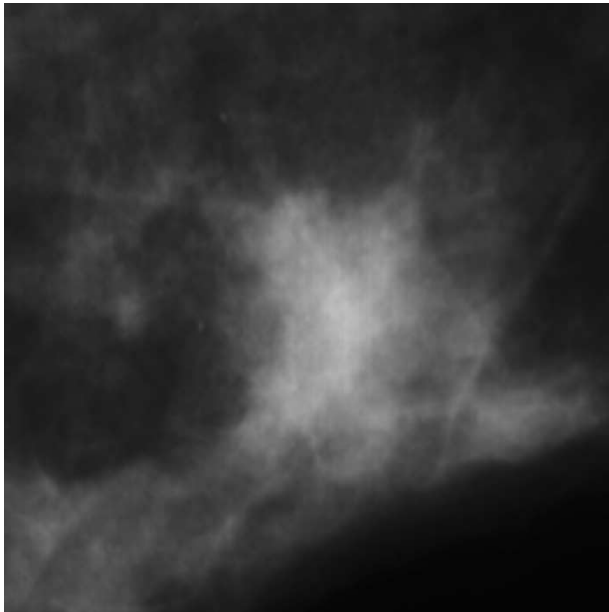
Fig. 28. Original version of a 700×700-pixel portion of a mammogram with a spiculated malignant tumor. Pixel size $= 62.5$ μm. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "Fuzzy fusion operators to combine results of complementary medical image segmentation techniques", Journal of Electronic Imaging, 12(3): 379–389, 2003. © SPIE and IS&T.

11.2. *Evaluation of the results of segmentation using the measure of fuzzyness*

In order to compare the result obtained by a method for segmentation with the corresponding contour of the mass drawn by the radiologist, we aggregated the result of segmentation with the reference contour using the fuzzy fusion method described above. The result of fusion was evaluated by a normalized measure of fuzzyness computed as

$$\bar{f}(A) = \frac{\sum_{x \in X}\{1 - |2\,A(x) - 1|\}}{|X|},\tag{5}$$

where A is a fuzzy set representing the result of aggregation, X is the set of pixels representing the entire image frame, and $A(x)$ is the degree of membership of the pixel x in X. 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.

Results of contour extraction by classical region growing with fuzzy-set preprocessed images: The agreement between the results of segmentation and the contours drawn by the radiologist was measured objectively by computing
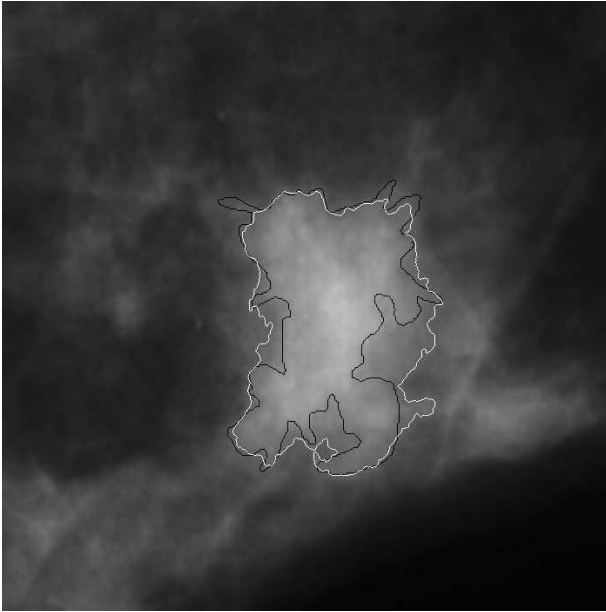
Fig. 29.   The white line shows the contour extracted by the closed-contour detection method for the mammogram in Fig. 28. The black line represents the boundary drawn by a radiologist (shown for comparison). Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and 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.

the measure of fuzzyness as in Eq. 5. The values obtained for the 47 mammograms tested were in the range $(0.13, 0.85)$, with the mean and standard deviation being 0.42 and 0.17, respectively. The measure of fuzzyness was less than 0.5 for 34 out of the 47 cases. In most cases where the measure of fuzzyness 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 fuzzyness, it is important to note that, for a spiculated lesion, there is no definite number of spicules that would characterize the lesion as malignant. The segmentation and fusion methods 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; see Sec. 5.1 for the results of pattern classification of masses.

Results of fuzzy region growing: The fuzzy regions obtained for the 47 mammograms tested were compared objectively with the corresponding contours drawn by the radiologist, by computing the measure of fuzzyness as in Eq. 5. 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 fuzzyness was smaller than 0.5 in 27 of the 47 cases analyzed. Regardless of this measure of agreement, it was
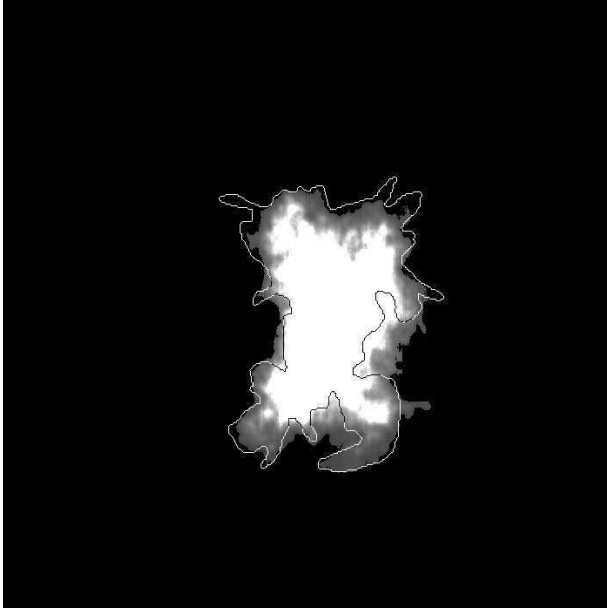
Fig. 30.   Result of fuzzy region growing for the tumor displayed in Fig. 28. The contour drawn by the radiologist is superimposed for comparison. Reproduced with permission from D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo, and J. E. L. Desautels, "Fuzzy fusion operators to combine results of complementary medical image segmentation techniques", Journal of Electronic Imaging, 12(3): 379–389, 2003. © SPIE and IS&T.

found that the fuzzy regions segmented contained adequate information to facilitate discrimination between benign masses and malignant tumors, as described in Sec. 5.2.

11.3.   *Evaluation of the results of fuzzy fusion using the measure of fuzzyness*

In order to realize an objective evaluation of the results of fusion, we used the contour drawn by an expert radiologist as the reference. 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 the fusion methods, denoted by RS_c, RS_r, and RF_r, respectively. We applied the fusion operator with the reliability measures both equal to unity (i.e. $C_r = C_c = 1.0$) 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 Figs. 29, 30, and 31, the measures of fuzzyness obtained were $f(RS_c) = 14{,}774.00$, $f(RS_r) = 14{,}245.85$, and $f(RF_r) = 9{,}710.82$, respectively, using Eq. 4. The aggregation or fusion of the two segmentation results
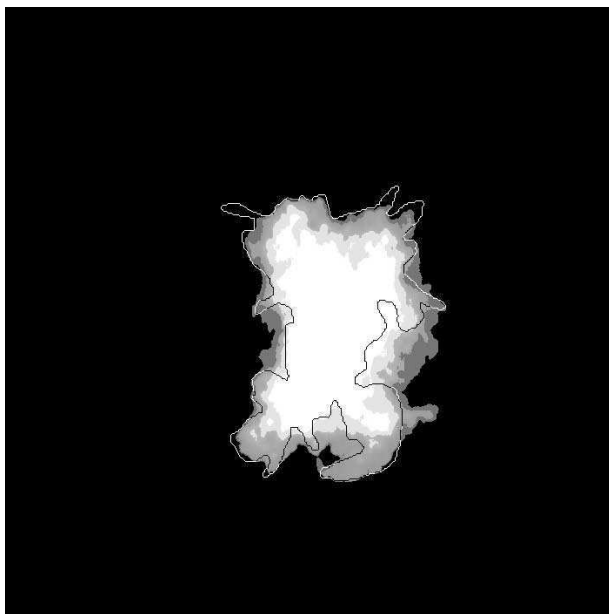
Fig. 31.   Result of the fusion operator applied to the segmented regions in Figs. 29 and 30, with $C_r = 1.0$, $C_c = 0.9$. The original image of the tumor is displayed in Fig. 28. The contour drawn by the radiologist is superimposed for comparison. Reproduced with permission from D. Guliato, R. M. 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.

presents less uncertainty than either, yielding a better result as expected. The corresponding measures for the results in Figs. 8, 12, and 27 are $f(RS_c) = 8,096.00$, $f(RS_r) = 9,223.73$, and $f(RF_r) = 7,905.44$, respectively. A smaller value of $f(RF_r)$ for the result of fusion than the same measure for the results of contour and region segmentation [shown as $f(RS_c)$ and $f(RS_r)$, respectively, in Table 3] indicates that the fusion result is better than the other two.

The methods were tested with 14 mammographic images of biopsy-proven cases; see Table 3. The values of C_c and C_r used to obtain the results of fusion for the 14 mammograms are also listed in the table. Note, however, that C_c and C_r were maintained equal to unity when computing the measure of fuzzyness 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 obtained provide good evidence that the fusion operator obtains regions with a higher degree of certainty than the results of the individual segmentation methods. Evaluation of the effect

Table 3. Measures of fuzzyness for the results of segmentation and fusion for 14 mammograms. 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.

Mammogram	C_r, C_c used to get fusion result	$f(RS_c)$	$f(RS_r)$	$f(RF_r)$	Fusion result is better
spic-s-1	1.0, 0.9	14,774.00	14,245.85	9,710.82	yes
circ-fb-010	1.0, 1.0	8,096.00	9,223.73	7,905.44	yes
spx111m	1.0, 1.0	5,130.00	9,204.63	4,680.67	yes
spic-fh0	1.0, 0.6	28,938.00	23,489.54	21,612.45	yes
circ-x-1	1.0, 0.8	6,877.00	2,990.71	3,862.76	no
spic-fh2	0.8, 1.0	45,581.89	38,634.00	34,969.52	yes
circ-fb-005	1.0, 1.0	26,176.00	34,296.26	25,084.75	yes
circ-fb-012	1.0, 0.9	16,170.00	15,477.88	12,693.21	yes
spic-db-145	1.0, 0.9	8,306.00	7,938.23	7,658.20	yes
circ-fb-025	1.0, 0.6	56,060.00	44,277.83	49,093.39	no
spic-fb-195	1.0, 1.0	11,423.00	12,511.86	10,458.00	yes
spic-s-112	1.0, 0.6	31,413.00	17,784.31	12,838.99	yes
spic-s-401	1.0, 0.6	13,225.00	11,117.99	11,195.90	no
circ-fb-069	1.0, 1.0	46,835.00	53,321.57	38,832.50	yes

of the fusion operator on the accuracy of the classification of the tumors as benign or malignant will be performed in our future projects.

12. Discussion and Conclusion

We have described a method for the extraction of the boundaries of breast masses in mammograms using a fuzzy-set-based preprocessing step to enhance the mass ROI and a region growing procedure. Contours provided by the method have demonstrated good agreement with reference contours drawn by a radiologist. Previous attempts to apply region growing methods directly to mammographic images of masses had failed; the success of the approach presented here lies in preprocessing the images using fuzzy set theory to obtain the membership degrees of the pixels in a mass ROI.

A survey of the current medical image processing literature indicates growing interest in interactive user-guided segmentation methods. The methods presented in this chapter are being used in an interactive procedure for the segmentation of medical images. We are working on the development of procedures to determine the seed region and optimal values for the parameters used in the methods (β, the threshold for contour detection, $\Delta\mu_{\max}$, and $\Delta CV_{\max}$). An automatic method that has been recently proposed to detect breast masses and tumors[10] may be used to provide seed regions and statistical parameters for the segmentation methods described in the present work. We intend to extend the definition of the mass

feature vector used in the fuzzy-set-based preprocessing step to include the use of more properties of the mass ROI, such as local gradient, texture, and variance.

The fuzzy region growing algorithm works with greater stability than others using the same approach owing to the fact that the region growing procedure checks all pixels within a given distance from the seed for inclusion in the region before enlarging the scope of search for further region growing. The algorithm uses the new acceptance concept of *accepted with restriction*, which avoids undesirable fluctuations in the statistics of the region being grown. The fuzzy segmentation techniques 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 the boundaries of the ROIs. The fusion approach allows the postponement of the decision regarding a crisp boundary of an ROI to a higher level of image analysis; the approach may even obviate the need to obtain a crisp boundary if discriminant features can be derived from the fuzzy regions, as demonstrated in this work.

We have demonstrated the application of fuzzy segmentation to the problem of identification of breast masses in mammograms. The fuzzy nature of the transition of a malignant tumor from its core region to the surrounding tissues is preserved by the membership values in the fuzzy region. Parameters computed from the segmented regions have shown potential use in classification of the regions as benign or malignant, with a sensitivity of 0.8 and a specificity of 0.9. We are exploring the potential use of shape factors of mass contours as well as measures of edge sharpness and texture[7-10,43] to improve the accuracy of classification. Further work on the optimization of the procedures and the classification of a larger database of masses is in progress.

We have also introduced the concept of a fuzzy fusion operator oriented by a finite automaton that is designed to combine different sources of information using a set of basic fusion operators. We have applied a particular operator — based on six basic operators guided by a specific finite automaton — to the problem of mammographic image segmentation for the detection of breast tumors, in order to improve the results obtained *a priori* by two different image segmentation techniques. The same operator has been used to derive a degree of agreement of each result with a reference contour given independently by an expert radiologist. The results obtained present good evidence that the fusion operator can provide regions with a higher degree of certainty than the results of the individual segmentation methods. Procedures need to be developed to determine appropriate values of the reliability measures (C_c and C_r in the present work) for each method; one possible approach to determine their values would be to use measures of accuracy of the sources of information being fused (in the present study, the methods that provided the contour and the region).

We would like to emphasize that the fusion operator described is an abstract notion that generalizes the concept of most of the fusion operators presented in the literature. It should be possible to extend the method to the fusion of multiple

contours and regions. The measure of fuzzyness may also be applied to compare the results of segmentation with a "gold standard" reference contour or region provided by a radiologist. The methods should be applicable to the segmentation of ROIs in several types of medical images.

Acknowledgments

References

1. I. N. Bankman (ed.), *Handbook of Medical Imaging: Processing and Analysis* (Academic Press, London, UK, 2000).
2. M. Sonka and J. M. Fitzpatrick (eds.), *Handbook of Medical Imaging, Volume 2: Medical Image Processing and Analysis* (SPIE Press, Bellingham, WA, 2000).
3. Alberta Cancer Board, Alberta, Canada. *Screen Test: Alberta Program for the Early Detection of Breast Cancer — 1999/2001 Biennial Report*, 2001.
4. Alberta Cancer Board, Alberta, Canada. *Screen Test: Alberta Program for the Early Detection of Breast Cancer — http://www.cancerboard.ab.ca/screentest/*, 2004.
5. H.-O. Peitgen (ed.), *Proceedings of the 6th International Workshop on Digital Mammography*, Bremen, Germany (June 2002), Springer-Verlag.
6. 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**(5–6) (2003) 820–835.
7. R. M. Rangayyan, N. M. El-Faramawy, J. E. L. Desautels and O. A. Alim, Measures of acutance and shape for classification of breast tumors, *IEEE Transactions on Medical Imaging* **16**(6) (1997) 799–809.
8. 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**(5) (2000) 487–496.

9. N. R. Mudigonda, R. M. Rangayyan and J. E. L. Desautels, Gradient and texture analysis for the classification of mammographic masses, *IEEE Transactions on Medical Imaging* **19**(10) (2000) 1032–1043.

10. 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) (2001) 1215–1227.

11. M. Sameti and R. K. Ward, A fuzzy segmentation algorithm for mammogram partitioning, in *3rd International Workshop on Digital Mammography*, eds. K. Doi, M. L. Giger, R. M. Nishikawa and R. A. Schmidt, Chicago, IL (9–12 June 1996), pp. 471–474.

12. C. H. Chen and G. G. Lee, On digital mammogram segmentation and microcalcification detection using multiresolution wavelet analysis, *Graphical Models and Image Processing* **59**(5) (1997) 349–364.

13. 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) (1993) 1403–1416.

14. N. Karssemeijer, Detection of stellate distortions in mammograms using scale-space operators, in *Information Processing in Medical Imaging*, eds. Y. Bizais, C. Barillot and E. Di Paola, Kluwer Academic Publishers, The Netherlands (1995), pp. 335–346.

15. A. Laine, W. Huda, D. Chen and J. Harris, Segmentation of masses using continous scale representations, in *3rd International Workshop on Digital Mammography*, eds. K. Doi, M. L. Giger, R. M. Nishikawa and R. A. Schmidt, Chicago, IL (9–12 June 1996), pp. 447–450.

16. L. Miller and N. Ramsey, The detection of malignant masses by non-linear multiscale analysis, in *3rd International Workshop on Digital Mammography*, eds. K. Doi, M. L. Giger, R. M. Nishikawa and R. A. Schmidt, Chicago, IL (9–12 June 1996), pp. 335–340.

17. M. Zhang, M. L. Giger, C. J. Vyborny and K. Doi, Mammographic texture analysis for the detection of spiculated lesions, in *3rd International Workshop on Digital Mammography*, eds. K. Doi, M. L. Giger, R. M. Nishikawa and R. A. Schmidt, Chicago, U.S.A. (9–12 June 1996), pp. 347–350.

18. T. Matsubara, H. Fujita, T. Endo, K. Horita, M. Ikeda, C. Kido and T. Ishigaki, Development of mass detection algorithm based on adaptive thresholding technique in digital mammograms, in *3rd International Workshop on Digital Mammography*, eds. K. Doi, M. L. Giger, R. M. Nishikawa and R. A. Schmidt, Chicago, IL (9–12 June 1996), pp. 391–396.

19. M. A. Kupinski and M. L. Giger, Automated seeded lesion segmentation on digital mammograms, *IEEE Transactions on Medical Imaging*, **17**(4) (1998) 510–517.

20. R. P. Nikhil and K. P. Sankar, A review on image segmentation techniques, *Pattern Recognition* **26**(9) (1993) 1277–1294.

21. R. L. Cannon, J. V. Dave and J. C. Bezdek, Fuzzy C-Means clustering algorithms, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **PAMI-8**(2) (1986) 248–255.

22. M. C. Clark, D. B. Hall, D. B. Goldgof, L. P. Clarke, R. P. Velthuizen and M. S. Silbiger, MRI segmentation using fuzzy clustering techniques, *IEEE Engineering in Medicine and Biology* (November/December 1994), pp. 730–742.

23. D. Guliato, R. M. Rangayyan, J. A. Zuffo and J. E. L. Desautels, Detection of breast tumor boundaries using iso-intensity contours and dynamic thresholding, in *4th International Workshop on Digital Mammography*, eds. N. Karssemeijer,

M. Thijssen, J. Mendris and L. van Erning, Nijmegen, The Netherlands (June 1998), Kluwer Academic, Dordrecht, pp. 253–260.

24. D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo and J. E. L. Desautels, Segmentation of breast tumors in mammograms by fuzzy region growing, in *20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Hong Kong (29 October–1 November 1998), pp. 1002–1004.

25. 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) (2003) 369–378.

26. A. Rosenfeld and A. C. Kak, *Digital Picture Processing*, 2nd edn. (Academic Press, New York, NY, 1982).

27. R. M. Haralick and L. G. Shapiro, Survey—image segmentation techniques, *Computer Vision, Graphics and Image Processing* 29 (1985) 100–132.

28. F. Meyer and S. Beucher, Morphological segmentation, *Journal of Visual Communication and Image Representation* **1**(1) (1990) 21–46.

29. H. Asar, N. Nandhakumar and J. K. Aggarwal, Pyramid-based image segmentation using multisensory data, *Pattern Recognition* **23**(6) (1990) 583–593.

30. J. F. Haddon and J. F. Boyce, Image segmentation by unifying region growing and boundary information, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **12**(10) (1990) 929–948.

31. T. Pavlidis and Y.-T. Liow, Integrating region growing and edge detection, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **12**(3) (1990) 225–233.

32. L. Vincent and P. Soille, Watershed in digital spaces: An efficient algorithm based on immersion simulations, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **13**(6) (1991) 583–598.

33. Y. Xiaohan and J. Yla-Jaaski, Direct segmentation in 3D and its application to medical images, in *Proc. SPIE Image Processing*, volume 1898, (1993), pp. 187–192.

34. A. Hadjarian, J. Bala, S. Gutta, S. Trachiots and P. Pachowicz, The fusion of supervised and unsupervised techinques for segmentation of abnormal regions, in *4th International Workshop on Digital Mammography*, Nijmegen, The Netherlands (June 1998), Kluwer Academic Publishers, pp. 299–302.

35. D. Guliato, R. M. Rangayyan, W. A. Carnielli, J. A. Zuffo and J. E. L. Desautels, Fuzzy fusion of results of medical image segmentation, in *SPIE Conferece on Medical Imaging — Image Processing*, volume 3661, San Diego, CA (20–26 February 1999), pp. 1075–1084.

36. 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) (2003) 379–389.

37. L. A. Zadeh, Fuzzy sets, *Information and Control* **8** (1965) 338–353.

38. G. J. Klir and B. Yuon, *Fuzzy Sets and Fuzzy Logic* (Prentice Hall, Englewood Cliffs, NJ, 1995).

39. R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 2nd edn. (Addison-Wesley, MA, 1992).

40. J. E. Cabral Jr., K. S. White, Y. Kim and E. L. Effmann, Interactive segmentation of brain tumors in MR images using 3D region growing, in *Proc. SPIE, Image Processing*, volume 1898 (1993), pp. 171–181.

41. Y. Chang and X. Li, Adaptive region growing, *IEEE Transactions on Image Processing* **3**(6) (1994) 868–872.

42. D. Guliato, R. M. Rangayyan, F. Adorno and M. M. G. Ribeiro, Analysis and classification of breast masses by fuzzy-set-based image processing, in *6th International Workshop on Digital Mammography*, Bremen, Germany (June 2002), pp. 195–197.

43. O. Menut, R. M. Rangayyan and J. E. L. Desautels, Parabolic modeling and classification of breast tumors, *International Journal of Shape Modeling* **3** (1997) 155–166.
44. R. R. Yager, Connectives and quantifiers in fuzzy sets, *Fuzzy Sets and Systems* **40** (1991) 39–75.
45. I. Bloch, 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) (1996) 52–67.

This page intentionally left blank

ROC METHODOLOGY AND ITS APPLICATION IN BREAST CANCER DIAGNOSIS

SMADAR GEFEN

Department of Neuroanatomy and Biology
Drexel University College of Medicine, Philadelphia, PA 19129-1096
sgefen@drexel.edu

OLEH TRETIAK

Department of Electrical and Computer Engineering
Drexel University, Philadelphia, PA 19104
tretiak@cbis.ece.drexel.edu

DAVID GEFEN

Department of MIS Management, LeBow College of Business
Drexel University, Philadelphia, PA 19104
gefend@drexel.edu

ROC analyses are common in medical decision making because they produce a visual index and because of the ease with which they can be applied to machine learning and data mining. In the last two decades ROC curves have been used excessively to explore the relationship between the sensitivity and the specificity of a variety of diagnostic systems. This paper presents techniques for ROC analysis of a diagnostic system and discusses ways of improving the accuracy of ultrasound image diagnoses through a combination of computer-generated features, a patient's age, and a radiologist's diagnosis. The application of an ROC to the specific case of identifying breast cancer lesions is shown to significantly decrease false positive while maintaining a high degree of true positive diagnoses.

Keywords: ROC; breast cancer diagnosis.

1. Introduction

Receiver Operating Characteristics (ROC) curve is a statistical method that evaluates the accuracy of a dichotomous classification scheme. The ROC curve shows the tradeoff between the detection probability and the false-alarm probability of a detection system, or in medical terms, the tradeoff between the true-positive and false-positive of a diagnosis system. One objective of an ROC analysis is to identify a classification scheme which will maximize classification accuracy, i.e. achieve an acceptable tradeoff between positive hits and false alarms.[1,2] A second objective of an ROC analysis is as a means of comparison between classifiers. The efficacy of a classifier can be measured by the area under the ROC curve. Generally speaking,

a classifier is better when its ROC curve is more to the top left, showing more true positive and less false positive rates.

Performance analysis based on ROC curves has been used in many classification studies.[3–5] In this paper we demonstrate the usage of ROC analysis for the evaluation of a breast cancer diagnosis system. ROC analysis was used to quantify the discriminatory properties of eighteen different features and their combinations. The ROC summary index, A_z, and its confidence interval, $(\beta - \alpha)$, were estimated based on the approach developed in Gefen *et al.*[6]

Mammography is currently the preferred way to screen for breast cancer. But this method is not perfect, as benign lesions are still found in a high proportion of women with positive mammographic tests. Ultrasonic imaging is another technique which is widely used as a diagnostic modality because of its non-ionizing and non-invasive nature combined with its low cost. But a radiological visual inspection of an ultrasound imaging of the breast on its own is of limited diagnostic value.[7] A definitive diagnosis can be achieved through biopsy, but such a procedure is an expensive and highly traumatic procedure for the patient. Thus, improving screening methods is an imperative.

In this chapter we present a case study that demonstrates an ROC analysis that was used to improve the accuracy of ultrasound diagnoses through a combination of computer-generated features, patient age, and diagnosis by a radiologist. The desired goal was to achieve a reduction in the number of biopsies that result in a benign diagnosis. Decreasing the degree of false positive diagnoses while maintaining a high degree of true positive fraction can prevent unnecessary biopsies.

The case study presented here first examines the individual performance of each feature, in which the characteristics of each feature are reviewed and the efficacy of each feature is evaluated. Second, the features are fused using linear and quadratic discriminant functions in order to improve the overall classification performance. Performance evaluation results show that feature fusion leads to improved performance over any single feature. The performance of this combined set of features is then shown to be significantly better than a radiologist's pre-biopsy decision.

2. ROC Methodology

2.1. *Classification systems*

Classification is a common operation whenever a decision needs to be made automatically. In a classification system a set of features that embodies information regarding the state or category of a certain pattern is analyzed, where a pattern refers to any type of data that can be measured about the object being investigated. Classification is applied in many types of applications that involve pattern recognition. For example, voice recognition analyzes the audio signal and identifies spoken words. Another example is in the manufacturing industry where a classifier decides

if a component in an assembly line has a defect and should be removed or not. And in medical diagnostic systems, classification deals with a decision regarding the presence or the absence of a disease or whether, for instance, a lesion is malignant or benign. This study focuses on a binary classifier that uses a set of features that are derived from a given pattern; a binary classifier makes a True or False decision.

In general, several stages take place when designing a supervised classifier. These steps consist of features generation, features selection, classifier design, and classifier evaluation.

The first step is to define what type of features should be extracted from a given pattern. A vector of these features: $\mathbf{x} = [x_1, x_2, \ldots, x_l]$ is then derived from each given sample of the pattern. Useful features are those that best characterize the class that a sample pattern best represents. Good feature vectors, when applied to pattern samples, will generate points in the feature space that will be spatially separable into the different classes. On a practical level, it is also desirable that these features be efficiently computable, especially in applications where real-time performance is required.

However, there are often too many available features, and so it is necessary to select a subset of features. Working with a limited number of features is desirable not only because of computation time constraints but, most importantly, in order to assure classifier generality. A minimum ratio between the training dataset size and the number of features should be maintained in order to prevent an over-fitting of the classifier parameters to the specific training dataset. (Usually a ratio of 10–20 is recommended.) Since increasing the training set is often not possible, keeping a relative low number of features is essential. Moreover, it can be shown statistically that based on random noise alone, a large number of even unrelated parameters will result in a good classification.

Nevertheless, although two features may result in a good discrimination when considered separately, sometimes there will not be a significant improvement when considered together. This may be because the two features are correlated and therefore add little new additional information when combined. In statistical terms this is called multi-colinearity. Statistical methods to reduce the number of features, while retaining as much as possible of their combined class discriminatory information, are discussed by Theodoridis and Koutroumbas.[8]

Next, the classifier is designed with respect to some optimality criterion. With a supervised classifier, the classifier parameters are computed based on a training set of patterns in which the true class membership is known. In general, this feature space is divided into class regions that are separated by hyperplanes (in the linear case) or surfaces (in the nonlinear case). A classifier receives as input a vector of object features. Based on its location in the feature space, the object is then classified. When only two classes are involved, the classifier maps (fuses) a vector of features into one scalar value. This scalar is then compared against a threshold. If it is above the threshold, the associated pattern is classified as belonging to a certain class. Otherwise, the pattern is classified as the other class.

Figure 1 shows the general stages involved in classifier design and evaluation. The feature vectors $X_{[Nxl]}^{Training}$ are computed for the given training sample set $I_{i=1,\ldots,N}^{Training}$, where N is the size of training set and l is the number of features. Next, the classifier parameters, $\mathbf{C}(X_{[Nxl]}^{Training})$, are estimated. In the testing phase a testing sample set, $I_{i=1,\ldots,M}^{Testing}$, is used to compute a new feature vector, that in turn is fed into the classifier. The classifier then makes a classification, i.e. a decision regarding the class membership of each testing sample. Commonly-used performance measures can now be computed, as described in the below, for a diagnostic system.

Each time a decision is made by the classifier there are four possible outcomes:

(i) There is a disease and the system classifies the pattern as positive — True Positive (TP).
(ii) There is a disease and the system classifies the pattern as negative — False Negative (FN).
(iii) There is no disease and the system classifies the pattern as negative — True Negative (TN).
(iv) There is no disease and the system classifies the pattern as positive — False Positive (FP).

These outcomes can be represented in a contingency table as shown in Fig. 2.
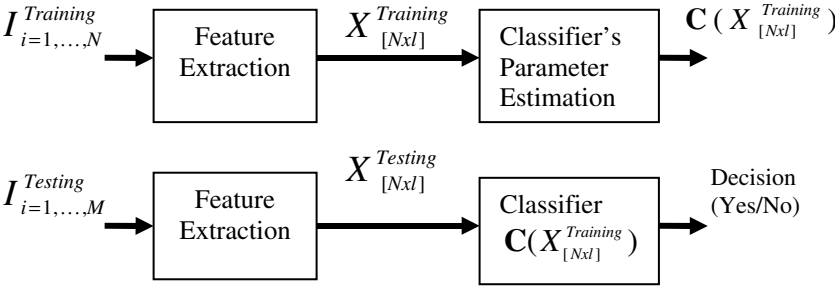
Fig. 1.   Classifier design (the training phase) and classifier evaluation (the testing phase).

		True Class	
		Positive (P)	Negative (N)
Classifier Decision	True	**True Positive (TP)**	**False Positive (FP)**
	False	**False Negative (FN)**	**True Negative (TN)**

Fig. 2.   Contingency table.

Two commonly-used metrics, derived from the contingency table, are the sensitivity and the specificity of a classification system. The system sensitivity, defined as the true positive rate (detection probability), is:

$$P_{TP} = \frac{TP}{TP + FN} \equiv \frac{TP}{P}.$$

The system specificity is the true negative rate

$$P_{TN} = \frac{TN}{FP + TN} \equiv \frac{TN}{N}$$

or $1 - P_{FP}$, where P_{FP} is the false positive rate (false-alarm probability)

$$P_{FP} = \frac{FP}{FP + TN} \equiv \frac{FP}{N}.$$

A classification system must balance between its detection rate and its false-alarm rate. In some procedures achieving the highest possible detection rate is desirable, even if this high detection rate comes at the expense of a high false-alarm. In these cases, treating a patient for a falsely detected disease may be tolerated. In other procedures, the risk and cost involved in treatment call for a system with a low false-alarm rate, sometimes even at the expense of a lower detection rate.

Other metrics, such as the precision

$$\frac{TP}{TP + FP},$$

and the accuracy

$$\frac{TP + TN}{P + N},$$

are also used in the literature to evaluate the performance of a diagnostic system.

2.2. ROC estimation

The ROC shows the detection rate, P_{TP}, as a function of the false-alarm rate P_{FP}. Figure 3 shows the ROC curves of different classifiers — $\alpha_0, \alpha_1, \alpha_2$, and α_3. Each point on the ROC curve represents a different classifier detection threshold. In one extreme case where $P_{TP} = P_{FP} = 0$, the classifier detection threshold is so high that it always makes a "negative" decision. In this case, if there is a disease, it will not be detected. Nonetheless, no false alarms will be made either. In the other extreme case where $P_{TP} = P_{FP} = 1$, the threshold is so low that a positive decision will always be made. Here, if there is a disease it will always be detected, but even if there is no disease the alarm will be raised. Obviously, both extremes are not good. A good classifier should have a detection rate close to one, $P_{TP} \Rightarrow 1$, and a false alarm close to zero, $P_{FP} \Rightarrow 0$. In practice, however, there is always a tradeoff between P_{TP} and P_{FP}. Based on the specific application requirements, the tradeoff point on the ROC curve between these two parameters can be chosen to set the threshold accordingly.
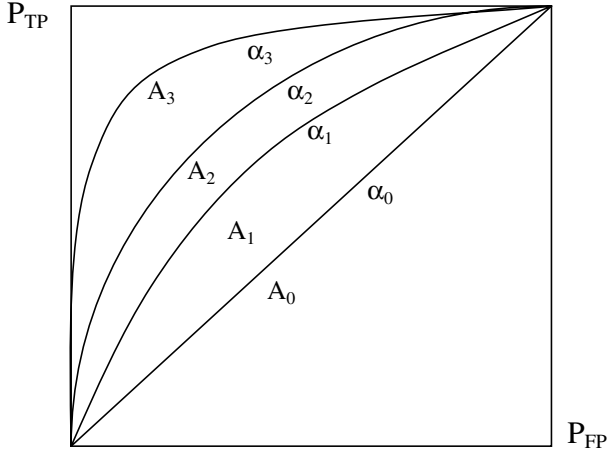
Fig. 3. Receiver Operating Characteristic (ROC) depicting the detection rate, P_{TP}, as a function of the false-alarm P_{FP}.

Another important service an ROC curve provides is a means to statistically compare classifiers. The area under the ROC curve, A_z, is commonly used as a summary index; the closer this area is to one, the better the classifier is. A classifier with an $A_z = 0.5$ is as good as a random guess. Note, however, that it is possible for a classifier to have a smaller A_z than another classifier and still have a lower P_{FP} for the same P_{TP}.

In this study an empirical ROC curve and an A_z were generated as follows. The empirical ROC was evaluated by: (1) pooling the feature data of malignant and benign cases, (2) sorting the feature data in ascending order, and (3) varying a threshold while classifying the data that are above the threshold as positive and data that are below the threshold as negative. Thus, for each threshold the True Positive Fraction (TPF) and the False Positive Fraction (FPF) define a point on the ROC curve. Once the ROC curve was given, the area underneath it was computed to estimate A_z.[3]

3. Imaging Methods

The challenges that engineers face in the research and development of medical imaging system are manifold. One of these challenges is to devise an imaging modality that will detect tumors and determine whether they are malignant or benign accurately at its early stages. In addition, it is also beneficial if these modalities are non-invasive and not painful. Especially in the case of breast cancer, the comfort and safety of patients during breast imaging and biopsy is imperative in order not to discourage participation in periodic screening and diagnosing.

Various imaging methods have been developed and proven useful for breast cancer detection. These methods include X-ray mammography, ultrasound imaging

(UI), magnetic resonance imaging (MRI), T-Scan, and nuclear medicine imaging (NMI). While X-ray mammography is the only modality that is approved for breast cancer screening by the U.S. Food and Drug Administration (FDA), UI, MRI, T-Scan, and NMI are recognized as complementary procedures to mammography in the diagnosis of breast cancer. Other imaging methods such as computerized thermal imaging are currently undergoing clinical testing to prove their efficacy as adjacent tools to mammography.

3.1. *Mammography*

X-rays, discovered by Wilhelm Konrad Roentgen in 1895, for which he won the first Nobel Prize for physics in 1901, have been in medical use since 1896.[9] By the 1970s the first mammography imaging systems that utilized X-rays to image breast tissues were manufactured. These systems have been vastly improved in the last decade, with the ratio between image quality and radiation dose reduced to a negligible risk level, considering the benefits of early breast cancer detection.

Replacing the film cassette with a digital receptor in the mammography system without reducing the image quality opened up many new options. Having the image displayed almost immediately on a computer screen significantly shortens examination time, allows for real-time stereotactic biopsy, enables remote and prompt consultation with other physicians, and brings about the benefits of digital information such as the ability to electronically manipulate images to enhance, magnify, archive and retrieve them.

While digital mammography at this stage cannot match the spatial resolution of standard mammography, which is crucial for the detection of small abnormalities, its better contrast emphasizes abnormalities. To date, studies that compared digital mammography to standard film mammography showed no significant statistical difference between the two modalities. For example, Bick *et al.*[10] compared the performance of a full-field digital mammography system (GE Senographe 2000D) with standard mammography and found that all abnormalities, including both masses and microcalcifications (often associated with early-stages of breast cancer) visible on standard film mammograms, were at least equally well visible with digital mammograms. However, to fully estimate the impact of digital mammography on breast cancer diagnosis, more cases from both technologies need to be compared.[11]

3.2. *MRI*

Breast magnetic resonance imaging (MRI) is a promising complementary method to mammography screening. MRI is effective in imaging dense breast tissue, in evaluating the extent of breast cancer, in staging and treating breast cancer, and in disease tracking after treatment. MRI uses powerful magnetic fields and radio waves to create images, such as of the breast. The main component of most MRI systems is a large cylindrical magnet. During the examination, a radio signal is turned on and off, and, subsequently, some of the energy is absorbed by different molecules

in the body, and some is echoed from the body. These echoes are continuously measured by the MRI scanner. A digital computer reconstructs these echoes into images.

One of the benefits of MRI is that it can easily acquire direct views of the breast in almost any orientation. In contrast, mammography requires re-orientation of the breast and of the mammography system for each view. In an MRI, a contrast material is usually injected into a vein in the arm before the examination in order to improve the quality of the images by providing a contrast medium.

Although MRI could be a significant supplementary tool to mammography in the diagnosis of breast cancer, there are still hurdles an MRI must overcome before it gains wider acceptance. While mammography can image calcifications, MRI cannot image these tiny calcium deposits. Also, MRI cannot always distinguish between cancerous and non-cancerous abnormalities, which may result in unnecessary breast biopsies. MRI screening is also more expensive, costing about ten times as much as mammogram screening, and more time consuming, taking three times as much time as it takes a mammogram. On the other hand, there are abnormalities that can be found with an MRI but that may not be visible with mammography or ultrasound, especially in the dense tissues of the breast, making an MRI guided biopsy system necessary. Currently, breast screening with an MRI is performed mostly at research centers and to augment the capabilities of standard mammography and breast ultrasound.

3.3. *Ultrasound imaging*

Breast ultrasound uses high-frequency waves to image the breast. High-frequency waves are transmitted from a transducer through the breast. The sound waves echo off the breast; this echo is picked up by the transducer and then translated by a computer into an image that is displayed on a computer monitor. The ultrasound unit contains a control panel, a display screen, and a transducer (resembling a microphone or computer mouse). Before the examination begins the patient is instructed to lie on a special table. The ultrasound technologist covers the part of the breast that will be imaged with a gel. The gel lubricates the skin and helps the transmission and reception of the sound waves.

When the examination begins, the radiologist glides the transducer over the breast. The transducer emits sound waves and picks up the echoes. The computer then analyzes these echoes and displays an image. The shape and intensity of the echoes depend on the density of the breast tissue. If a fluid-filled cyst is imaged, most of the sound waves will pass through the cyst and emit faint echoes. However, if a solid tumor is imaged, the sound waves will bounce off the tumor and the pattern of echoes will be translated by the computer into an image that the radiologist will recognize as indicating a solid mass. An ultrasound exam usually lasts between 20 and 30 minutes, but may take longer if the operator has difficultly finding the breast abnormalities being examined.

The advantage of ultrasound is in its high contrast image. But this comes at the cost of poor spatial resolution. Because of this, ultrasound imaging is used mostly as a supplement to mammography. Once a suspected lesion has been found with screening mammography, ultrasound is a useful tool, for instance to further explore if it is a cyst or a mass. Ultrasound is a non-expensive and non-invasive method that allows for a maximum degree of freedom of viewing orientation. This makes it a very efficient tool for imaging when a specific region is investigated or when guided biopsy is necessary. However, the disadvantage of ultrasound is that it is a subjective imaging method in which the quality of analysis depends on the sonographer's experience.

For example, Fig. 4 shows saggital (a) and horizontal (b) views of a Fibroadenoma (benign tumor) and saggital (c) and horizontal (d) views of an Invasive Ductal Carcinoma. Prior to biopsy, the first lesion was judged to be malignant while the second lesion was indeterminable.

4. Diagnosis of Breast Cancer

Considerable research has been devoted to examining the use of ultrasound for the diagnosis of breast lesions. In addition to the use of radiological interpretation, a number of research groups have been reporting advances on the development in automatic classifiers. A research group from the Department of Surgery at the China Medical College and Hospital in Taichung, Taiwan, and from the Department of Computer Science and Information Engineering at the National Chung Cheng University is conducting research on systems that diagnose based on correlation, texture, and morphological features derived from manually segmented ultrasonographic images.[12−15] Chen *et al.*[12] used digital ultrasonography to differentiate between benign and malignant solid breast nodules. An autocorrelation feature vector was extracted from an ROI that was located by a physician. The classifier in use in that study was a Learning Vector Quantization (LVQ) neural network. The evaluation method used was the *k-fold holdout* method. Results from 52 malignant and 88 benign cases produced an A_z of 0.9560 with a standard deviation of 0.0183 for $k = 10$.

Horsch *et al.*[16,17] proposed an algorithm for mass segmentation on ultrasound images. Horsch *et al.*[16] compared their proposed automatic delineation of a mass with a manual evaluation. An ROC analysis of the two classifiers' performance was done. The input of the linear discriminant classifier in use had four computer-generated features. Results from 94 malignant and 306 benign cases were an A_z of 0.90 for manual delineation and an A_z of 0.87 for automatic delineation.

Chou *et al.*[18] report on a system that computes features from a tumor's boundary based on stepwise logistic regression. Kitaoka *et al.*[19] report using histogram-based features for benign-malignant differentiation. Lefebvre *et al.*[20] report a sensitivity of 95% at a specificity of 80% for benign-malignant differentiation using three features computed from a region in the vicinity of the tumor
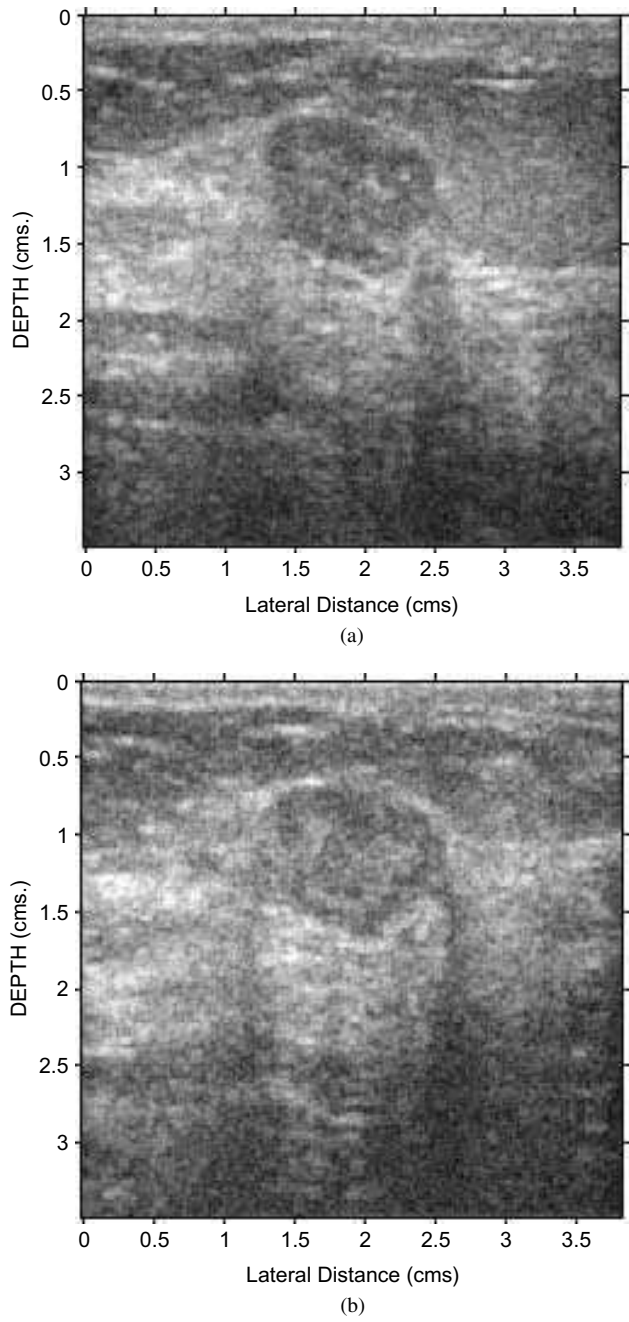
Fig. 4. Saggital (a) and horizontal (b) views of a fibroadenoma (benign tumor) and saggital (c) and horizontal (d) views of an Invasive ductal carcinoma. Prior to biopsy, the first lesion was judged to be malignant while the second lesion was indeterminate.
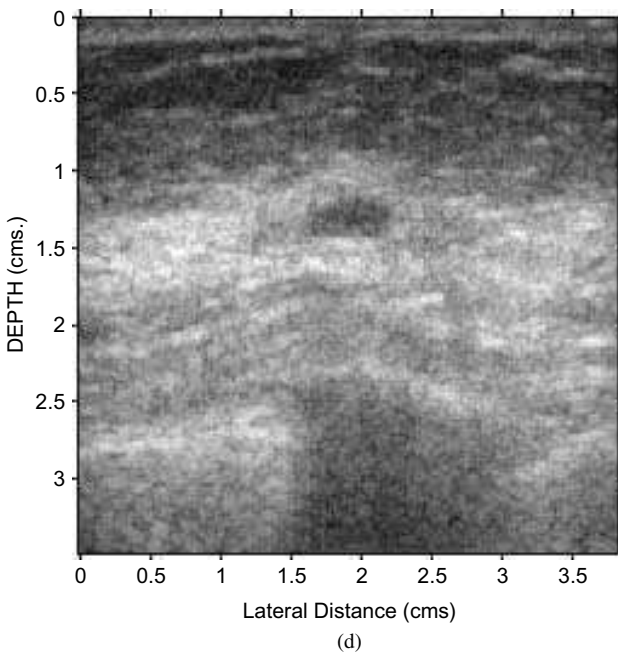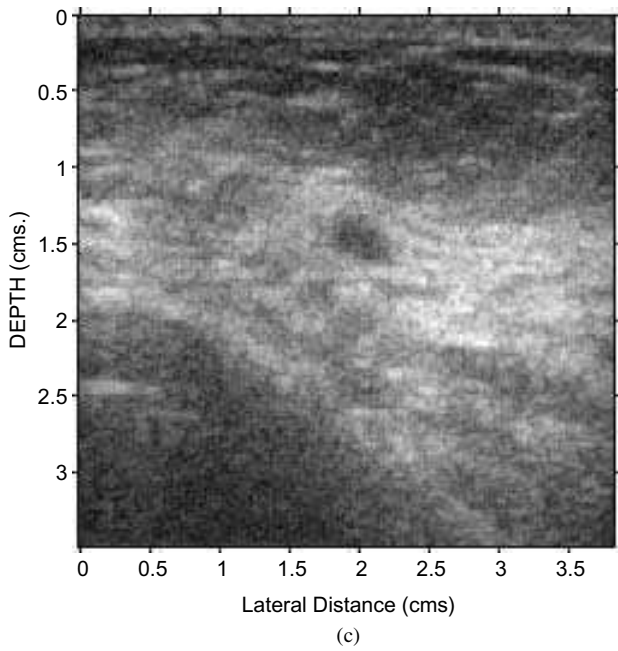
(c)

(d)

Fig. 4. (*Continued*)

boundary. Shivaramakrishnan *et al.*[21] used Haralick's[22] texture features, and found that they could increase specificity by providing a better characterization of solid lesions. The above research projects used conventional envelope images as the input to analysis. In contrast, in our case study the digitized RF data which contains amplitude and phase information is also given.

4.1. *Classification — A case study*

This case study classifies breast masses into malignant and benign tumors using three classes of features: (1) RF-signal-based features, (2) morphological features, and (3) auxiliary features. The study evaluates the discriminatory properties of each of these features and of a fused version of them. Fusion is done by linear and nonlinear classifiers. The human subject study presented here was approved by the Institutional Review Board of Thomas Jefferson University Hospital. Female patients who were scheduled for biopsy of a presumed solid breast mass were recruited for the study. Patients were scanned by a radiologist using the UM9 HDI system (Phillips, Bothell, WA, USA). The scans were obtained using the L10-5 ultrasonic transducer. The focal zone on the transmitted pulse was set to the center of each solid mass. The system used dynamic focusing on the received signal. Transmission frequencies, between 6–8 MHz, were determined by the depth of penetration.

RF data, used for the extraction of the RF-signal-based features, were sampled at 20 MHz with 12 bits quantization after applying analog time-gain control (TGC). Prior to processing, data were normalized. Data normalization is an amplitude scaling operation that accounts for gain changes in the RF data that had occurred as the receiving end aperture was dynamically expanded from a few elements at skin-line to a maximum of 128 elements. With normalization, data are scaled to compensate for this apparent depth-dependent change in gain.

The morphological features were computed in the vicinity of lesion boundaries and interiors using the envelope images. The lesion boundaries were manually outlined from the RF signal envelop image. A sample RF signal envelop image with 192 scanlines and 1367 RF sample points per scanline is shown in Fig. 5. The resolution of one data sample is 0.2 mm (across scanlines) and 0.04 mm (along scanlines). Along with each RF image, the center of the lesion and the rectangle that encloses the lesion's boundary were given (as indicated manually by the radiologist). The smallest lesion area was enclosed by a rectangle of a 4 mm lateral by 4 mm axial, while the largest lesion area was enclosed by a rectangle of a 24 mm lateral by 8 mm axial dimension. Orthogonal views of each mass were imaged. The feature values from different images of the mass were averaged to result in one representative feature value.

Study materials consisted of digitized RF ultrasound scans of mass breast lesions containing 39 positive (malignant) and 125 negative (benign) cases, as determined through pathology. The pathological results were composed of a variety of histopathological diagnostic categories. Eighteen features, F1 through F18, were considered. Features F1 through F8 were RF-signal-based and were computed
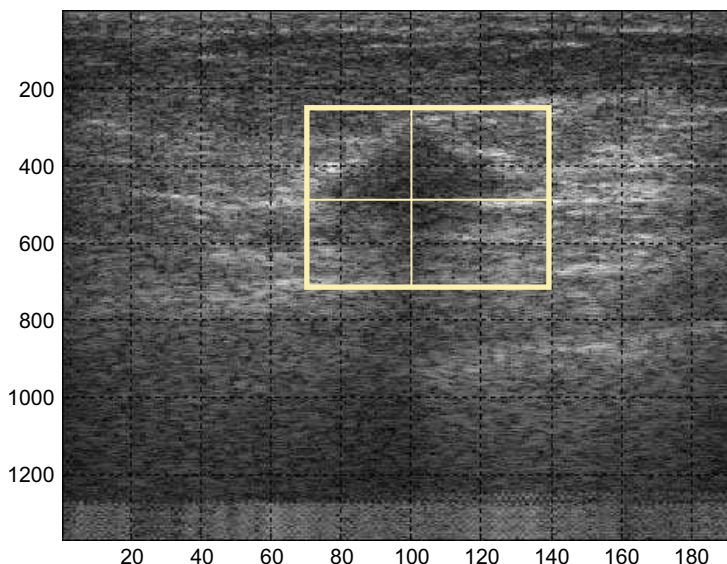
Fig. 5. RF signal envelope image with 192 scanlines and 1367 RF points per scanline. The resolution of one data sample is 0.2 mm across scanlines and 0.04 mm along scanline. The center of the lesion and the rectangle enclosing its boundary were determined manually by a radiologist.

Table 1. Tissue characterization features.

Features		Description
RF-signal-based	F1-F2	Nakagami model
Features	F3-F6	Generalized Spectrum model
	F7-F8	PowerLaw Shot Noise model
Morphological	F9	Mass boundary contrast
Features	F10	Mass orientation
	F11	Mass circularity
	F12	Mass compactness
	F13	Shape roughness
	F14	Mass boundary variation
	F16	Tumor Boundary Roughness (TBR)
Auxiliary Features	F17	Patient's age
	F18	Radiologist Diagnosis - LOS

directly from the RF scans. Features F9 through F16 were morphological features that were computed from the envelope image. Features F17 and F18 are auxiliary features, age and the radiologist's pre-biopsy decision regarding the presence of disease, respectively. Table 1 lists the eighteen features analyzed in this study.

4.2. *RF-signal-based feature*

Previous research examining the Nakagami distribution of the envelope of backscattered echo has determined its ability to classify benign and malignant lesions in

B-scan images of breast tissue.[23] The Nakagami distribution for the envelope R of
the backscattered echo is given as

$$f(r) = \frac{2m^m r^{2m-1}}{\Gamma(m)\Omega^m} \exp\left(\frac{-m}{\Omega}r^2\right)u(r).$$

where $u(\cdot)$ is the unit step function, $\Gamma(\cdot)$ is the gamma function, m is the Nakagami
parameter indicating the statistics of the envelope, and Ω is a scaling parameter of
the average energy in the backscattered echo. These parameters were normalized so
that they would be insensitive to operator-gain settings, time-gain-compensation
settings, and machine-dependent effects. Techniques of frequency diversity and compounding have been investigated for performing this normalization.[24]

Briefly, frequency diversity was created by splitting the spectrum of the
backscattered echo using two bandpass filters of second order and at a bandwidth of
1 MHz. The center frequencies of the filters were set at 5.75 MHz and 6.75 MHz. The
quadrature demodulated envelopes of the frequency diverse signals had a Nakagami
distribution with non-identical parameters: m_1, Ω_1, m_2, and Ω_2. After the creation
of diversity in the frequency domain, compounding was performed by combining
the envelopes of the frequency-diverse signals through a weighted summation as

$$R_{eff} = \sqrt{\frac{m_1}{\Omega_2}}R_1 + \sqrt{\frac{m_2}{\Omega_1}}R_2$$

where R_{eff} is the compounded envelope and R_1 and R_2 are the envelopes of the
frequency diverse signals. The analytical expressions for the parameters of the
Nakagami distributed compounded envelope, namely m_{eff} and Ω_{eff}, were derived as

$$m_{eff} \approx m_1 + m_2$$

$$\Omega_{eff} = m_1\frac{\Omega_1}{\Omega_2} + m_2\frac{\Omega_2}{\Omega_1} + 2\frac{\Gamma(m_1 + 0.5)\Gamma(m_2 + 0.5)}{\Gamma(m_1)\Gamma(m_2)}.$$

The details of the derivation process are given in Dumane and Shankar[24] where it
is shown that the parameters m_{eff} and Ω_{eff} are insensitive to scaling and are therefore normalized. These parameters were estimated based on data sample extracted
from the tumor interior ROI (see Fig. 6). The analysis window size was of 20 scanlines (4 mm) and 100 RF points per scanline (4 mm). Concatenation of the data
segments yielded 2000 sample points available for parameter estimation. In order
to avoid correlation among the samples, which might have been caused due to over
sampling during data collection, every other sample was chosen to perform the
estimation.

Another set of features was derived with the Generalized Spectrum (GS)
method.[25-27] GS is defined over the bi-frequency plane by:

$$G(f_1, f_2) = \mathrm{E}[Y(f_1)Y^*(f_2)],$$

where $\mathrm{E}[\cdot]$ is the expected value operator, $Y(f_i)$ is the spectral component of a local
RF segment at frequency f_i, and superscript * indicates a complex conjugate. For
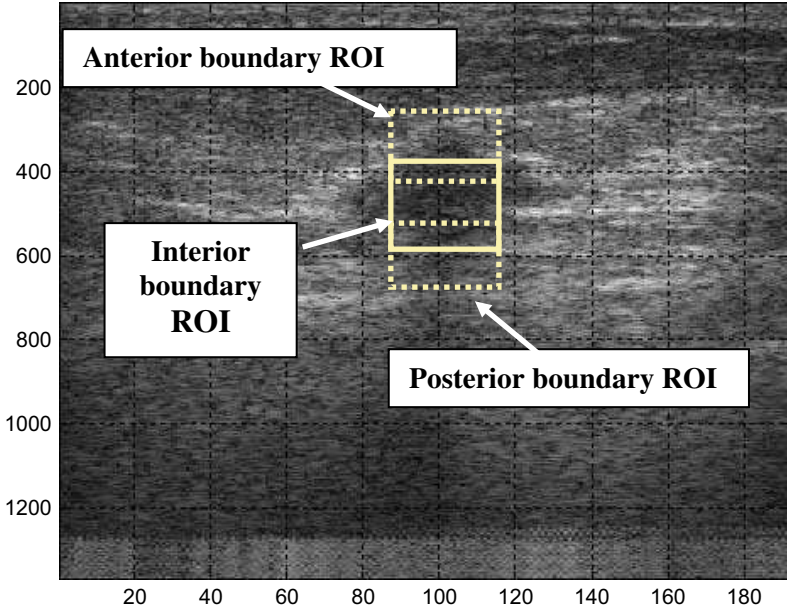
Fig. 6. The location of the three ROIs: Tumor *Front-Edge* ROI, Tumor *Interior* ROI, and Tumor *Back-Edge* ROI.

a stationary random process $G(f_1, f_2) = 0$ whenever $f_1 \neq f_2$. In other words, the off-diagonal values of GS are a measure of the non-stationarity of the signal. The GS was estimated from Discrete Fourier Transform (DFT) segments of the RF data along the axial direction: the outer products of the DFT segments were averaged over an analysis window. Two normalization methods were used to mitigate the effects of attenuation. The energy-normalized GS (EGS) was computed by scaling the DFT of the segment to unit energy. The system-normalized GS (SGS) was computed by scaling all DFT values to unit magnitude. The EGS preserves the spectral magnitude while SGS preserves only the spectral phase patterns.

The GS of each window was used to compute several features from the Collapsed Average (CA) of the GS, defined as

$$\hat{C}(h) = \frac{1}{M(h)} \left| \sum_{m-n=h} \hat{G}(n, m) \right|,$$

where $\hat{G}(n, m)$ is the discrete GS estimate, $M(h)$ is the number of GS off-diagonal elements satisfying $m - n = h$, and h is an index for the frequency difference corresponding to each off-diagonal in the bi-frequency plane. The CA reflects consistencies in the phase patterns of the spectrum due to non-stationarities.

The analysis windows were $2\,\text{mm}$ lateral by $4\,\text{mm}$ axial (10 by 100 samples). Three rectangular ROI were selected for each breast mass. These ROI captured the tumor *anterior boundary*, the tumor *interior*, and the tumor *posterior boundary*

(see Fig. 6). All ROI in the dataset were larger than the analysis window so that multiple windows were applied within the ROI with a 50% overlap in both the lateral and axial directions. This led to numerous estimates of the feature. The reported feature for each ROI was the value at the 75 percentile in magnitude (75% order statistic).

Both simple (F3 and F4) and complex features (F5 and F6) were examined. Simple features are single quantities extracted from the CA estimate that relate to a specific dominant scatterer property within the analysis window. Complex features are formed by linearly combining several simple features. Further details concerning these features' extraction can be found in Ref. 28.

Previous research suggested that a Power-Law Shot Noise (PLSN) model for the ultrasound RF echo might also discriminate between malignant and benign masses.[29–31] This model is applicable when (1) the tissue is considered as a collection of Poisson point scatterers that are embedded in a uniform medium, and when (2) the attenuation of the pulse envelope can be assumed to be power-law decaying. The power-law exponent β is related to tissue attenuation. Hence, the RF echo, $r(t)$, modeled as the sum of all reflections, is the backscattered signal that is received by the transducer,

$$r(t) = \sum_i a_i (t - t_i)^{-\nu} u(t - t_i) \cos(\omega_c(t - t_i) + \phi_i)$$

where t is the time of observation, t_i's are the times of occurrence of independent reflections that are assumed to be random events taken from a non-homogeneous Poisson point process of rate $\lambda(t)$, ω_c is the center frequency of the transducer, $u(t)$ is the unit step function, a_i is the envelope amplitude, and ϕ_i is a random phase. Thus, the generalized power spectrum of both the in-phase, $r_s(t)$, and quadrature $r_c(t)$, components of the RF echo follows $1/f^\beta$ behavior with $\beta = 2(1 - \nu)$:

$$S(f) = \lambda E[a^2] \Gamma^2(\beta/2)(2\pi f)^{-\beta}$$

where $\Gamma(\cdot)$ is the gamma function, and $E[a^2]$ denotes the expected value of the square of the random amplitude a. Expressions for the power spectrum of the envelope were derived in Kutay *et al.*,[31] where, through numerical evaluation, they were shown to have a power-law form with exponent β_{env}. This result has also been justified based on extensive simulations.[31]

The following procedure was followed in order to estimate β_c, β_s, and β_{env}. The quadrature $r_c(t)$, the in-phase, $r_s(t)$, and the envelop $|r(t)|$ data were computed from the RF echo signal, $r(t)$, as follows:

$$r_c(t) = \cos(\omega_c t) r(t) + \sin(\omega_c t) \hat{r}(t)$$
$$r_s(t) = \cos(\omega_c t) \hat{r}(t) - \sin(\omega_c t) r(t)$$
$$|r(t)| = \sqrt{r_c^2(t) + r_s^2(t)}$$

where $\hat{r}(t)$ is the Hilbert transform of $r(t)$. Periodograms of $r_c(t)$, $r_s(t)$, and $|r(t)|$ were calculated, and estimates $\hat{\beta}_c$, $\hat{\beta}_s$, and $\hat{\beta}_{env}$ were found from these using:

$$(\hat{\beta}, \hat{A}) = \arg\min_{\beta, A} \left[\mu \frac{\int_{f_1}^{f_2} (\log(S(f)) + \beta \log(f) - \log A)^2 d(\log f)}{\int \log(S(f))^2 d(\log f)} \right.$$

$$\left. + (1 - \mu) \frac{\int_{f_1}^{f_2} (S(f) - A/f^\beta)^2 \, df}{\int S(f)^2 \, df} \right]$$

where $S(f)$ is the periodogram spectrum estimation, μ is a weight parameter (set to 0.1) and $[f_1 f_2]$ is the frequency band over which the matching is performed. More details regarding the estimation of the spectral exponents β_c, β_s, and β_{env} can be found in Kutay *et al.*[31,32]

The following process was followed in order to accommodate for depth localization and long data records. Analysis windows of 75 by 12 samples (2.4 mm lateral and 3 mm axial) were extracted. Samples were appended one to another to form a record of 900 samples. This appended sample was then used for signature extraction. Sixty overlapping windows were examined inside each tumor interior ROI. $\beta \equiv (\beta_c + \beta_s)/2$, and β_{env} were computed for each one of these windows. Features F7 and F8 are the mean values of β and β_{env} across these sixty windows.

4.3. *Morphological features*

When a radiologist examines a tumor, among other things, he or she considers its sonographic characteristic. The Breast Ultrasound Lexicon Subcommittee of an Expert Working standardized descriptors to characterize sonographic features of breast masses.[33] Based on this characterization, several morphological features were studied here to examine if they could automatically discriminate between benign and malignant lesions.

Mass shape can be used to characterize the type of the tumor. Spherical or globular lesions are likely to be simple cysts while elliptical tumors may be Fibroadenomas. On the other hand, if the tumor shape is undefined (neither circular nor elliptic) it may be invasive carcinoma. The circularity of the boundary of a mass can be evaluated by the ratio between the mean radial distance (the distance between the center-of-mass and points on the boundary) and its standard deviation — μ_r/σ_r.[34] In addition, the mass orientation may indicate that a mass is Fibroadenomas, for instance when its long axis parallels the skin. The characteristic of a mass boundary is also indicative regarding its nature. Usually a benign tumor shows a well defined interface between mass and tissue. The likelihood of malignancy increases as the mass boundary becomes less well defined. A thick indistinct echogenic rim may indicate an invasive cancer. Computing the contrast along the boundaries provides a feature that reflects the mass interface attribute.[22]

In general, the more irregular the shape of the mass is, the more likely it is to be a malignant tumor. Simple cysts typically have a smooth and regular shape, while

invasive cancer may have irregular margins, spiculations, and microlobulations. These characteristics can be quantified, for instance by computing the compactness of the mass, the ratio between the square mass perimeter, and mass area.[35] The shape's roughness can also indicate shape regularity, which is the difference between the fourth and second normalized moments of inertia of radial distance.[35] Another feature that measures shape smoothness is the standard deviation of the radial distance.[22] Finally, Tumor Boundary Roughness (TBR), which is the proportion of radial lines that are being crossed by more than one boundary point, may also indicate a malignant tumor.[35]

4.4. *Auxiliary features*

The chance of getting breast cancer increases as a woman grows older. A woman over the age of 60 is at the greatest risk. Therefore, using age as a feature could contribute significantly to diagnosis accuracy. A second feature in this class is the radiologist's decision. Accounting for this, the images were reviewed retrospectively by an experienced radiologist in order to determine the Level of Suspicion (LOS). The LOS feature takes discrete values on a 1–5 scale: 1 (normal), 2 (benign), 3 (indeterminate), 4 (probably malignant), and 5 (malignant). The characteristics of each mass used to determine the LOS included shape, margination, echogencity, internal homogeneity, presence of shadowing, calcification, and architectural distortion.

4.5. *Classifiers and diagnosis analysis*

Each of the features above can be used independently to perform diagnosis. But a fusion of all the features or a subset of them can improve the discriminatory capability as compared with each single feature considered independently. Classifiers are one way of doing this. Classifiers map a vector of features into a single representative scalar value. In this case study, a linear classifier and a quadratic classifier were considered.[8] The linear discriminant function defined as

$$f_l = \mathbf{v} \cdot \mathbf{x} + w$$

(and also referred to as the Fisher's discriminant function). The quadratic discriminant function defined as

$$f_q = (\mathbf{x} - \mathbf{m}_n)^T \mathbf{C}_n^{-1}(\mathbf{x} - \mathbf{m}_n) - (\mathbf{x} - \mathbf{m}_p)^T \mathbf{C}_p^{-1}(\mathbf{x} - \mathbf{m}_p)$$

where $\mathbf{x}$ is the observation vector (feature vector), $\mathbf{v} = 2(\mathbf{m}_p - \mathbf{m}_n)^T \mathbf{C}^{-1}$, and $w = \mathbf{m}_n^T \mathbf{C}^{-1} \mathbf{m}_n - \mathbf{m}_p^T \mathbf{C}^{-1} \mathbf{m}_p$. $\mathbf{m}_p$ and $\mathbf{m}_n$ are the means of the observation vectors for the positive and negative cases, respectively. $\mathbf{C}_p$ and $\mathbf{C}_n$ are the covariance of the observation vector for the positive and for the negative cases, respectively; and $\mathbf{C}$ is a combination of $\mathbf{C}_p$ and $\mathbf{C}_n$ weighted by the *a priori* probabilities, P_p and P_n. The *a priori* probabilities were estimated based on the positive and negative case frequencies: $\mathrm{P}_p = \mathrm{N}_p/(\mathrm{N}_p + \mathrm{N}_n)$ and $\mathrm{P}_n = \mathrm{N}_n/(\mathrm{N}_p + \mathrm{N}_n)$.

When estimating the above means and covariance of the observation data, the case study analysis allowed for the possibility of a partial dataset. A partial dataset occurs when some feature values are not present in all the feature vectors. Typically, an incomplete dataset is allowed when features may be rejected because of either technical or situational reasons (outliers).

Following is the description of the estimation computation of $\mathbf{m}_p$, $\mathbf{m}_n$, $\mathbf{C}_p$, and $\mathbf{C}_n$ under this condition of an incomplete dataset. Let S be the full set, either a positive or a negative case set, of N observation vectors. Let S_i be the subset of observation vectors for which the ith feature is present. Similarly, let S_{ij} be the subset of observation vectors for which the ith feature and the jth feature are present. Thus: $S_{ij} = S_i \cap S_j$. The sample mean and sample covariance can then be computed from the incomplete dataset as follows:

$$\hat{\mathbf{m}}_i = \frac{1}{N_i} \sum_{n \in S_i} \mathbf{x}_{i,n} \quad N_i = O(S_i)$$

$$\hat{\mathbf{C}}_{ij} = \frac{1}{N_{ij} - 1} \sum_{n \in S_{ij}} (\mathbf{x}_{i,n} - \hat{\mathbf{m}}_i)(\mathbf{x}_{j,n} - \hat{\mathbf{m}}_j); \quad N_{ij} = O(S_{ij}).$$

Next, when computing the discriminant function, either linear or quadratic, in the equations above, only the portions of $\hat{\mathbf{m}}_p$, $\hat{\mathbf{m}}_n$, $\hat{\mathbf{C}}_p$ and $\hat{\mathbf{C}}_n$ that corresponded to the feature values which were available in the feature vector being classified, were substituted.

The next step was to evaluate the classifier's output. An ROC curve was used for performance evaluation. The summary diagnostic performance index is the area under the ROC curve. In performance evaluation based on the empirical ROC, the confidence interval of A_z, $(\beta - \alpha)$ was estimated by using a bootstrap technique.[36] One thousand sets of positive and negative case data were formed by sampling with replacement from the experimental set. The empirical ROC and A_z for each set was then calculated. The 95% confidence interval for A_z was derived from this collection of measurements by finding the minimum interval $[\alpha, \beta]$ for which $P_{A_z}(\alpha \leq A_z < \beta) = 0.95$.

Though the estimated area under the empirical ROC is an unbiased estimate of A_z when the sample features are statistically independent,[3] this is not the case for estimates of discriminant function performances when the same samples are used for testing and training. Two techniques were used to monitor the bias: the resubstitution method, which has an optimistic bias (estimated A_z higher than the actual), and the leave-one-out method, which has a pessimistic bias.[37] In resubstitution, the same set of cases is used for training (computing classifier parameters) and for testing (estimating the ROC). In a leave-one-out analysis a subset of all but one case is used for training. The classifier output, f_l or f_q, is then computed for the case that was taken out. This analysis is repeated for all the cases. The result is a generated table of statistics that are used for the ROC analysis. Both the resubstitution

and leave-one-out methods were embedded in the bootstrap procedure to obtain confidence intervals for A_z.

4.6. *Experimental results*

The individual features and the combined subsets of individual features were evaluated to determine their capability to distinguish between malignant and benign cases. When fusing a set of features into a single representative feature, using a discriminant function, the case study analysis examined four different cases: (1) linear discriminant with resubstitution, (2) linear discriminant with leave-one-out, (3) quadratic discriminant with resubstitution, and (4) quadratic discriminant with leave-one-out.

Table 2 shows the results for the RF-signal-based features, containing the estimate A_z, its 95% confidence interval, $[\alpha\text{-}\beta]$, and the FPF at a TPF of 97.5%. The reported values were calculated with both empirical ROC and with a bootstrap. The observed values of A_z range from 0.61 to 0.91. The values of FPF range from 0.91 to 0.32. Table 3 shows the results of a fusion of all the RF-signal-based features. It can be seen that when employing a linear classifier and when using a resubstitution (leave-one-out) analysis method the A_z is 0.96 (0.94). The confidence intervals in this case show a substantial overlap. The corresponding values for the quadratic classifier are 0.97 (0.87) and the union of the confidence intervals ranges from 0.79 to 1. This indicates that even though the performance of the quadratic classifier may be superior, it is impossible at this stage to prove this with the available study population.

The performance of both classifiers exceeds that of the individual features. For instance, based on the leave-one-out analysis, the linear classifier has A_z in the

Table 2. A_z and confidence interval, $(\beta\text{-}\alpha)$, for individual RF-signal based features.

	F1	F2	F3	F4	F5	F6	F7	F8
A_z	0.75	0.73	0.61	0.76	0.91	0.84	0.85	0.85
α	0.63	0.62	0.46	0.63	0.85	0.74	0.79	0.78
β	0.85	0.83	0.74	0.88	0.97	0.93	0.91	0.91
FPF (TPF = 0.975)	0.83	0.80	0.91	0.90	0.32	0.69	0.58	0.63

Table 3. A_z and confidence interval, $(\beta\text{-}\alpha)$, for combined RF-signal based features (F1–F8).

	Linear classifier		Quadratic classifier	
	Resub.	Leave-one-out	Resub.	Leave-one-out
A_z	0.96	0.94	0.97	0.87
α	0.94	0.91	0.95	0.79
β	0.99	0.98	1.00	0.95
FPF (TPF = 0.975)	0.16	0.22	0.26	0.89

range of 0.91 to 0.98 at a significance level of 0.05. The benefit of fusion is especially evident when considering the reduction in FPF. The probability of performing a biopsy on a benign tumor ranges between 32% and 91% for independent RF-signal based features. Nevertheless, when combining the features with a linear classifier and applying resubstitution (leave-one-out), the probability for applying unnecessary biopsy is reduced to 16% (22%).

Similarly, Table 4 shows the results for classification when based on morphological features. In this case A_z ranges between 0.54 and 0.7. For the linear classifier, the resubstitution and leave-one-out A_z are 0.78 and 0.74. The corresponding values for the quadratic classifier are 0.83 and 0.65. At least with the data in this case study, it is apparent that the morphological features, independently or combined, are inferior to the RF-signal-based features. Table 5 shows the results of a fusion of all the morphological features.

The results of classification based on the auxiliary features are shown in Tables 6 and 7. Table 6 shows that based on the patient's age the classification yields an A_z of 0.81. Classification based on a radiologist's diagnosis yields an A_z of 0.85. The fusion of the two features results in an A_z of 0.9 (0.86), in the case where linear

Table 4. A_z and confidence interval, $(\beta\text{-}\alpha)$, for morphological features.

	F9	F10	F11	F12	F13	F14	F15	F16
A_z	0.70	0.66	0.54	0.56	0.63	0.57	0.59	0.55
α	0.60	0.56	0.42	0.44	0.53	0.45	0.51	0.44
β	0.81	0.76	0.64	0.68	0.74	0.67	0.72	0.65
FPF (TPF = 0.975)	0.92	0.92	0.94	0.96	0.96	0.93	0.92	0.95

Table 5. A_z and confidence interval, $(\beta\text{-}\alpha)$, for combined morphological features (F9–F16).

	Linear classifier		Quadratic classifier	
	Resub.	Leave-one-out	Resub.	Leave-one-out
A_z	0.78	0.74	0.83	0.65
α	0.69	0.63	0.75	0.51
β	0.88	0.85	0.91	0.76
FPF (TPF = 0.975)	0.9	0.95	0.73	0.96

Table 6. A_z and confidence interval, $(\beta\text{-}\alpha)$, for auxiliary features.

	F9	F10
A_z	0.81	0.85
α	0.72	0.77
β	0.89	0.93
FPF (TPF = 0.975)	0.76	0.79

Table 7. A_z and confidence interval, $(\beta\text{-}\alpha)$, for combined auxiliary features (F17–F18).

	Linear classifier		Quadratic classifier	
	Resub.	Leave-one-out	Resub.	Leave-one-out
A_z	0.90	0.86	0.90	0.89
α	0.83	0.81	0.84	0.81
β	0.96	0.95	0.96	0.96
FPF (TPF = 0.975)	0.73	0.79	0.71	0.82

Table 8. A_z and confidence interval, $(\beta\text{-}\alpha)$, for combined features (F5–F8, F17, F18).

	Linear classifier		Quadratic classifier	
	Resub.	Leave-one-out	Resub.	Leave-one-out
A_z	0.97	0.97	0.98	0.96
α	0.95	0.94	0.97	0.92
β	1.0	0.99	1.0	0.99
FPF (TPF = 0.975)	0.28	0.34	0.11	0.51

classifier with resubstitution (leave-one-out) is used, and an A_z of 0.9 (0.89) when a quadratic classifier with resubstitution (leave-one-out) is used.

Finally, fusion of a subset of the features presented above was analyzed. A subset of features with $A_z > 0.8$ was selected (F5–F8, F17, F18). Table 8 shows the results.

References

1. J. P. Egan, *Signal Detection Theory and ROC Analysis* (Academic Press, New York, 1975).
2. J. A. Swets, R. M. Dawes and J. Monahan, in *Scientific American* (2000), pp. 82–87.
3. D. Bamber, The area above the ordinal dominance graph and the area below the receiver operating characteristic graph, *Journal of Mathematical Psychology* **12** (1975) 387–415.
4. C. E. Metz, ROC methodology in radiologic imaging, *Investigative Radiology* **21** (1986) 720–733.
5. B. S. Garra *et al.*, Quantitative ultrasonic-detection of parenchymal structural-change in diffuse renal-disease, *Investigative Radiology* **29** (1994) 134–140.
6. S. Gefen *et al.*, ROC analysis of ultrasound tissue characterization classifiers for breast cancer diagnosis, *IEEE Transaction on Medical Imaging* **22** (2003) 170–177.
7. D. B. Kopans, *Breast Imaging* (Lippincott Raven, Philadelphia - New York, 1997).
8. S. Theodoridis and K. Koutroumbas, *Pattern Recognition* (Academic Press, 1999).
9. A. Thomas, The invisible light, *Journal of The Radiology History and Heritage Charitable Trust* **15** (2001).
10. U. Bick, F. Diekmann, S. Grebe, B. Hamm and L. Charit'e, in *ECR* (Germany, 2000).
11. P. Malhotra *et al.*, in *SPIE* (University of South Florida, 2002).
12. D. R. Chen, R. F. Chang and Y. L. Huang, Computer-aided diagnosis applied to US of solid breast nodules by using neural networks, *Radiology* **213** (1999) 407–412.

13. D. Chen, R. Chang and Y. Huang, Breast cancer diagnosis using self-organizing map for sonograpy, *Ultrasound in Med. & Bio.* **26** (2000) 405–411.

14. D. Chen *et al.*, in *Seminars in Ultrasound, CT, and MR* (2000), pp. 308–316.

15. D. Chen, W. Kuo, R. Chang, W. Moon and C. Lee, Use of the bootstrap technique with small training sets for computer-aided diagnosis in breast ultrasound, *Ultrasound in Med. & Bio.* **28** (2002) 897–902.

16. K. Horsch, M. L. Giger, L. A. Venta and C. J. Vyborny, Automatic segmentation of breast lesions on ultrasound, *Med. Phys.* **28** (2001) 1652–1659.

17. K. Horsch, M. Giger, L. Venta and C. Vyborny, Computerized diagnosis of breast lesions on ultrasound, *Med. Phys.* **29** (2002) 157–164.

18. Y. Chou *et al.*, Stepwise logistic regression analysis of tumor contour features for breast ultrasound diagnosis, *Ultrasound in Med. & Bio.* **27** (2001) 1493–1498.

19. F. Kitaoka *et al.*, Internal echo histogram examination has a role in distinguishing malignant tumors from benign masses in the breast, *Clinical Imaging* **25** (2001) 151–153.

20. F. Lefebvre, M. Meunier, F. Thibault, P. Laugier and G. Berger, Computerized ultrasound B-scan characterization of breast nodules, *Ultrasound in Med. & Bio.* **26** (2000) 1421–1428.

21. R. Sivaramakrishna, K. Powell, M. Lieber, W. Chilcote and R. Shekhar, Texture analysis of lesions in breast ultrasound images, *Computerized Medical Imaging and Graphics* **26** (2002) 303–307.

22. R. Haralik, A measure of circularity of digital figures, *IEEE Trans. Systems, Man Cybernetics SMC-3* (1974) 453–465.

23. P. M. Shankar *et al.*, Classification of ultrasonic B mode images of breast masses using the Nakagami distribution, *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.* **48** (2001) 569–580.

24. V. A. Dumane and P. M. Shankar, Use of frequency diversity and Nakagami statistics in ultrasonic tissue characterization, *IEEE Trans. Ultrason., Ferroelect., Freq. Cont.* **48** (2001) 1139–1146.

25. K. D. Donohue, F. Forsberg, C. W. Piccoli and B. B. Goldberg, Analysis and classification of tissue with scatterer structure templates, *IEEE Trans. on Ultrasonic, Ferroelectrics and Frequency Control* **46** (1999) 300–310.

26. K. D. Donohue and L. Huang, in *IEEE Proceedings of International Conference on Acoustics, Speech and Signal Processing* (Salt Lake City, UT, 2001).

27. K. D. Donohue, L. Huang, T. Burks, F. Forsberg and C. W. Piccoli, Tissue classification with generalized spectrum parameters, *Ultrasound in Medicine and Biology* **27** (2001) 1505–1514.

28. L. Huang, K. D. Donohue, V. Genis and F. Forsberg, Duct detection and wall spacing estimation in breast tissue, *Ultrasonic Imaging* **22** (2000) 137–152.

29. A. P. Petropulu, T. Golas and G. Viswhanathan, in *IEEE Int. Symp. on Time-Freq. Time-Scale Anal.* (Pittsburgh PA, 1998).

30. A. P. Petropulu, J. C. Pesquet, X. Yang and J. Yin, Power-law shot noise and its relationship to long-memory α-stable processes, *IEEE Truns. on Sig. Process.* **48** (2000) 1883–1892.

31. A. M. Kutay, A. P. Petropulu and C. W. Piccoli, On modeling biomedical ultrasound RF echoes using a power-law shot noise model, *IEEE Trans. Ultrason. Ferroelect. Freq. Contr.* **48** (2001) 953–968.

32. A. M. Kutay, A. P. Peropulu and C. Piccoli, Tissue characterization based on the power-law shot noise model, *to appear in Pattern Recognition Letters* (2002).

33. C. R. B. Merritt, B. B. Goldberg and E. B. Mendelson, *Breast Imaging Reporting and Data System: Ultrasound* (Thomas Jefferson University Hospital, Philadelphia), pp. 1–21.
34. L. Shen, R. M. Rangayyan and J. E. Desautels, Application of shape analysis to mammographic calcifications, *IEEE Trans. on Medical Imaging* **13** (1994) 263–274.
35. J. A. Mann, All measuring of the roughness of interfaces, *Chemometrics Intel. Lab. Systems* **12** (1991) 169–180.
36. B. Efron and R. J. Tibshirani, *An Introduction to the Bootstrap* (Chapman & Hall, 1993).
37. K. Fukunaga, *Introduction to statistical Pattarn Recognition* (Academic Press, 1990).

PARAMETRIC SHAPE RECONSTRUCTION IN INVERSE PROBLEMS: FUNDAMENTAL PERFORMANCE BOUNDS AND ALGORITHMS

JONG CHUL YE

Department of BioSystems
Korea Advanced Institute of Science and Technology (KAIST)
373-1 Gusong-Dong, Yuseong-Gu, Daejeon 305-701, Korea
jongcye@biosys.kaist.ac.kr

We address the problem of estimating object boundaries from noisy measurements in inverse problems, when the boundaries are parameterized by a finite number of unknown variables. We review the existing reconstruction algorithms from different imaging modalities and discuss fundamental performance bounds of the parametric approaches using Cramér-Rao lower bounds and asymptotic global confidence regions.

Keywords: Parametric shape reconstruction; the Cramér-Rao inequality; asymptotic global confidence region.

1. Introduction

The problem of estimating object boundaries from noisy measurements is encountered in applications such as computed tomography (CT), image deconvolution, synthetic aperture radar (SAR), and nonlinear inverse scattering. If there exist sufficient number of measurements, the boundary can be extracted directly from the reconstructed images. However, in many applications there is generally insufficient information to recover details of the boundary. This situation is typical of an ill-posed image reconstruction problem.

In such problems, the boundary is represented mathematically as a curve in $\mathbb{R}^2$ (or a surface in $\mathbb{R}^3$) and is often parameterized by a finite number of unknown variables. Such a parametric formulation (for instance using B-splines[1] or Fourier descriptors[2,3]) is a first step towards constructing a stable boundary estimation algorithm. This type of parametrization imposes smoothness on the reconstructed boundary and provides the basis for, e.g. an efficient maximum likelihood estimation algorithm.

Several estimation algorithms have been investigated for different imaging modalities. Kirsch formulated the inverse acoustic scattering problems in parametric framework and implemented Gauss–Newton type estimation algorithms to estimate the target boundaries.[4] Zheng and Doerschuck investigated 3-D viral reconstruction from averaged Fourier transform magnitude by parameterizing the virus structure with spherical harmonics.[39] Most recently, Schmid *et al.* formulated the parametric

reconstruction problem in a more general framework such that they can simultaneously estimate both parameter and the model orders.[40] In all these approaches, superior reconstruction of the boundaries have been obtained even from severely under-sampled measurement data.

The superior reconstruction performance of the parametric formulation has been also theoretically verified by computing the fundamental bounds — the Cramér–Rao lower bound (CRB).[5] The CRB provides an unbeatable performance limit for any unbiased estimator, and hence can be used to investigate the fundamental limits of parameter estimation problems, or as a baseline for assessing the performance of a specific estimator.[6] The CRB can be used for optimum design of the measurement system, e.g. by selecting sampling points to minimize a functional of the CRB matrix, such as its trace or determinant.[7] While CRB's are available for estimation of signal parameters such as target location,[8–11] direction-of-arrival (DOA),[7,12–14] and size and orientation of a scatterer,[11,15,16] only recently has this type of analysis been conducted for estimation of target shapes in general inverse problems.[5,17]

The CRB's computed for shape estimation can also be used to compute a global uncertainty region around the boundary,[17] providing an easily interpreted geometric display of boundary uncertainty. A related idea has been applied to tomographic reconstruction by Hanson *et al.*[18] using Monte-Carlo simulations, which are limited to a particular estimator and time-consuming to construct. In contrast, the global confidence region can be easily and quickly constructed using the CRB covariance matrix, even before the construction of an estimator is attempted.

This chapter is organized as follows. Section 2 formulates the parametric shape estimation problem formally and introduces some real applications. In Sec. 3, we briefly review the existing maximum likelihood parametric shape estimation algorithms. Section 4 discusses importance of the Cramér–Rao lower bound and the global confidence regions to evaluate the fundamental performance of parametric shape estimation, and Sec. 5 derives the general formula of the domain derivatives — the essential mathematical tools for fundamental performance analysis. Section 6 extends our theory to an object composed of multiple domains. The CRB is then computed for several illustrative examples in Sec. 7. Conclusions are presented in Sec. 8.

2. The Shape Estimation Problem

Consider a real-valued image f consisting of a constant-valued 2-D object and a known background density $f_2(x, y)$:

$$f(x, y) = \begin{cases} f_1, & (x, y) \in D \\ f_2(x, y), & (x, y) \in \mathbb{R}^2 \backslash D \end{cases}. \tag{1}$$

The intensity f_1 and region D are unknown, whereas $f_2(x, y)$ is known for all $(x, y) \in \mathbb{R}^2$. This scenario models an object of constant but unknown intensity and

unknown shape, partly occluding (or replacing) a *known* background. This situation arises in applications requiring the imaging of localized changes or anomalies in a reference object, or the imaging of objects on a homogeneous background. The object is thus completely defined by its image value f_1 and its boundary $\Gamma = \partial D$. The support set D need not be a connected region, so the formulation includes the case of multiple objects.

2.1. *Parametric boundary model*

Since the estimation of Γ from a limited number of noisy samples is generally an ill-posed inverse problem, a possible remedy is to represent the boundary Γ as a known function with a finite number of unknown parameters:

$$\Gamma = \{\mathbf{s}(u; \boldsymbol{\phi}), u \in I\}, \tag{2}$$

where $\boldsymbol{\phi} = (\phi_1, \ldots, \phi_K) \in \mathbb{R}^K$ is an unknown parameter vector, and $I \subset \mathbb{R}$ an interval. Because object boundaries correspond to closed contours, I is a closed interval, and $\mathbf{s}(u; \boldsymbol{\phi})$ a periodic function of u, with period equal to the length $|I|$. In particular, we use the series expansion model

$$\mathbf{s}(u; \boldsymbol{\phi}) = [x(u; \boldsymbol{\phi}), y(u; \boldsymbol{\phi})]^T = \sum_{i=1}^{K} \phi_i \mathbf{b}_i(u), \quad u \in I, \tag{3}$$

where $\mathbf{b}_i(u) \in \mathbb{R}^2$ is the i-th basis function. Parameterizations such as Fourier descriptors (FD),[2,3] B-splines[1] and wavelet descriptors[19,20] are special cases of this model and have been widely used for shape representation.

Throughout the paper, we assume that the boundary Γ is a 1-manifold in $\mathbb{R}^2$, or of class C^1. In practice, this regularity condition is not too restrictive and implies that the boundary Γ is sufficiently smooth without any cusps and crossings. Equivalently, a parametrized-manifold $\mathbf{s}(u)$, $u \in I$ should be continuously differentiable with respect to u, its inverse should be continuous to avoid any crossing, and $\dot{x}(u)^2 + \dot{y}(u)^2 \neq 0$ for all $u \in I$ to prevent any cusps, where $\dot{x}(u)$ and $\dot{y}(u)$ denote the derivatives of the x and y components of $\mathbf{s}(u)$ with respect to u. These regularity conditions are in addition to the earlier condition that $\mathbf{s}(u)$, $u \in I$ be periodic with period $|I|$.

Under the series expansion model (3), in order to satisfy these regularity conditions, the basis function $\mathbf{b}_i \colon I \to \mathbb{R}^2$, $i = 1, \ldots, K$ should be continuously differentiable. The conditions for avoiding intersecting boundaries and cusp are, however, more difficult to impose for each basis function $\mathbf{b}_i$, because these are global properties contributed by the linear combination of all basis functions. Therefore, we will assume that the parameters ϕ_i, $i = 1, \ldots, K$ are chosen such that the resultant boundary $\Gamma = \{\mathbf{s}(u; \boldsymbol{\phi}), u \in I\}$ does not have crossings and cusps.

2.2. *Forward problem*

Let $g = \mathcal{H}f$ be a general linear integral transformation of f, defined by

$$g(s,t) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x,y) h(x,y,s,t) \, dx \, dy, \quad (s,t) \in \Omega \qquad (4)$$

where $h : \mathbb{R}^2 \times \Omega \to \mathbb{R}$ is a known kernel, and Ω a subset of $\mathbb{R}^2$. We assume that f and h satisfy appropriate conditions so that g is well defined. For example, one such set of conditions is that f is bounded, and $h(\cdot, \cdot, s, t)$ is absolutely integrable for any $(s,t) \in \Omega$. Examples of the linear operator (4) are as follows.

2.2.1. *Computed tomography*

Consider tomographic reconstruction from line-integral projections or the samples of the 2-D Radon transform. The Radon transform g of a 2-D function f, is defined as the collection of line integrals of f along lines, indexed by these lines' angle s from the x-axis and at distance t from the origin:

$$g(s,t) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x,y) \delta(x \cos(s) + y \sin(s) - t) \, dx \, dy,$$
$$-\infty < t < \infty, \quad 0 \le s < \pi, \qquad (5)$$

where $\delta(\cdot)$ is the Dirac impulse. The function f is assumed to be such that (5) is well-defined.[21]

This reconstruction problem (and several variations thereof[21]) is of interest in numerous applications,[22] of which medical x-ray CT is the best known. It also arises in non-destructive evaluation of metal castings for the presence of cracks or bubbles, or monitoring nuclear reactor cores. In geophysics it may be an appropriate model for estimating the shape of an underground reservoir.

2.2.2. *Fourier imaging*

The second problem we consider is Fourier imaging.[23,24] This problem arises in applications such as synthetic aperture radar (SAR),[25,26] diffraction tomography (DT),[8] magnetic resonance imaging (MRI),[27] and other image formation modalities.[23] In Fourier imaging, measurements are samples of the 2-D Fourier transform:

$$g(s,t) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x,y) \, e^{-j2\pi(sx+ty)} dx \, dy, \quad (s,t) \in \mathbb{R}^2. \qquad (6)$$

The same sufficient conditions as in the computed tomography problem are applicable to the Fourier imaging problem too. As in the tomography problem, the known background $f_2(x,y)$ in these applications may be obtained from baseline reference scans. For example, in the case of transient change detection (response to a stimulus in functional MRI, or the appearance of a target in SAR) an accurate image of the static background $f_2(x,y)$ may be easily obtained, because of the essentially

unlimited time available for data acquisition. In other applications, such as imaging on a sky background, a constant background may be assumed.

2.2.3. *Deconvolution*

The last problem we consider is deconvolution — a ubiquitous problem in all areas of science and engineering. Suppose the image $f(x, y)$ of (1) is blurred with a shift-invariant point spread function $h(x, y)$. The noiseless blurred image $g(x, y)$ is given by

$$g(s, t) = (h * *f)(s, t) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) h(s - x, t - y) \, dx \, dy,$$

$$(s, t) \in R_f \subset \mathbb{R}^2 \tag{7}$$

where $**$ is the 2-D convolution, and R_f denotes the image domain, respectively. For the convolution $g(s, t)$ to be well defined, it is sufficient that one of f and h be bounded and the other absolutely integrable, which is satisfied in all practical applications. For instance, if the point spread function is Gaussian with width parameter ν, then (7) takes the form

$$g(s, t) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{1}{2\pi\nu^2} \, e^{-\frac{(s-x)^2 + (t-y)^2}{2\nu^2}} f(x, y) \, dx \, dy. \tag{8}$$

Deconvolution is a key problem in imaging systems and seismic inversion.[23] As in the other imaging problems, the known background $f_2(x, y)$ in these applications may be obtained from baseline reference scans. For example, in the case of change detection in an optical surveillance problem, an accurate image of the static background $f_2(x, y)$ may be easily obtained from the previous scenes.

3. Parametric Shape Estimation Algorithms

3.1. *Maximum likelihood estimation*

Suppose $g(s, t)$ is sampled at a finite number M of positions $\{s_m, t_m\}_{m=1}^{M}$. The estimation problem we consider is to estimate the object boundary Γ from noisy measurements $\{y_m\}_{m=1}^{M}$ of the samples $g_m = g(s_m, t_m)$, $m = 1, \ldots, M$. The measurement model is specified by a conditional probability density function (pdf) $p_{Y|\mathbf{g}}(\mathbf{y}|\mathbf{g})$, where $\mathbf{y} = [y_1, \ldots, y_M]^T$ denotes a particular realization of the random vector Y and $\mathbf{g} = [g_1, \ldots, g_M]^T$ is the noise-free sample vector. In view of the parametric representation (2) and (3) of the object, we denote the m-th noise-free sample $g(s_m, t_m)$ by $g_m = g_m(\boldsymbol{\theta})$, where

$$\boldsymbol{\theta} = [f_1, \phi_1, \ldots, \phi_K]^T \tag{9}$$

in order to show explicitly the dependence on the unknown parameter vector. Then, the statistics of Y can be described by

$$p_{Y|\mathbf{g}}(\mathbf{y}|\mathbf{g}) = p_{Y|\mathbf{g}(\boldsymbol{\theta})} (\mathbf{y}|\mathbf{g}(\boldsymbol{\theta})), \tag{10}$$

which we denote, for brevity, by $p(\mathbf{y}|\boldsymbol{\theta})$. Note that this formulation includes the case where the observation involves a nonlinear transformation of $\mathbf{g}$. For example, in low dose X-ray imaging or positron emission tomography, the measurement y_m is Poisson distributed with rate $\lambda_m = \lambda_T \exp(-g_m)$.

Now, for a given k, one can determine the maximum likelihood parameter estimation of the parameters, $\hat{\boldsymbol{\theta}}$, as the solution of the following optimization problem:

$$\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta} \in \mathbb{R}^{k+1}} \log p(\mathbf{y}|\boldsymbol{\theta}) \tag{11}$$

where $\log p(\mathbf{y}|\boldsymbol{\theta})$ is the log-likelihood function.

In order to implement (11), Schmid *et al.*[40] invoked the function *fminunc.m* from MATLAB optimization toolbox. This function performs an unconstrained nonlinear optimization based on a quasi-Newton method. Initial experiments have shown that the algorithm exhibits instabilities: it is sensitive to initial guess, especially for low signal-to-noise ratio (SNR). In this case, the algorithm often results in a self intersecting boundary. To avoid the problem, they proposed to include an additional penalty term in (11) that penalizes rough boundaries:

$$\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta} \in \mathbb{R}^{k+1}} \left[\log p(\mathbf{y}|\boldsymbol{\theta}) + \lambda \int_I \left([\dot{x}(u)]^2 + [\dot{y}(u)]^2 \right)^{1/2} dt \right], \tag{12}$$

where λ is the regularization parameter. Consider a sequence of regularization problems in the form of (12) parameterized by a sequence of parameters λ_l gradually diminishing to zero. Under this setting, the solution to the problem in (12) converges to the solution of the non-regularized problem in (11), as $\lambda_l \to 0$. They implemented the procedure described above recursively, that is, the solution to the optimization problem (12) with larger λ is used as an initial guess to the optimization problem with a smaller value of λ. They showed that this lead to a stable solution and helps avoid the self intersecting boundaries.

3.2. *Complexity regularized estimator*

One hidden, but important, assumption of the maximum likelihood estimation framework (11) is that the "optimum" model order K is known *a priori*. However, in real application scenario, the correction model order K usually needs to be estimated from the data. Hence, the remaining problem, is to determine the "correct" order K of the parametric boundary model (3). In general, increasing K, allows to capture increasing spatial detail in the represented boundary. However, with limited and noisy data, increasing K beyond a certain point leads to diminishing returns, because of increase in the estimation error of more parameters from a fixed amount of data. Thus, once the parametric model has been picked for boundary, regularization reduces to selection of the model order K.

Schmid *et al.*[40] used Rissanen's MDL criterion to determine the optimum model order K. The MDL criterion is stated as follows:

$$\hat{K}(\mathbf{y}) = \arg \max_{K=1,2,\dots} \left[\log p(\mathbf{y}|\boldsymbol{\theta}^{(k+1)}) + L(\boldsymbol{\theta}^{(k+1)}) \right] \tag{13}$$

where $L(\boldsymbol{\theta}^{(k+1)})$ is the penalty term equal to the code length of $k+1$ dimensional parameter vector $\boldsymbol{\theta}^{(k+1)}$.

A variety of coding scheme can be considered to describe the penalty term in (13). Under the practical condition of a large sample size, the description length of each parameter can be approximated by $1/2 \log M$, with M being the number of data samples. Assuming that every model parameter is equally penalized, this leads to $L(\boldsymbol{\theta}^{(k+1)}) \sim k + 1/2 \log M$. Schmid *et al.* showed that the model order estimation using (13) showed a good trade-off between bias and variance of shape estimation.[40]

4. Fundamental Performance Limit for Parametric Shape Estimation

4.1. *The Cramér–Rao inequality*

As shown in (11), given data $\mathbf{y}$ and its statistical description by $p(\mathbf{y}|\boldsymbol{\theta})$, the shape reconstruction problem has been reduced to a statistical parameter estimation problem. As is well-known, the variance of an unbiased estimator of $\boldsymbol{\theta}$ is subject to a fundamental limit — the Cramér–Rao lower bound (CRB). Moreover, under appropriate regularity conditions, the maximum likelihood estimator (MLE) achieves the CRB asymptotically.[6] Hence, the CRB is not only a lower bound, but is also useful for predicting the large sample or high SNR performance of the MLE, or other asymptotically efficient estimators. The accuracy of the shape estimate can in turn be assessed from the accuracy of the parameter estimate. In particular, one can then construct global confidence regions for the object boundary.[17]

Here, we first concentrate on the derivation of a procedure for computing the CRB. According to the Cramér–Rao inequality, subject to some regularity conditions on the conditional pdf $p_{Y|\boldsymbol{\theta}}$, the $(K+1) \times (K+1)$ covariance matrix of the estimation error $\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}$ for the unknown parameter $\boldsymbol{\theta}$ is bounded from below as[6]

$$\text{Cov}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \geq \mathbf{C}_{\boldsymbol{\theta}} \triangleq (\mathbb{I}_{\boldsymbol{\theta}})^{-1}, \tag{14}$$

for *any unbiased* estimate $\hat{\boldsymbol{\theta}}$ of $\boldsymbol{\theta}$. Here, the matrix inequality notation $A \geq B$ indicates that $A - B$ is positive semidefinite, with A and B being Hermitian positive semidefinite matrices. In (14), the Fisher information matrix, $\mathbb{I}_{\boldsymbol{\theta}}$, is the $(K+1) \times (K+1)$ matrix

$$\mathbb{I}_{\boldsymbol{\theta}} = E \left[\nabla_{\boldsymbol{\theta}} \ln p(\mathbf{y}|\boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}}^T \ln p(\mathbf{y}|\boldsymbol{\theta}) \right] \tag{15}$$

where $\ln p(\mathbf{y}|\boldsymbol{\theta})$ is the log-likelihood function, and $\nabla_{\boldsymbol{\theta}}$ denotes gradient with respect to $\boldsymbol{\theta}$.

For any pdf $p(\mathbf{y}|\boldsymbol{\theta})$ for which the Fisher information matrix is well-defined, it follows from the chain rule that the entries of $\mathbb{I}_{\boldsymbol{\theta}}$ in (15) are (possibly nonlinear) functions of $g_m(\boldsymbol{\theta})$ and the derivatives $\frac{\partial g_m(\boldsymbol{\theta})}{\partial \theta_i}$, $i = 1, \dots, K+1$, $m = 1, \dots, M$. For examples, if the noisy measurement Y_m is a Gaussian random variable with mean $g_m(\boldsymbol{\theta})$ and variance σ_m^2, under the additional assumption of independence of measurements at different locations, the (i, j)-th element of the Fisher information matrix is given by

$$(\mathbb{I}_{\boldsymbol{\theta}})_{i,j} = \sum_{m=1}^{M} \frac{1}{\sigma_m^2} \frac{\partial g_m(\boldsymbol{\theta})}{\partial \theta_i} \frac{\partial g_m(\boldsymbol{\theta})}{\partial \theta_j}, \quad i, j = 1, \dots, K+1, \tag{16}$$

whereas for the complex circular Gaussian noise with mean zero and variance $E|Y_m - g_m|^2 = 2\sigma_m^2$ we have

$$(\mathbb{I}_{\boldsymbol{\theta}})_{i,j} = \sum_{m=1}^{M} \frac{1}{\sigma_m^2} \mathrm{Re}\left[\left(\frac{\partial g_m(\boldsymbol{\theta})}{\partial \theta_i} \right)^{*} \frac{\partial g_m(\boldsymbol{\theta})}{\partial \theta_j} \right], \quad i, j = 1, \dots, K+1, \tag{17}$$

where $Re[\cdot]$ denotes the real part and the superscript $*$ denotes the conjugation.

While (16)–(17) appear simple, techniques for computing the derivatives $\left\{ \frac{\partial g_m(\boldsymbol{\theta})}{\partial \theta_i} \right\}$ for models of the form (4) and (2)–(3) have not been studied in the literature, except for special cases.[11,16,28] In Sec. 5, we will develop a general technique to compute those quantities in a generic linear inverse problem.

4.2. *From CRB's to global confidence regions*

In practice, because $\mathbf{s}(u; \boldsymbol{\theta})$ describes the geometry of an object, one is interested in assessing the quality of estimates of $\mathbf{s}(u; \boldsymbol{\theta})$ in easily interpreted geometric terms. Rather than the quality of estimates of $\boldsymbol{\theta}$ itself, what is needed is a global quality measure for the entire boundary $\{\hat{\mathbf{s}}(u; \boldsymbol{\theta}), \forall\, u \in I\}$. The CRB $\mathbf{C}_{\boldsymbol{\theta}}$ computed by the techniques of this chapter can be used as described in Ref. 17, to construct small-size *global confidence regions* in the asymptotic regime where the estimate is unbiased, efficient, and Gaussian.

More specifically, under appropriate regularity conditions,[29,30] the maximum-likelihood estimator (MLE) of (11) is asymptotically unbiased, efficient and Gaussian-distributed (or best asymptotically normal (BAN))

$$\hat{\boldsymbol{\theta}} \sim \mathcal{N}(\boldsymbol{\theta}, \mathbf{C}_{\boldsymbol{\theta}}) \tag{18}$$

where the $(K+1) \times (K+1)$ covariance matrix $\mathbf{C}_{\boldsymbol{\theta}} = E[\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}][\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}]^T$ is the inverse of the Fisher information matrix. Examples of asymptotics in the observation model include collecting an infinite number of snapshots, or an increasing number of sensor elements, with independent sensor noise realizations.

Similar properties apply to smooth functions of $\boldsymbol{\theta}$ such as $\mathbf{s}(u; \boldsymbol{\theta})$. In particular, by the invariance property of the MLE, the asymptotic distribution of $\mathbf{s}(u; \hat{\boldsymbol{\theta}})$ is likewise BAN

$$\mathbf{s}(u; \hat{\boldsymbol{\theta}}) \sim \mathcal{N}(\mathbf{s}(u; \boldsymbol{\theta}), \mathbf{C_s}(t)) \tag{19}$$

where the 2×2 covariance matrix

$$\mathbf{C_s}(u) = E[\mathbf{s}(u; \hat{\boldsymbol{\theta}}) - \mathbf{s}(u; \boldsymbol{\theta})][\mathbf{s}(u; \hat{\boldsymbol{\theta}}) - \mathbf{s}(u; \boldsymbol{\theta})]^T \tag{20}$$

is given by

$$\mathbf{C_s}(u) = [\nabla_{\boldsymbol{\theta}} \mathbf{s}(u; \theta)]^T \mathbf{C_{\boldsymbol{\theta}}} [\nabla_{\boldsymbol{\theta}} \mathbf{s}(u; \boldsymbol{\theta})] = \mathbf{B}(u)^T \mathbf{C_{\boldsymbol{\theta}}} \mathbf{B}(u) \tag{21}$$

where

$$\mathbf{B}(u) = \nabla_{\boldsymbol{\theta}} \mathbf{s}(u; \theta) \in \mathbb{R}^{(K+1) \times 2} \tag{22}$$

is the gradient of $\mathbf{s}(u; \boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$. In the remainder of the paper, we assume the appropriate regularity conditions are satisfied, and the estimator operates in the asymptotic regime, so that the distribution (18) can be used.

For a *fixed* u, according to classical estimation theory,[29] a confidence region for at confidence level $\alpha \in [0, 1]$ is any subset of $\mathbb{R}^2$ such that

$$\Pr\{\mathbf{s}(u) \in \hat{\mathcal{U}}_\beta(u)\} = \alpha. \tag{23}$$

While there are infinitely many choices of $\hat{\mathcal{U}}_\beta(u)$ for each specified α, Wilks and Daly[32] showed that the smallest size confidence region on the average (because is random) is

$$\hat{\mathcal{U}}_\beta(u) = \{\mathbf{x} \in \mathbb{R}^2 : (\mathbf{x} - \hat{\mathbf{s}}(u))^T \mathbf{C_s}^{-1}(u) (\mathbf{x} - \hat{\mathbf{s}}(u)) \leq \beta^2\} \tag{24}$$

for an appropriate β such that $\Pr\{\chi_2^2 \leq \beta^2\} = \alpha$, where $\mathbf{C_s}(u)$ is the covariance matrix for $\mathbf{s}(u; \hat{\boldsymbol{\theta}})$ given by (21), and χ_2^2 denotes the chi-square distribution with two degrees of freedom. For each u, $\hat{\mathcal{U}}_\beta(u)$ is an ellipse centered on the boundary point $\hat{\mathbf{s}}(u)$.

It is now desired to construct a global confidence region $\hat{\mathcal{U}}_\beta$ for the entire function $\{\mathbf{s}(u), \forall u \in I\}$. A possible design is obtained by moving $\hat{\mathcal{U}}_\beta(u)$ along the boundary

$$\hat{\mathcal{U}}_\beta = \bigcup_{u \in I} \hat{\mathcal{U}}_\beta(u). \tag{25}$$

The region $\hat{\mathcal{U}}_\beta$ forms a tube around the estimated boundary $\hat{\mathbf{s}}(u)$. Because $\hat{\mathbf{s}}(u)$ is a random process, $\hat{\mathcal{U}}_\beta$ is a random set. One important quantity for the global

confidence analysis is to find a method to calculate the global confidence level for the so-formed confidence region $\hat{\mathcal{U}}_\beta$

$$\gamma \overset{\triangle}{=} \Pr\{\mathbf{s}(u) \in \hat{\mathcal{U}}_\beta, \forall u \in I\}. \tag{26}$$

Clearly, γ is smaller than the probability in (23). Another important probability of theoretical interest is

$$\Pr\{\hat{\mathbf{s}}(u) \in \mathcal{U}_\beta, \forall u \in I\}. \tag{27}$$

where $\hat{\mathbf{s}}(u)$ is the estimate of $\mathbf{s}(u)$, and the deterministic confidence region $\mathcal{U}_\beta$ is defined as

$$\mathcal{U}_\beta = \bigcup_{u \in I} \mathcal{U}_\beta(u). \tag{28}$$

where

$$\mathcal{U}_\beta(u) = \{\mathbf{x} \in \mathbb{R}^2 : (\mathbf{x} - \mathbf{s}(u))^T \mathbf{C}_{\mathbf{s}}^{-1}(u) (\mathbf{x} - \mathbf{s}(u)) \leq \beta^2\}. \tag{29}$$

Again, the region $\mathcal{U}_\beta$ is a tube, but this time centered around the true boundary. While the probability in (26) is the *a posteriori* probability that the true shape lies in a confidence region generated around the MLE, the *a priori* probability in (27) focuses on predicting the fundamental uncertainty region for any asymptotically normal and efficient estimator. Therefore, in general they are different, and both are useful in practice.

The probabilities (26) and (27) are difficult to compute owing to the overlaps of individual ellipses $\hat{\mathcal{U}}_\beta(u)$ (resp., $\mathcal{U}_\beta(u)$) for all $u \in I$. We therefore wish to determine lower bounds on (26) and (27) that are reasonably tight. Proposition 1 provides one such tight lower bound.

Proposition 1. *Suppose* $\mathbf{C}_{\boldsymbol{\theta}}^{1/2}\mathbf{B}(u)$ *has full column rank, and*

$$\mathbf{A}(u) = \mathbf{C}_{\boldsymbol{\theta}}^{1/2}\mathbf{B}(u)\mathbf{B}(u)^T\mathbf{C}_{\boldsymbol{\theta}}^{1/2}. \tag{30}$$

Then, we have the following two inequalities:

$$\Pr\{\hat{\mathbf{s}}(u) \in \mathcal{U}_\beta, \forall u \in I\} \geq 1 - (1 + \kappa_0\beta)e^{-\frac{\beta^2}{2}} \tag{31}$$

$$\Pr\{\mathbf{s}(u) \in \hat{\mathcal{U}}_\beta, \forall u \in I\} \geq 1 - (1 + \kappa_0\beta)e^{-\frac{\beta^2}{2}} \tag{32}$$

where the constant κ_0 *is given by*

$$\kappa_0 = \int_I \left(\frac{\operatorname{trace}\{\mathbf{P}_{\mathbf{A}}^\dagger \dot{\mathbf{A}}\mathbf{A}^\dagger \mathbf{A}^\dagger \dot{\mathbf{A}}\mathbf{P}_{\mathbf{A}}^\dagger\}}{2\pi}\right)^{1/2} du \tag{33}$$

where $\mathbf{A}^\dagger$ *denotes the pseudo-inverse of* $\mathbf{A}$; $\dot{\mathbf{A}}$ *denotes the derivative of* $\mathbf{A}(u)$ *with respect to* u; *and* $\mathbf{P}_{\mathbf{A}}^\dagger$ *denotes the projection on the orthogonal complement of the range space of* $\mathbf{A}$.

Proof. See Ref. 17.   □

Such lower bounds are useful since one can generate somewhat larger confidence regions, and still guarantee the required probability that the shape lies within that region. We illustrate the construction of such confidence regions in the numerical examples.

5. The Domain Derivatives for Linear Inverse Problems

Combining (1) and (4) and separating the contribution of the domain D from the integral in (4), we obtain

$$g_m(\boldsymbol{\theta}) = c(s_m, t_m) + \int_D Z(s_m, t_m, x, y) dx\, dy$$

$$\stackrel{\triangle}{=} J(D)(s_m, t_m),\ 1 \le m \le M, \tag{34}$$

where

$$c(s_m, t_m) = \int_{\mathbb{R}^2} f_2(x, y) h(x, y, s_m, t_m) dx\, dy \tag{35}$$

$$Z(s_m, t_m, x, y) = (f_1 - f_2(x, y)) h(x, y, s_m, t_m). \tag{36}$$

Our goal is to compute the derivatives $\partial g_m(\boldsymbol{\theta})/\partial\theta_i$. The idea is to proceed using a two-step approach akin to the chain rule: (i) find the change (deformation) of $D(\boldsymbol{\theta})$ produced by an infinitesimal change in θ_i; (ii) then find the change in $\int_D Z dS$ produced by the corresponding infinitesimal deformation of D. Because D is represented by its boundary $\Gamma(\boldsymbol{\theta})$ in our parametrization, it is more convenient to consider the corresponding deformations of $\Gamma(\boldsymbol{\theta})$. The first step is easy. In view of (2)–(3), a change in θ_i to $\theta_i + t$ produces the deformed boundary

$$\Gamma_t = \Gamma + t\mathbf{b}_i = \{\mathbf{z}|\mathbf{z} = \mathbf{s}(u) + t\mathbf{b}_i(u), u \in I\}, \tag{37}$$

where $\mathbf{b}_i$ is the i-th basis function in the linear model (3) and $\mathbf{s}(u)$ is the parametric representation of boundary given by (3).

The second step thus requires to compute the change in $\int_D Z dS$ produced by deformation of Γ to Γ_t for infinitesimal t. This is known as the *domain derivative* or the *shape derivative*, and has been studied extensively in structural mechanics.[33] Using Proposition A.1 in Appendix A, we have

$$\frac{\partial g_m(\boldsymbol{\theta})}{\partial\theta_i} = \int_I (f_1 - f_2[x(u), y(u)]) h[x(u), y(u), s_m, t_m] \mathbf{b}_i^T(u) \mathbf{n}(u) \tau(u) du, \tag{38}$$

where

$$\tau(u) = \sqrt{\dot{x}(u)^2 + \dot{y}(u)^2} = \left\| \sum_{i=1}^M \theta_i \dot{\mathbf{b}}_i(u) \right\|, \tag{39}$$

$\dot{\mathbf{b}}_i(u)$ denotes the derivative of $\mathbf{b}_i$ with respect to u, and $\|\cdot\|$ denotes the Euclidean norm. In (38), the outer-normal vector $\mathbf{n}$ at $(x(u), y(u))$ is given by

$$\mathbf{n}(u) = \frac{1}{\tau(u)} \begin{bmatrix} \dot{y}(u) \\ -\dot{x}(u) \end{bmatrix} = \frac{\sum_{i=1}^M \theta_i [\dot{b}_i^y(u), -\dot{b}_i^x(u)]^T}{\left\| \sum_{i=1}^M \theta_i \mathbf{b}_i(u) \right\|} \tag{40}$$

and so

$$\mathbf{b}_i^T \mathbf{n} = \frac{\sum_{j=1}^M \theta_j \left(b_i^x \dot{b}_i^y - b_i^y \dot{b}_i^x \right)}{\left\| \sum_{i=1}^M \theta_i \dot{\mathbf{b}}_i \right\|}, \tag{41}$$

where b_i^x and b_i^y denote the x and y components of $\mathbf{b}_i$, respectively.

The derivative with respect to $\theta_{K+1}(= f_1)$ is given by

$$\frac{\partial g_m(\boldsymbol{\theta})}{\partial \theta_{K+1}} = \frac{\partial g_m(\boldsymbol{\theta})}{\partial f_1} = \int_D h(x, y, s_m, t_m) \, dx \, dy. \tag{42}$$

The following alternative expression, involving a 1D instead of an area integral, is obtained using Green's theorem:

$$\int_D h(x, y, s_m, t_m) dx \, dy = \frac{1}{2} \int_I \left(\int_0^{x(u)} h(z, y(u), s_m, t_m) \, dz \right) \dot{x}(u)\tau(u) du$$

$$- \frac{1}{2} \int_I \left(\int_0^{y(u)} h(x(u), z, s_m, t_m) dz \right) \dot{y}(u)\tau(u) du. \tag{43}$$

This expression is particularly convenient when the boundary is parameterized by (42). Even when it cannot be evaluated in closed form, its numerical evaluation requires, in general, less computation for given accuracy than that for the area integral in (42).

An important though somewhat expected observation that follows from (38), is the following. Although the measurements y_m, and in fact the existence of the transforms that define them, often depend on all of the background f_2, the Fisher information matrix only depends on the values of the background f_2 on the boundary Γ of the domain D.

5.1. *Computed tomography*

For computed tomography problem, (38) can be further simplified by exploiting the properties of the delta function. Combining (4), (5), and (38) it follows that

$$\frac{\partial g_m(\boldsymbol{\theta})}{\partial \theta_i} = \int_I \delta(F_m(u))G(u)du \tag{44}$$

where $\delta(\cdot)$ denotes the Dirac impulse and

$$F_m(u) = x(u)\cos(s_m) + y(u)\sin(s_m) - t_m \tag{45}$$

$$G(u) = (f_1 - f_2[x(u), y(u)])\mathbf{b}_i^T(u)\mathbf{n}(u)\tau(u). \tag{46}$$

Note that in (44) if the equation $F_m(u) = 0$ has no solution on I, the integral is trivially zero. However, when $F_m(u) = 0$ does have a solution, multiple such solutions may exist, all contributing to the integral. Lemma 2 provides the formula for the integral accounting for all these contributions.

Lemma 2. *Let functions $F : I \to \mathbb{R}$ be continuously differentiable and $G : I \to \mathbb{R}$ be continuous, where $I \subseteq \mathbb{R}$. Suppose furthermore $F(u)$ has L isolated roots $u_l \in I$,*

$l = 1, \ldots, L$ such that $F(u_l) = 0$ and $F'(u_l) \neq 0$. Then, if F does not have roots on boundary points of I, or if I is a closed interval and F periodic on I with period equal to its length $|I|$, we have

$$\int_I \delta(F(u))G(u)du = \sum_{l=1}^{L} \frac{G(u_l)}{|F'(u_l)|}, \tag{47}$$

where $\delta(\cdot)$ is the Dirac delta function, F' denotes the derivative of F, and $|\cdot|$ is the absolute value.

Proof. See Appendix B. □

Now we apply Lemma 2 to (44). Let $L(m)$ denote the number of solutions of $F_m(u) = 0$, and $u_l^m \in I$ denote the l-th such solution, i.e. (by (45))

$$x(u_l^m; \boldsymbol{\theta})\cos(s_m) + y(u_l^m; \boldsymbol{\theta})\sin(s_m) - t_m = 0. \tag{48}$$

Refering to Fig. 1, we identify $u_l^m \in I$ as the l-th intersection of the boundary Γ with the line indexed by (s_m, t_m). Using Lemma 2 and (48), we have

$$\frac{\partial g_m(\boldsymbol{\theta})}{\partial \theta_i} = \sum_{l=1}^{L(m)} \frac{\Delta f(u_l^m)\mathbf{b}_i(u_l^m)^T \mathbf{n}(u_l^m)\tau(u_l^m)}{|\cos(s_m)\dot{x}(u_l^m) + \sin(s_m)\dot{y}(u_l^m)|}, \quad i = 1, \ldots, K, \tag{49}$$

where

$$\Delta f(u_l^m) = f_1 - f_2(x(u_l^m), y(u_l^m)). \tag{50}$$

By (42), the derivative with respect to f_1 is given by

$$\frac{\partial g_m(\boldsymbol{\theta})}{\partial \theta_{K+1}} = \frac{\partial g_m(\boldsymbol{\theta})}{\partial f_1} = \int_D \delta(x\cos(s_m) + y\sin(s_m) - t_m)dx\, dy, \tag{51}$$

which is a line integral projection of the indicator function of the domain D. Refering again to Fig. 1, this quantity equals the sum of the lengths of the chords formed by the intersection of the line parameterized by (s_m, t_m) with the domain D. The
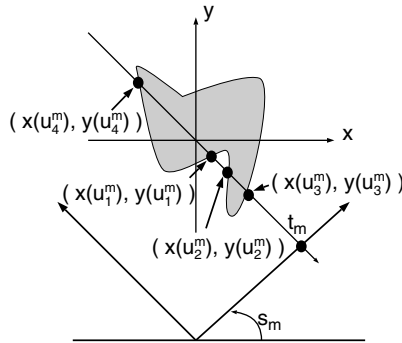

Fig. 1.   A set of $\{(x(u_l^m), y(u_l^m))\}_{l=1}^{L(m)}$ of line-boundary intersections for a tomography problem.

lengths can be computed by solving the appropriate equations (48) to find the intersection of the line parameterized by (s_m, t_m) with the parametrized boundary. The Fisher information matrix is then obtained by substituting (49) and (51) into (16).

6. Partitioned Density with Multiple Boundaries

6.1. *Partitioned density*

This section extends our previous results to more general image configurations than (1), which assumes f is constant over D. Suppose the image can be partitioned into two disjoint regions,

$$f(x,y) = \begin{cases} f_1(x,y), & (x,y) \in D \\ f_2(x,y), & (x,y) \in \mathbb{R}^2 \backslash D \end{cases}, \tag{52}$$

where the background image $f_2(x,y)$ is known, and the unknown domain D is a union of sub-domains D_j, $j = 1, \ldots, L$:

$$D = \bigcup_{j=1}^{L} D_j. \tag{53}$$

Each boundary $\Gamma_j = \partial D_j$ is parameterized as

$$\Gamma_j = \{\mathbf{s}_j(u; \boldsymbol{\theta}_j) : u \in I = [0, 2\pi)\}, \tag{54}$$

where $\boldsymbol{\theta}_j \in \mathbb{R}^{K_j}$ denotes a separate unknown parameter vector. Thus, the entire domain D is parameterized by the parameter vector $\boldsymbol{\theta} = [\boldsymbol{\theta}_1^T, \ldots, \boldsymbol{\theta}_K^T]^T$ of dimension $K = \sum_{j=1}^{L} K_j$. In general, the sub-domains D_j need not be disjoint, hence the domain D can be further partitioned into P disjoint regions:

$$D = \bigcup_{k=1}^{P} \Omega_k, \quad \Omega_j \cap \Omega_k = \varnothing, \quad j \neq k. \tag{55}$$

Furthermore, for each sub-domain D_j, there exists an index set $Q(j) \subset \{1, \cdots, P\}$ such that

$$D_j = \bigcup_{k \in Q(j)} \Omega_k. \tag{56}$$

We require $f_1(x,y)$ to be constant over each set Ω_k:

$$f_1(x,y) = \sum_{k=1}^{P} f_1^k \chi_{\Omega_k}(x,y) \tag{57}$$

where χ_{Ω_k} denotes the indicator function of the set Ω_k, and f_1^k is a constant. The partitioned density model (57) is quite general and can also serve an approximation to continuous densities.

6.2. *Computation of the CRB*

Combining the object model (52) and the measurement equation (4), we can represent $g(x, y)$ with respect to each domain D_j:

$$
\begin{aligned}
g(s, t) = &\int_{D_j} f_1(x, y) h(x, y, s, t) dx\, dy \\
&+ \int_{\mathbb{R}^2 \setminus D} f_2(x, y) h(x, y, s, t) dx\, dy \\
&+ \int_{D \setminus D_j} f_1(x, y) h(x, y, s, t) dx\, dy.
\end{aligned} \tag{58}
$$

From the partitioned density in (57), we define:

$$
f_{1,j}(x, y) = \sum_{k \in Q(j)} f_1^k \chi_{\Omega_k}(x, y) \tag{59}
$$

$$
f_{1,j^c}(x, y) = \sum_{k \in \{1, \ldots, P\} \setminus Q(j)} f_1^k \chi_{\Omega_k}(x, y). \tag{60}
$$

Using (58), (59) and (60), we then have

$$
\frac{\partial g(\boldsymbol{\theta})}{\partial \theta_i^{(j)}} = \delta J(D, \mathbf{b}_i^{(j)}) = \int_{\Gamma_j} Z_j \langle \mathbf{b}_i^{(j)}, \mathbf{n} \rangle d\Gamma, \tag{61}
$$

where $\theta_i^{(j)}$ and $\mathbf{b}_i^{(j)}$ denote the i-th element of θ_j and the corresponding basis function, respectively, and

$$
Z_j(x, y, s, t) = (f_{1,j}(x, y) - f_{1,j^c}(x, y)) h(x, y, s, t). \tag{62}
$$

Furthermore, the derivative with respect to the pixel value f_1^k is given by

$$
\frac{\partial g(\boldsymbol{\theta})}{\partial f_1^k}(s, t) = \int_{\Omega_k} h(x, y, s, t) dS. \tag{63}
$$

7. Numerical Examples

7.1. *Estimation of tumor on MRI scan of human brain*

We now turn to a more realistic imaging problem motivated by image reconstruction from sparse Fourier samples.[34–37] Consider the 128×128 MRI scan of a human brain with a tumor in Fig. 2(a). The image has 256 gray levels and was taken from a Harvard University medical database.[38] We parameterize the boundary of the tumor using Fourier descriptors (FD)[2,3]:

$$
x(u) = a_0 + \sum_{i=1}^{L} (a_i \cos(iu) + a_{L+i} \sin(iu))
$$

$$
y(u) = b_0 + \sum_{i=1}^{L} (b_i \cos(iu) + b_{L+i} \sin(iu)), \tag{64}
$$

where $L = 15$. In order to overcome the ambiguity due to the starting point of the contour, in (64) the constraint $a_{L+1} = b_1$ is imposed.[5] Hence, the resulting parameter vector $\boldsymbol{\theta}$ is:

$$\boldsymbol{\theta} = [a_0, a_1; \cdots a_{2L}; b_0, b_2, \ldots; b_{2L}]^T \in \mathbb{R}^{4L+1}. \tag{65}$$

The tumor is assumed to have constant intensity.

Suppose that 64 uniformly spaced Fourier samples are taken. This corresponds to $64/(128^2) = 0.4\%$ of the Nyquist sampling rate that would be required to avoid spatial aliasing for this 128×128 image. Suppose furthermore we have a full reference MRI scan of the healthy brain, and know *a priori* the intensity of the tumor, and the number of FD coefficients. Then, the imaging problem can be formulated as shape estimation of the tumor on the known background. Note that in this simulation the known background image has *inhomogeneous* density, and the boundary of the tumor is *not* star-shaped. Using the techniques described in Sec. 5, we calculated the CRB for the unknown FD coefficients.

The CRB matrix $\mathbf{C}_{\boldsymbol{\theta}}$ is 61×61, and even its 61-element diagonal is unwieldy, and hard to interpret. Instead we have applied the global confidence region technique of Ref. 17 to the example of Fig. 2(a) and used the computed $\mathbf{C}_{\boldsymbol{\theta}}$ to compute, in turn, the 98% asymptotic global confidence region, which is readily visualized. Consider an example where the Fourier samples are corrupted with additive complex Gaussian noise with a standard deviation σ equal to 20% of the rms value of the noise-free measurements. (We denote the corresponding signal-to-noise-ratio by SNR = 5.) Fig. 2(b) illustrates the resulting asymptotic global confidence region, in which asymptotically any unbiased estimate of the boundary will lie with 98% probability. This bound suggests that accurate estimates are possible even at low

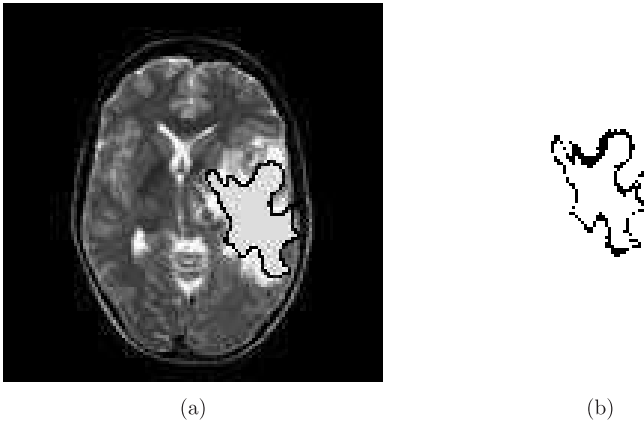(a)   (b)

Fig. 2. (a) MRI scan of a human brain with tumor, (b) 98% global confidence region (indicated by the black region) for the boundary estimate using Fourier data corrupted by complex Gaussian additive noise at SNR = 5.

sampling rates, if we have a parametric description of the tumor shape and know both the density of the tumor and the MRI scan of the healthy brain.

7.2. *Estimating boundaries of synthetic phantom*

Now we discuss the fundamental performance limit for estimating partitioned densities. Consider the 128×128 simulated phantom in Fig. 3(a). The phantom consists of four circular boundaries Γ_j, $j = 1, \ldots, 4$ (for the disks $D_j, j = 1, \ldots, 4$), which are parameterized by:

$$\mathbf{s}_j(u; \boldsymbol{\theta}_j) = \begin{bmatrix} x_j \\ y_j \end{bmatrix} + r_j \begin{bmatrix} \cos(u) \\ \sin(u) \end{bmatrix}, \quad j = 1, \ldots, 4. \tag{66}$$

where $\boldsymbol{\theta}_j = [r_j, x_j, y_j]^T$. The true values of the parameters are given by:

$$\boldsymbol{\theta}_1 = \begin{bmatrix} 0.80 \\ 0 \\ 0 \end{bmatrix}, \quad \boldsymbol{\theta}_2 = \begin{bmatrix} 0.76 \\ 0 \\ -0.02 \end{bmatrix}, \quad \boldsymbol{\theta}_3 = \begin{bmatrix} 0.21 \\ -0.20 \\ 0.35 \end{bmatrix}, \quad \boldsymbol{\theta}_4 = \begin{bmatrix} 0.21 \\ 0 \\ -0.45 \end{bmatrix} \tag{67}$$

For this problem, $D_3, D_4 \subset D_2 \subset D_1$, and the domain $D = \bigcup_{j=1}^{4} D_j$ is partitioned as follows:

$$D = \bigcup_{k=1}^{4} \Omega_k \tag{68}$$

where

$$\Omega_1 = D_1 \setminus D_2, \quad \Omega_2 = D_2 \setminus (D_3 \cup D_4), \quad \Omega_3 = D_3, \quad \Omega_4 = D_4. \tag{69}$$

The partitioned image $f_1(x, y)$ in the domain D is piecewise constant and given by

$$f_1 = \chi_{\Omega_{11}} + 0.2\chi_{\Omega_{12}} + 0.4\chi_{\Omega_{13}} + 0.4\chi_{\Omega_{14}}, \tag{70}$$
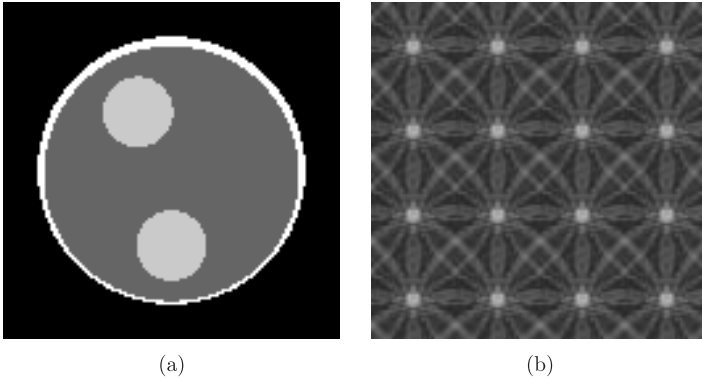

(a)   (b)

Fig. 3. (a) Synthetic phantom image; (b) direct Fourier inversion using 1024 uniformly spaced Fourier samples.

where the intensities for each domain Ω_k

$$\mathbf{f} = [1, 0.2, 0.4, 0.4] \tag{71}$$

can be estimated in case *a priori* knowledge of them are not available.

The first imaging example is again image reconstruction from sparse Fourier samples.[34–37] Suppose, that 1024 uniformly spaced Fourier samples are taken, which corresponds to $1024/(128 \times 128) = 6.25\%$ of the Nyquist rate. Direct Fourier inversion (Fig. 3(b)) exhibits strong aliasing artifacts. Suppose we know *a priori* that the image consists of four circles. Then, the imaging problem can be formulated as the estimation of the center locations and the radii of each region, and the CRB for the unknown parameter vectors can be obtained using the techniques described in the previous section. The resulting CRB is a 12×12 matrix (12 parameters) and is rather difficult to interpret. We therefore do not show it explicitly.

We have applied the asymptotic global confidence region technique to the example of Fig. 3(a) and computed the 95% asymptotic global confidence region for an example where the Fourier samples are corrupted with additive complex Gaussian noise at SNR $= 10\,$dB. First we assume that intensity (71) is known. Fig. 4(a) illustrates the resulting asymptotic global confidence region, in which asymptotically any unbiased estimate of the boundary will lie with 95% probability. This bound suggests that accurate estimates are possible even under at low sampling rates, if we have *a priori* knowledge of the number of domains and their densities. In addition, Fig. 4(a) tells us that the estimates of small region boundaries are more uncertain while the boundaries of large regions can be estimated very accurately. In Fig. 4(b), we also computed the 95% global confidence regions assuming unknown intensity. Interestingly, the uncertainty contribution from the unknown pixel values are not so significant that we can obtain nearly the same global confidence regions.

The second imaging example is a computed tomography problem. Under the same boundary parameterization (66), if each x-ray always crosses the object

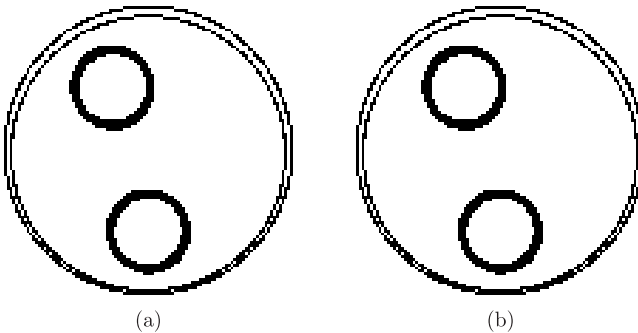(a)   (b)

Fig. 4.   Illustration of 95% global confidence regions for parametric estimation of the boundaries of the object in Fig. 3 from 1024 uniformly spaced Fourier samples and samples with additive complex Gaussian noise at SNR $= 10\,$dB assuming (a) known pixel values, and (b) unknown pixel values, respectively.

boundary twices, then the derivatives of $g_m(\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}_j = [r_j, x_j, y_j]^T$ are given by

$$\frac{\partial g_m(\boldsymbol{\theta})}{\partial r_j} = \Delta f_j \left(\sum_{l=1}^{2} \frac{1}{\left| \sin \left(s_m - u_{m,j}^l \right) \right|} \right) \tag{72}$$

$$\frac{\partial g_m(\boldsymbol{\theta})}{\partial x_j} = \Delta f_j \left(\sum_{l=1}^{2} \frac{\cos \left(u_{m,j}^l \right)}{\left| \sin \left(s_m - u_{m,j}^l \right) \right|} \right) \tag{73}$$

$$\frac{\partial g_m(\boldsymbol{\theta})}{\partial y_j} = \Delta f_j \left(\sum_{l=1}^{2} \frac{\sin \left(u_{m,j}^l \right)}{\left| \sin \left(s_m - u_{m,j}^l \right) \right|} \right), \tag{74}$$

where $u_{m,j}^l$ is computed by:

$$r_j \cos \left(u_{m,j}^l - s_m \right) = t_m - x_j \cos(s_m) - y_i \sin(s_m). \tag{75}$$

Suppose furthermore that the Radon transform is sampled uniformly both in angle s and in radial direction t with N_a and N_r samples, respectively, such that the total number of samples is $N_r \times N_a = 1024$. We again use the asymptotic global confidence region technique and compute the 95% asymptotic global confidence region for the same synthetic phantom example. Here we assume that the Radon transform samples are corrupted with additive Gaussian noise at SNR = 27 dB.

Figures 5(a) and 5(b) illustrate the resulting asymptotic global confidence regions for $N_r = N_a = 32$ assuming both known and unknown intensities, respectively. As observed in Fig. 5(a), the boundary estimates of the large regions are nearly perfect while the boundary estimates of the small regions suffer from some uncertainty. If *a priori* knowledge of intensities of (71) is not available, the uncertainty of the boundary estimates increase as shown in Fig. 5(b). Especially, the uncertainties of the boundaries of the thin ring surrounding the object increase
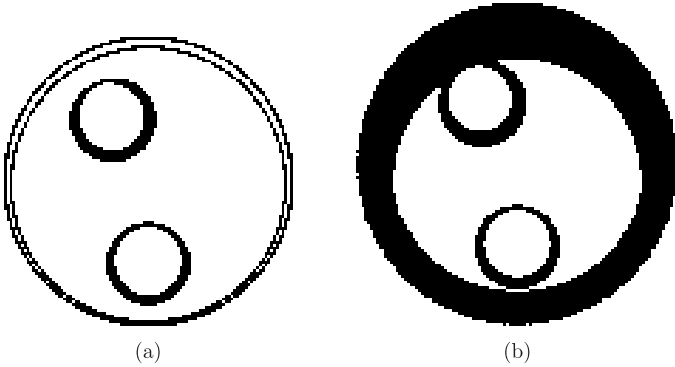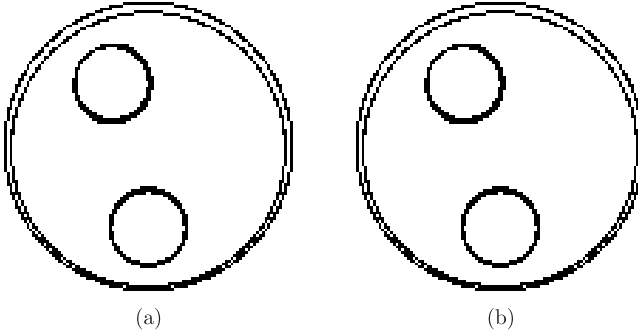

(a) (b)

Fig. 5.   Illustration of the 95% global confidence regions for parametric estimation of the boundaries of the object in Fig. 3 from 1024 uniformly spaced Radon samples with sampling parameters of $N_a = N_r = 32$ with additive real Gaussian noise at SNR = 27 dB assuming (a) known pixel values, and (b) unknown pixel values, respectively.

Fig. 6. Illustration of the 95% global confidence regions for parametric estimation of the boundaries of the object in Fig. 3 from 1024 uniformly spaced Radon samples with sampling parameters of $N_a = 16$ and $N_r = 64$ with additive real Gaussian noise at SNR = 27 dB assuming (a) known pixel values, and (b) unknown pixel values, respectively.

significantly. Note the single wide confidence region in Fig. 5(b) for the two boundaries of the thin ring surrounding the object. This is the union of the individual regions for each of the concentric circles. Compared to Fig. 5(a), we can deduce that the uncertainty of the outer ring estimate is mainly due to the uncertainty of the image value inside the narrow ring. Therefore, the single wide confidence region should be interpreted as uncertainty of the location of the ring as whole rather than suggesting that the outer boundary will be estimated to be interior to the inner boundary of the ring with some probability.

However, if we increase the radial and decrease the angular sampling rates such that $N_r = 64$ and $N_a = 16$, accurate estimation of boundaries can be achieved for both known and unknown intensities as shown in Figs. 6(a) and 6(b). Note that both boundaries of the outermost ring can be nearly perfectly estimated while the other uncertainty regions are significantly reduced compared to those in Figs. 5(a) and 5(b). Apparently, radial resolution is more important than angular resolution in this example. As observed from this discussion, the CRB and the global confidence region analysis can be used to design the optimal sampling pattern.

The last imaging example is a deconvolution problem with circular Gaussian point spread function given by (8). We again assume that 1024 uniformly spaced samples are taken from blurred image. Under the same boundary parameterization (66), the 95% global confidence regions are illustrated in Figs. 7(a) and 7(b) for both known and unknown intensities, with Gaussian width parameter $\nu = 0.2$. These figures show that *a priori* knowledge of the intensities greatly improves the estimation performance. We also computed the global confidence region for a different value of the width parameter, in Figs. 8(a) and 8(b). Comparing Figs. 7(a) and 8(a) reveals that, as might be expected, a smaller width parameter ν (less blurring) generally gives better boundary estimation performance. This trend is broken however, for sufficiently small ν (little or no blurring), as shown in Fig. 8(b)
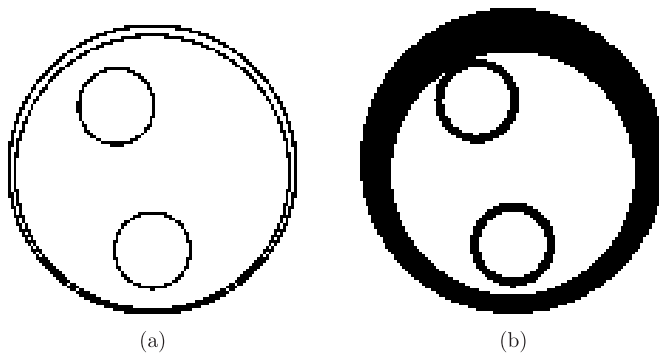
(a)   (b)

Fig. 7.   Illustration of the 95% global confidence regions for parametric estimation of the bound-
aries of the object in Fig. 3 from 1024 uniform samples of blurred image with Gaussian point
spread function at real Gaussian noise level of SNR = 15 dB and Gaussian point spread function
width parameters given by $\nu = 0.2$ assuming (a) known pixel values, and (b) unknown pixel values,
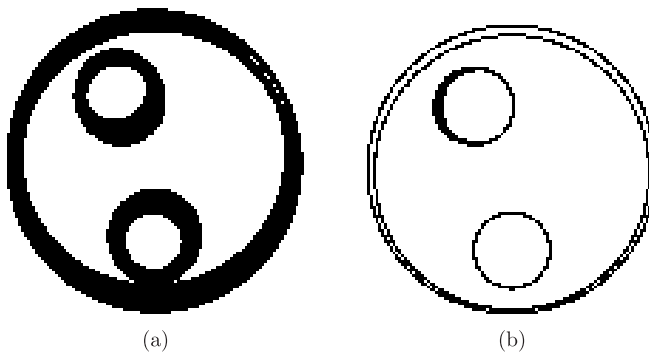respectively.

(a)   (b)

Fig. 8.   Illustration of the 95% global confidence regions for parametric estimation of the bound-
aries of the object in Fig. 3 from 1024 uniform samples of blurred image with Gaussian point spread
function at real Gaussian noise level of SNR = 15 dB assuming known pixel values. Gaussian point
spread function width parameters are given by (a) $\nu = 0.4$ and (b) $\nu = 0.00165$, respectively.

as compared to Fig. 7(a), where the uncertainty in the boundary of one of the inner
circles is larger for less blurring.

The apparent paradoxical behavior, where "better measurements" appear to
lead to "worse estimates" can be explained in the following way. If ν is much
smaller than the distance between the sampling points (or their majority) and the
object boundaries, then the measurements become essentially decoupled from the
boundaries. Small changes in the boundary will only produce vanishing changes in
the measurement. This leads to large values of the CRBs on the boundary estimates,
and to the wide associated confidence regions. These results provide a bound and
asymptotic performance prediction for *unbiased* estimators. Better estimates may
be provided by biased estimators (e.g. by bracketing the boundary position between

the sampling points), but their performance is not bounded or predicted by the CRB and associated confidence regions.

8. Conclusions

This chapter has reviewed parametric shape reconstruction algorithms and introduced a general method to compute Cramér–Rao bounds and the asymptotic global confidence regions for parametric shape estimation in linear inverse problems, such as computed tomography, Fourier imaging, deconvolution, and etc.

Although the imaging map is linear, the dependence of the measurements on the shape of the object, as described by its boundaries, is highly nonlinear. If object boundaries are parameterized using a finite number of unknown parameters, superior reconstruction can be obtained even under highly limited samples. Furthermore, such superior performance can be theoretically verified by computing Cramér–Rao bounds and the asymptotic global confidence region as described in this chapter.

Our approach is quite general and can be easily adapted to the particular form of any linear inverse problem with possibly nonlinear observations.

Appendix A. The Domain Derivative

Combining the object model (1) and the noise-free measurement equation (4) yields

$$g(s,t) = f_1 \int_D h(x, y, s, t) dx\, dy$$
$$+ \int_{\mathbb{R}^2 \setminus D} f_2(x, y) h(x, y, s, t) dx\, dy, \quad (s, t) \in \Omega. \tag{A.1}$$

Equation (A.1) then defines a mapping $J: \{D\} \rightarrow \{g\}$ from the set of domains $\{D\}$, or equivalently, boundaries $\{\Gamma\}$, to the space of functions $\{g\}$. This mapping admits the general form:

$$g = J(D) = f_1 \int_D Z_1\, dS + \int_{\mathbb{R}^2 \setminus D} Z_2\, dS = c + \int_D Z\, dS, \tag{A.2}$$

where $dS = dx\, dy$, Z_1, Z_2, and $Z = f_1 Z_1 - Z_2$ are known functions on $\mathbb{R}^2$, D is the unknown object support, and $c = \int_{\mathbb{R}^2} Z_2\, dS$ is a function independent of D.

Given our parametrization of the boundary, we can rewrite (A.2) to display explicitly the dependence on $\boldsymbol{\theta}$,

$$g(\boldsymbol{\theta}) = J[D(\boldsymbol{\theta})] = c + \int_{D(\boldsymbol{\theta})} Z\, dS. \tag{A.3}$$

One essential step required to compute $\nabla_{\boldsymbol{\theta}} g(\boldsymbol{\theta})$ is produced by the functional changes $\int_D Z dS$ with respect to the deformation of Γ to Γ_t for infinitesimal t (see (37)). This is known as the *domain derivative* or the shape derivative, and has been studied extensively in structural mechanics.[33]

Proposition A.1. *Let $D = D(\boldsymbol{\theta})$ be a domain with boundary Γ of class C.[1] Suppose the boundary Γ is deformed as in (37). Then, the domain function (A.3) is shape differentiable with respect to boundary deformation $\mathbf{b}_i$, with domain derivative with respect to parameter θ_i:*

$$\frac{\partial g}{\partial \theta_i} = \delta J(D; \mathbf{b}_i) = \int_I Z(u)\left(\mathbf{b}_i^T \mathbf{n}\right)(u)\tau(u)du, \quad i = 1, \ldots, K, \tag{A.4}$$

where $Z(u) \triangleq Z[\mathbf{s}(u)], \tau(u) = \sqrt{\dot{x}(u)^2 + \dot{y}(u)^2}$ where $\dot{x}(u)$ and $\dot{y}(u)$ are the derivatives of the components of $\mathbf{s}(u) = [x(u), y(u)]^T$, and $\mathbf{n}$ denotes the outer-normal vector of Γ.

Proof. See Ref. 41. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad\square$

Appendix B. Proof of Lemma 2

For any $\epsilon > 0$, define

$$U_\epsilon = \bigcup_i (u_i - \epsilon, u_i + \epsilon).$$

Then, because the integrand is identically zero on $I \backslash U_\epsilon$,

$$\int_I \delta\left(F(u)\right)G(u)du = \sum_{i=1}^L \int_{u_i-\epsilon}^{u_i+\epsilon} \delta\left(F(u)\right)G(u)du. \tag{B.1}$$

Next, because the roots are isolated, F continuously differentiable and $F'(u_i) \neq 0$, $\epsilon > 0$ can be chosen such that u_i is the only root of $F(u)$ and $F'(u) \neq 0$ for $u \in (u_i-\epsilon, u_i+\epsilon)$, for all i. It follows that $F(u)$ is strictly monotone for $u \in (u_i-\epsilon, u_i+\epsilon)$, and has an inverse $H_i(\alpha) = u$, with $\alpha = F(u)$, and $H_i(0) = u_i$. We can therefore use the change of variables $\alpha = F(u)$,

$$\int_{u_i-\epsilon}^{u_i+\epsilon} \delta\left(F(u)\right)G(u)du = \int_{F(u_i-\epsilon)}^{F(u_i+\epsilon)} \delta(\alpha)G\left(H_i(\alpha)\right)\frac{d\alpha}{F'(H_i(\alpha))} \tag{B.2}$$

$$= \int_{a_i}^{b_i} \delta(\alpha)G(H_i(\alpha))\frac{d\alpha}{F'(H_i(\alpha))}, \tag{B.3}$$

where, because of the continuity of F and its monotonicity on the interval $u \in (u_i - \epsilon, u_i + \epsilon)$, $a_i = F(u_i - \epsilon) < 0$ and $b_i = F(u_i + \epsilon) > 0$ if $F'(u_i) > 0$, with

opposite signs otherwise. Using the sifting property of $\delta(\cdot)$ we obtain

$$\int_{u_i-\epsilon}^{u_i+\epsilon} \delta\left(F(u)\right)G(u)du = \frac{G(H_i(0))}{|F'\left(H_i(0)\right)|} = \frac{G(u_i)}{|F'(u_i)|}. \tag{B.4}$$

Subtituting (B.4) into (B.1) concludes the proof.

References

1. C. K. Chui, *Multivariate Splines* (SIAM, Philadelphia, 1988).
2. C. T. Zahn and R. Z. Roskies, Fourier descriptors for plane closed curves, *IEEE Trans. Computers* (3) (1972) 269–281.
3. E. Persoon and K. S. Fu, Shape discrimination using Fourier descriptors, *IEEE Trans. Systems Man and Cybernetics* **7**(3) (1977) 170–179.
4. A. Kirsch, The domain derivative and two applications in inverse scattering, *Inverse Problems* **9**(1) (1993) 81–96.
5. J. C. Ye, Y. Bresler and P. Moulin, Cramér–Rao bounds for 2-D target shape estimation in nonlinear inverse scattering problems with application to passive radar, *IEEE Trans. Antennas and Propagat.* **49**(5) (2001) 771–783.
6. H. L. VanTrees, *Detection, Estimation and Modulation Theory, Part I: Detection, Estimation, and Linear Modulation Theory* (John Wiley & Sons, Inc., New York, 1968).
7. S. F. Yau and Y. Bresler, Worst case Cramér–Rao bounds for parametric estimation of superimposed signals with applications, *IEEE Trans. Signal Processing* **40**(12) (1992) 2973–2986.
8. A. J. Devaney and G. A. Tsihirintzis, Maximum likelihood estimation of object location in diffraction tomography, *IEEE Trans. Signal Processing* **39**(3) (1991) 672–682.
9. G. A. Tsihirintzis and A. J. Devaney, Maximum likelihood estimation of object location in diffraction tomography, Part II: strongly scattering objects, *IEEE Trans. Signal Processing* **39**(6) (1991) 1466–1470.
10. A. Schatzberg, A. J. Devaney and A. J. Wittens, Estimating target location from scattered field data, *Signal Process* **40**(2–3) (1994) 227–237.
11. D. J. Rossi and A. S. Willsky, Reconstruction from projections based on detection and estimation of objects — Part I: Performance analysis, *IEEE Trans. Acoustic Speech and Signal Processing* **32**(4) (1984) 886–897.
12. P. Stoica and A. Nehorai, MUSIC, maximum likelihood, and Cramér–Rao bound, *IEEE Trans. Acoustic Speech and Signal Processing* **37**(5) (1989) 720–741.
13. P. Stoica and A. Nehorai, Performance study of conditional and unconditional direction-of-arrival estimation, *IEEE Trans. Acoustic Speech and Signal Processing* **38**(10) (1990) 1783–1795.
14. A. L. Swindlehurst and P. Stoica, Maximum likelihood methods in radar array signal processing, *Proc. IEEE* **86**(2) (1998) 421–441.
15. P. S. Naidu and A. Buvaneswari, A study of Cramér–Rao bounds on object shape parameters from scattered field, *IEEE Trans. Signal Processing* **47**(5) (1999) 1478–1481.
16. D. J. Rossi, A. S. Willsky and D. M. Spielman, Object shape estimation from tomographic measurements — a performance analysis, *Signal Process* **18**(1) (1989) 63–87.
17. J. C. Ye, Y. Bresler and P. Moulin, Asymptotic global confidence regions in parametric shape estimation problems, *IEEE Trans. Information Theory, Special Issue on Information Theoretic Imaging* **46**(5) (2000) 1881–1895.

18. K. M. Hanson, G. S. Gunningham and R. J. McKee, Uncertainty assessment for reconstructions based on deformable geometry, *Int. J. Imaging Systems & Technology* **8**(6) (1997) 506–512.

19. G. C. H. Chuang and C. C. Jay Kuo, Wavelet descriptor of planar curves: theory and applications, *IEEE Trans. Image Processing* **5**(1) (1996) 56–70.

20. P. Wunsch and A. F. Laine, Wavelet descriptors for multiresolution recognition of handprinted characters, *Pattern Recognition* **28**(8) (1995) 1237–1249.

21. F. Natterer, *The Mathematics of Computerized Tomography* (Teubner, Stuggart and Wily, New York, 1986).

22. S. R. Deans, *The Radon Transform and Some of its Applications* (Wiley, New York, 1983).

23. R. E. Blahut, *Theory of Remote Image Formation* (Cambridge University Press, 2000).

24. R. N. Bracewell, *Two-Dimensional Imaging* (Prentice-Hall, Englewood Cliffs, NJ, 1995).

25. D. C. Munson-Jr., J. D. O'Brien and W. K. Jenkins, A tomographic formulation of spotlight-mode synthetic aperture radar, *Proc. IEEE* **71**(8) (1983) 917–925.

26. M. Soumekh, *Fourier Array Imaging* (Prentice-Hall, Englewood Cliffs, NJ, 1994).

27. Z. P. Liang and P. C. Lanterbur, *Principle of Magnetic Resonance Imaging: A Signal Processing Perspective* (IEEE Press, New York, 2000).

28. A. O. Hero, R. Piramuthu, J. A. Fessler and S. R. Titus, Minimax estimation computed tomography using high resolution anatomical side information and B-spline models, *IEEE Trans. Information Theory* **45**(3) (1999) 920–938.

29. M. Kendall and A. Stuart, *The Advanced Theory of Statistics*, 4th ed. (MacMillan Publishing Co., New York, 1977).

30. H. V. Poor, *An Introduction of Signal Detection and Estimation.* 2nd ed. (Springer-Verlag, New York, 1994).

31. I. A. Ibragimov and R. Z. Has'minskii, *Statistical Estimation: Asymptotic Theory* (Springer-Verlag, New York, 1979).

32. S. S. Wilks and J. F. Daly, An optimum property of confidence regions associated with the likelihood function, *Ann. Math. Statist.* **20**(10) (1939) 225–235.

33. J. Sokolowski and J. Zolesio, *Introduction to Shape Optimization: Shape Sensitivity Analysis* (Springer-Verlag, New York, 1991).

34. R. Venkataramani and Y. Bresler, Perfect reconstruction formulas and bounds on aliasing error in sub-Nyquist nonuniform sampling of multiband signals, *IEEE Trans. Information Theory* **46**(6) (2000) 2173–2183.

35. Y. Bresler and P. Feng, Spectrum-blind minimum-rate sampling and reconstruction of 2-D multiband signals, in *Proc. 3rd IEEE Int. Conf. Image Processing, ICIP'96*, Lausanne, Switzerland, vol. I (September 1996), pp. 701–704.

36. R. Venkataramani and Y. Bresler, Further results on spectrum blind sampling of 2D signals, *Proc. IEEE Int. Conf. Image Proc., ICIP*, Chicago, Vol. 2 (October 1998), pp. 752–756.

37. Y. Bresler, M. Gastpar and R. Venkataramani, Image compression on-the-fly by universal sampling in Fourier imaging systems, in *Proc. 1999 IEEE Information Theory Workshop Detection, Estimation, Classification, and Imaging*, Santa-Fe, NM (February 1999), pp. 48–48.

38. *Metastatic bronchogenic carcinoma*, http://www.med.harvard.edu/AANLIB/home.htm.

39. Yibin Zheng and P. Doerschuk, 3-D image reconstruction from averaged Fourier transform magnitude by parameter estimation, *IEEE Trans. Image Processing* (1988) no. 11. 1561–1570.

40. N. Schmid, Y. Bresler and P. Moulin, Complexity regularized shape estimation from noisy Fourier data, *Proc. IEEE Int. Conf. Image Proc.* ICIP'02, Rochester, NY, pp. 453–456.

41. J. C. Ye, Y. Bresler and P. Moulin, Cramér-Rao bound of parametric shape estimation in inverse problems, *IEEE Trans. Image Processing* (2003), no. 1. pp. 71–84.

CHAPTER 6

WAVELET TECHNIQUES IN REGION-BASED DIGITAL DATA COMPRESSION AND THEIR APPLICATION IN DIGITAL MAMMOGRAPHY

MÓNICA PENEDO

Laboratorio de Imagen Médica. Unidad de Medicina Experimental
Hospital General Universitario Gregorio Marañón
28007 Madrid, Spain
monica@mce.hggm.es

WILLIAM A. PEARLMAN

Electrical, Computer Systems Engineering Department
Rensselaer Polytechnic Institute
Troy-NY, USA

Mammography is the most effective method for the early detection of breast cancer. However, the high-resolution requirements associated to preserve small-size or low-contrast lesions have limited the widespread use of digital mammography and have hindered the use of computers as a second opinion in the detection of this disease. Data compression provides an important and viable solution to this problem. In this chapter we describe region-based methods for wavelet compression that represent the state-of-the-art in image compression performance. We begin by reviewing means for coding regions of interest with higher fidelity than the background through scaling of the corresponding sets of wavelet coefficients, including the JPEG2000 Maxshift method. Then, we describe briefly the SPIHT (Set Partitioning in Hierarchical Trees) image compression algorithm and its particular ROI (region of interest) adaptation, called OBSPIHT, to digital mammography. Next, we present a quantitative comparison of OBSPIHT, SPIHT, and JPEG2000 in the compression of the breast region. As we will see, the results indicate that region-based compression methods are clearly superior to regular methods for digital mammography. An extensive clinical study that ensued to compare OBSPIHT with JPEG2000 will show that the radiologically important features are well preserved even at substantial compression ratios.

Keywords: Digital mammography; wavelet techniques.

1. Mammography and Breast Cancer Detection

Breast cancer represents a leading cause of cancer death among women in developing countries. Despite that this disease represents a sanitary problem of great importance, the risk factors associated to breast cancer are not much modifiable or susceptible of primary prevention. Today, the efforts for the treatment of breast cancer are aimed toward establishing a diagnostic decision as early as possible. Breast cancer survivor success is strongly related to the stage of cancer at diagnosis. The earlier the breast cancer could be detected, the greater the possibility of treatment

success and the fewer women that would die or suffer unnecessary mutilation. For this reason, the secondary prevention, early detection, plays a fundamental role in the control of this disease.[1]

It has been shown that the most effective method for the detection of breast tumors before they manifest clinical symptoms is mammography. However, the lack of methodological procedures for predicting those women susceptible of suffering the disease has lead to run breast cancer screening programs that evaluate asymptomatic populations using mammography. Therefore, mammography has become the medical image modality used as a screening and diagnostic tool for the detection of breast cancer.

2. Digital Mammography

The main advantage of digital medical imaging systems is the possibility of splitting the acquisition, processing and display processes into separate stages, and offering the possibility of independently improving each stage. Digital image format allows a more efficient health care environment that integrates the hospital information system (HIS), which manages demographic and clinical history of the patients with the picture archiving and communication system (PACS). Moreover, digital medical imaging permits the utilization of advanced applications such as computer-aided detection (CAD) schemes to assist radiologists in the detection of radiological features that could point to different pathologies. Medical images in digital format constitute an obligatory requisite toward the implementation of teleradiology, offering the possibility of a second opinion consultation from an expert. These advantageous features favor the replacement of conventional screen film imaging by digital technology in radiology departments.

As other image modalities, mammography benefits from the digital format. However, small changes in the performance of an image modality used for screening tasks will affect a considerable number of individuals if the screening population is large. This fact has motivated that the commercialization of digital mammography devices require more extensive testing than any other x-ray technique for only diagnostic purposes.[2] Currently, there are four full-field digital mammography systems approved for clinical use by the U.S. Food and Drug Administration (FDA).[3] Each of these systems is based in different detector technologies: the FCR 5000MA Computed Radiography for Mammography (Fuji Medical Systems, Stamford, CT, USA), the SenoScan Full-Field Digital Mammography System (Fischer Imaging Corporation, Denver, CO, USA) that uses CCD (charged couple device) technology, the Senographe 2000D (GE Medical Systems, Milwaukee, WI, USA) that uses flat-panel detectors based on amorphous silicon, and the Lorad Selenia Full Field Digital Mammography System (Hologic Inc., Bedford, MA, USA) based in amorphous selenium technology. There is also a direct x-ray acquisition system based on silicon dioxide detector under clinical trials, the Sectra MicroDose Mammography system (Sectra, Linköping, Sweden).

However, the main advantage of conventional mammography is the image quality obtained at low cost. For this reason, the widespread use of digital mammography devices has been conditioned to evaluate whether the digital modality provides a better diagnostic performance over the conventional technology. Research works carrying out comparative studies have shown that digital mammography systems significantly improve detection of subtle lesions such as microcalcifications, visualization of low-contrast details or dense breast parenchymal tissue, and definition of skin.[4–9] Another study comparing detection of breast cancer using hard copy images or soft copy images displayed on a high-resolution monitor found no significant differences in diagnostic interpretation or reporting speed between both modalities.[10] The image processing methods included in digital mammography systems permit filtering, contrast enhancement, or other types of digital image manipulation that avoids repeating any mammogram with inadequate quality,[11] in contrast to conventional techniques where images are rejected in case of poor contrast or under/over radiation exposure. Those encouraging results promise the implementation of the digital technology in mammography in the near future.

3. Wavelet Compression in Digital Mammography

Transition from conventional to digital mammography also depends on two parameters of importance: the spatial and grey-level resolutions. Both pixel size and pixel depth are factors that critically affect the visibility of small-low contrast objects or signals, which often are relevant information for diagnosis. They must provide a diagnostic accuracy in digital images equivalent to that of conventional films.

For instance, it has been shown that isolated clusters of microcalcifications are one of the most frequent radiological features of asymptomatic breast cancer. A careful search for the clustered microcalcifications that may herald an early-stage cancer should be done on all mammograms.[1] Microcalcifications frequently appear as small sized-low contrast radiopacities.[12] Because of this, a typical mammogram must be digitized at a resolution of about 4000×5000 pixels with $50\,\mu\text{m}$ spot size and 12 bits, resulting in approximately $40\,\text{Mb}$ of digital data.[13] These high-resolution requirements also retard the implementation of digital mammography systems due to the increase in processing, transmission time and storage capacity cost associated with these high-resolution images. It is clear that advances in technologies for transmission or storage are not sufficient to solve this problem. Efficient data-compression schemes to reduce digital data have been investigated as an important and viable solution to this problem.

In medical image compression, and generally in data compression, the wavelet transform (see Appendix) emerged as a tool for new compression methods that outperformed the existing techniques mainly based on the discrete cosine transform (DCT).[14–16] Studies applying wavelet-based methods on medical images reported better qualitative performance than the DCT-based JPEG (Joint Photographic Experts Group) standard for image compression.[17–19] With time, the

wavelet transform became a technique frequently used when compressing medical images.[20,21] In digital mammography, research studies are analogously presented evaluating DCT[22–26] and wavelet compression techniques,[27,28] both reversible or lossless (exact reconstruction of the original image after compression) and irreversible or lossy (some information is lost during the compression process) methods.

More recently, embedded wavelet methods appeared with the possibility of progressive encoding of digital images.[29,30] These embedded methods are appropriate for applications such as telemammography, transmitting a first approximation of the mammogram and progressively increasing the pixel accuracy until a radiologist could diagnose with the reconstructed image quality achieved. Methods such as the well-known SPIHT (Set Partitioning in Hierarchical Trees) and the new emerging standard for still image compression, the JPEG2000, are embedded techniques based on the wavelet transform that have been applied to digital mammography.[31–34]

However, lossless techniques provide only modest reduction in file size. In images of high-resolution format used in screening programs as digital mammography, significant reduction of transmission and storage costs is mainly obtained using irreversible compression methods, although with detriment to the quality of the reconstructed image.

4. Evaluation of Compression Methods in Digital Mammography

Although some controversy still exists about the use of lossy compression in medical imagery, the U.S. Food and Drug Administration (FDA) allows medical image management devices to use irreversible compression if it is guaranteed that the information loss does not affect diagnostic performance. Research studies need to be performed, since the particular compression ratio that would be acceptable depends on the radiological features of the medical image modality and the compression technique used.[20] Some evaluation studies have determined that loss of information does not compromise diagnostic interpretation when compression rates are limited to certain ranges.[24,26,31,34–37]

4.1. *FDA and ACR-NEMA specifications for medical image compression*

The FDA demands that medical management device manufacturers identify any standard data compression scheme used for image communications and storage. When those devices use non-standard compression methods, such algorithms should be described in detail, while providing any available technical report. In either case, standard and non-standard methods, manufacturers should specifically state the compression ratios to be employed.[38]

Following FDA specifications, the operating instructions of medical image management devices utilizing irreversible methods should include description of compression effects and examples of the resulting image quality after the loss of

information. Moreover, all images compressed with an irreversible method should be labeled with the compression ratio used. Manufacturers should conduct laboratory tests for communications and storage devices using non-standard irreversible image compression techniques. For such devices, the normalized mean square error should be reported for each compression level employed by the device.

Teleradiology is defined by the American College of Radiology (ACR) as the electronic transmission of radiological images from one location to another for the purposes of interpretation and/or consultation.[39] The ACR Technical Standard for Teleradiology allows the use of reversible and irreversible data compression techniques to increase transmission speed and reduce storage requirements. Always under the supervision of a qualified physician, it is possible to utilize any compression ratio but the loss of information should not significantly decrease the image quality for clinical diagnostic purposes.

4.2. *Quantitative evaluation*

The evaluation of a compression method can be specified by different measures such as the distortion introduced, the memory requirements for its execution, the relative complexity and speed of the algorithm, the compression ratio, etc.

For example, the compression ratio (which is defined as the quotient between the number of bits necessary to represent the original image and the number of bits needed for the compressed image) determines the level of compression achieved by the method. Another common measure for reporting compression performance is the average number of bits required to represent a single sample of the compressed image. This is generally referred to as compression rate, which is measured in bits per pixel (bpp). Both compression ratio and the rate of the compression method are related by the expression:

$$Compression\ Ratio = \frac{rate\ of\ the\ original\ image}{rate\ of\ the\ compressed\ image} \tag{1}$$

Another important issue is to establish the quality of the reconstructed image, i.e. the distortion introduced in the compression process. As the compression ratio increases, the quality of the resulting image is degraded. A parameter for measuring the degree of distortion introduced is the peak signal-to-noise ratio (PSNR), which is often measured in a logarithmic scale

$$PSNR([dB]) = 10\log_{10}\left(\frac{A^2}{MSE}\right) \tag{2}$$

where A is the peak amplitude of the original image, and MSE is the mean squared-error between the original and the reconstructed image:

$$MSE = D_{MSE}(X - \hat{X}) = \frac{1}{N}\sum_{n=1}^{N}(x_n - \hat{x}_n)^2 \tag{3}$$

with x_n and $\hat{x}_n$ the original and the reconstructed image samples, respectively, and N the number of samples in each image.

Both PSNR and MSE are the most common parameters used for measuring the image quality of reconstructed images after compression. Although these parameters could be of limited value for medical images, the FDA demands to measure and report the normalized mean squared error (defined as the MSE normalized by the variance of the original image) for each compression factor utilized by those medical management devices using lossy compression methods.[38]

4.3. *Qualitative evaluation*

In digital mammography (and in general in digital medical imaging), the evaluation of a lossy compression method must consider whether the information loss is diagnostically significant for the specific clinical issue to be addressed. It is therefore necessary to investigate and determine the compression range where the compression method can be used without significant degradation of the medical image quality for human or machine interpretation.

4.3.1. *Important radiological signs in digital mammography*

The detection of primary radiological signs of breast cancer in compressed digital mammograms must be evaluated for determining whether it is possible to compress a digital mammogram before using it for diagnosis. In mammography, clusters of microcalcifications and masses are the most relevant radiological signs indicating the presence of breast cancer:

(i) Microcalcifications are small deposits of calcium with size between 0.1 and 2.0 mm. They can appear isolated or gathered in clusters (Fig. 1). While in most cases individual microcalcifications are not clinically significant, clustered microcalcifications appear in 30%–50% of breast cancers.[1] Malignancy of microcalcifications is mainly determined based on their radiological characteristics, such as type of clustering, number of microcalcifications within the cluster, their distribution, size or shape.

Identification and analysis of microcalcifications are difficult tasks for radiologists. Most of the difficulty is due to their irregular shape, their variation in size, and the inhomogeneous background of the parenchymal tissue in which
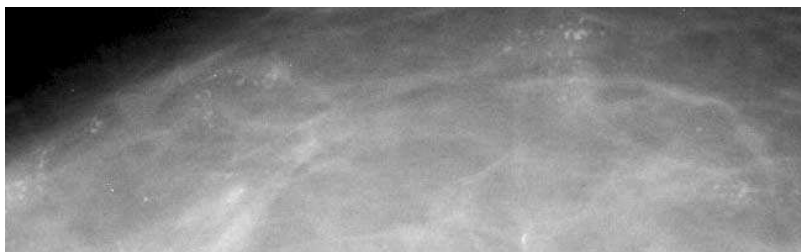

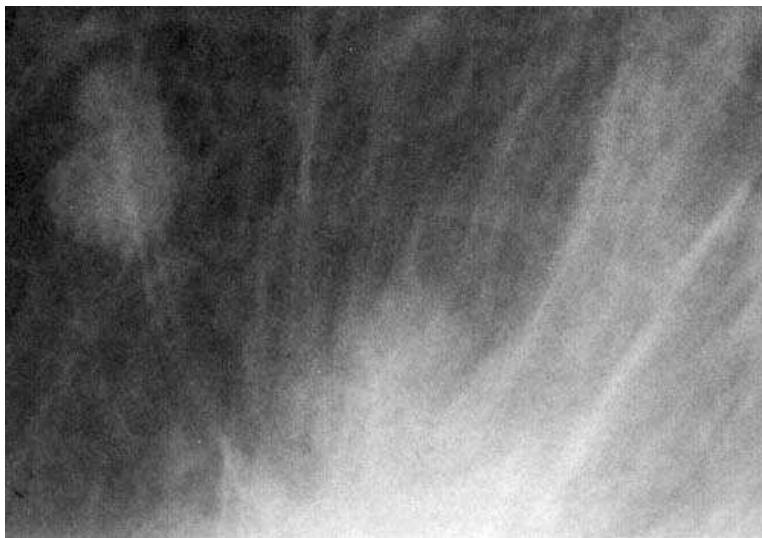Fig. 1.   Mammogram region containing several clusters of microcalcifications.

Fig. 2.   Mammogram region containing a malignant mass with speculated shape and a benign mass of rounded shape.

they are located, resulting in a poor contrast of the microcalcifications with the surrounding structures. Because of their clinical relevance and potential subtlety, detection of microcalcifications has been a primary focus of attention for the radiology community.[40]

(ii) Masses are three-dimensional structures that may represent a localizing sign of breast cancer. Malignant masses usually have spiculated or ill-defined margins and irregular shapes. Spiculated and ill-defined margins are due to the infiltrative process of the tumour in the tissue. On the other, benign masses are usually characterized for having sharp, well-defined margins and they generally present round regular shapes. In many cases, malignant masses are embedded in the glandular tissue, so that their contrast with the surrounding is very low. The detection of masses in mammograms is difficult because of their low contrast, and their variation in shape and size (Fig. 2).

The high image quality requirements for the detection of these important radiological signs in breast cancer have made digital mammography one of the medical image modalities where compression applications need to be evaluated with more care.

4.3.2. *Evaluation experiments*

First developed in statistical decision theory, Receiver Operating Characteristic (ROC) analysis is widely used to evaluate and compare the performance of diagnostic techniques in medical imaging.[41] In an ROC experiment, an observer or radiologist views a group of images and rates each image (typically in a 5-point or

continuous rating scale) according to his/her confidence about the presence (positive) or absence (negative) of a specified disease or radiological sign. Observer responses yield an ROC curve plotting sensitivity (the fraction of actually-positive cases correctly classified as positive) versus 1-specificity (the fraction of actually-negative cases incorrectly classified as positive). The area under the ROC curve is used as an objective measure of performance of the combination image modality and observer in detecting the disease or radiological sign. However, ROC experiments are limited to binary decisions, without considering signal location.

There are currently three generalizations of the ROC methodology to overcome this limitation, each requiring a different task of the observer for collecting detection and localization data. The Localization ROC (LROC) approach only permits zero or one signal per image and forces the observer to make a location response of a suspicious region even if the observer considers the image as negative.[42] In the region of interest (ROI) approach to location-specific ROC analysis,[43–45] the observer is forced to localize the signal within a quadrant or region of the image. This experiment imposes a reading paradigm different from the clinical practice. Alternatively, in the Free-response ROC (FROC) experiment,[46,47] the observer determines the number of responses or suspicious regions per image and localizes them, assigning a rate for each response. The FROC paradigm allows any number of signals per image, which appears to be a closer representation of the search strategy followed by radiologists in clinical practice.

The performance of radiologists reading digital mammograms reconstructed after compression can be investigated by means of these observer performance studies. Evaluation of a compression method is performed comparing the figures of merit (for example, the areas under the curve) obtained from the analysis of the collected data for each radiologist when using original digital mammograms and digital mammograms reconstructed after compression.

5. Studies in Wavelet Compression of Digital Mammography

In the literature, there are some research works evaluating lossy compression in digital mammography with computer-aided detection (CAD) systems or observer performance studies.

Good *et al.* applied the old standard JPEG to sixty digital mammograms, evaluating the performance of eight observers in the detection of masses and clusters of microcalcifications by ROC (Receiver Operating Characteristic) analysis.[24] In a similar study, Zheng *et al.*[26] assessed the performance of a CAD system in the detection of primary signs of breast cancer in 952 digital mammograms reconstructed after JPEG compression. Both studies obtained no statistical difference in results for detecting masses, but a significant difference was found for the detection of clusters of microcalcifications when using compressed images.

As wavelet transform emerged as a more effective technique for image compression than JPEG, other evaluation studies appeared in the literature applying wavelet-based methods to digital mammography.[31,37] Kocsis *et al.*[37] assessed

the detection of clusters of microcalcifications in sixty-eight mammograms compressed with a public domain wavelet-based method. Four observers were included in their ROC study obtaining that a visually lossless threshold is found at 40:1. Perlmutter *et al.*[31] evaluated fifty-seven digital mammograms compressed with the SPIHT (Set Partitioning In Hierarchical Trees) algorithm. In their work, a designed protocol simulating clinical practice was used to assess whether lossy compression yields differences to fundamental decisions made in screening. Although not based in accuracy of lesion detection, sensitivity and specificity were calculated for each observer and compression level, concluding that no statistical significant differences existed between analog or digital originals and compressed mammograms at 80:1 compression ratio.

In the last years, the JPEG2000 compression standard, based on wavelet transform, has emerged providing new features with promising applications in medical image compression. The JPEG2000 standard has been evaluated in digital mammography. Sung *et al.*[33] have assessed JPEG2000 at several compression ratios in twenty low-resolution digitized mammograms. They included three radiologists in the ROC study and determined no differences in lesion detectability with ratios up to 20:1. Suryanarayanan *et al.*[34] have investigated the effect of JPEG2000 in ten contrast-detail phantoms containing circular gold disks of different diameter (0.1–3.2 mm) and thickness (0.05–1.6 μm). The phantoms were acquired in a clinical full-field digital mammography system and compressed at different ratios (10, 20, and 30:1). They included seven observers in the study, finding no significant difference in perception up to 20:1 except for the disks of 1mm of diameter.

6. Region-based Wavelet Compression of Medical Images

Encoding of arbitrarily shaped regions inside an image at different quality levels can help to develop medical image compression methods that focus on those regions that are important for diagnostic purposes. Such methods provide the possibility of adequately compressing the regions with diagnostic relevance with better quality than the rest of the image.

Generally, in region-based compression methods the whole image is transformed (for example using the wavelet transform) and those coefficients associated to the region of interest are coded at higher precision (up to lossless) than the rest of the image. After segmenting the image into important regions (automatically or manually), coding of only ROI-related coefficients can be accomplished either after a certain compression bit rate of the entire image has been reached,[48] or before the information associated to the background,[49] as in the region-based coding methods of the JPEG2000 standard: the Maxshift method and a general-scaling based method.[50,51]

The growing interest in manipulating visual objects in digital images or video has led to new region-based or object-based wavelet coding techniques, which describe the images in terms of arbitrary contours and image samples inside the contour, coding them separately.[52,53] These new techniques apply shape-adaptive

wavelet transformation only to the image samples inside the object of interest[54] instead of the full image, as does the region-based methods in JPEG2000. Then, the object shape information is coded (lossless or lossy) before the pixels of the object. According to the technique used for coding the shape information, the object-based coding methods are classified into block-based or contour-based techniques. In the first, a binary image representing the object shape is encoded as in a conventional image coding method. In the latter, the shape information coding is performed along the edge or contour of the object. Among the contour-based techniques, a chain code[55] is one of the most frequently used methods for lossless coding of the object shape information. In this method, instead of encoding the absolute position of each contour point, the relative position between two consecutive points is encoded. Therefore, an initial point in the contour object and the necessary links following this point to describe the contour of the object are enough for reconstructing the object shape.

6.1. *JPEG2000 Maxshift ROI coding*

In the Maxshift ROI coding adopted in JPEG2000 Part-1,[56,57] an entire image is transformed and only the coefficients associated with the region of interest are scaled up through a particular number of bit-shifts (called scaling value s) given by the largest number of non-empty magnitude bit-planes of the coefficients (Fig. 3a). The bit-planes of coefficients are encoded plane by plane to let the ROI have higher fidelity than the rest of the image. For lossless coding of the region of interest, it is necessary to code those surrounding coefficients that take part in the reconstruction of the sample image values within the ROI. Figure 4 shows an example in one
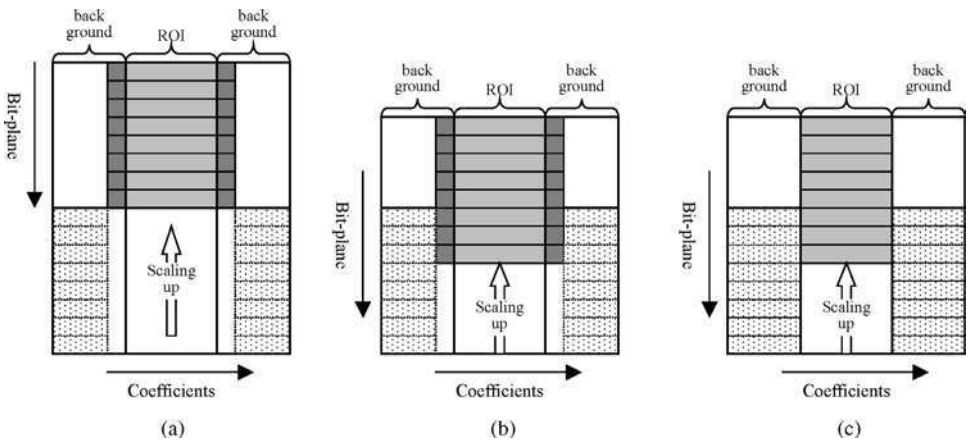
Fig. 3.   Coefficients bit-shifting in ROI coding: (a) Maxshift ROI, (b) scaling-based ROI method, and (c) object-based coding with shape-adaptive wavelet transform. The wavelet coefficients surrounding the ROI that are needed for the reconstruction of the image samples within the ROI are marked with dark hatch.
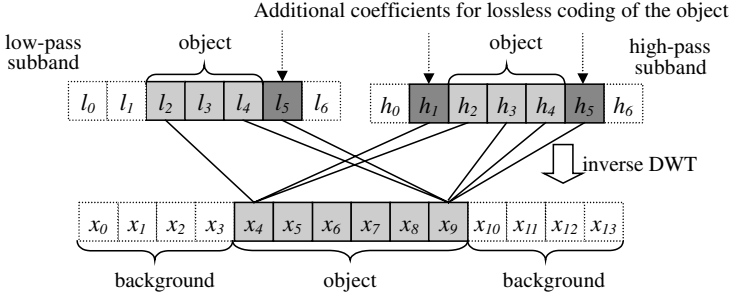
Fig. 4.   The inverse wavelet transform with the reversible 5/3 wavelet filter. Additional coefficients that are needed for the reconstruction of the image samples within the object are marked with dark hatch.

dimension of filtering corresponding high-pass and low-pass subbands using the reversible 5/3 wavelet filter[57] to reconstruct the image samples by

$$x_{2n} = l_n - \left\lfloor \frac{h_{n-1} + h_n + 2}{4} \right\rfloor \tag{4}$$

$$x_{2n+1} = h_n + \left\lfloor \frac{l_n + l_{n+1}}{2} \right\rfloor \tag{5}$$

Considering the inverse wavelet transform described above, it can be seen in Fig. 4 that coefficients l_5, h_1 and h_5 as well as the coefficients within the object are necessary to reconstruct samples from x_4 to x_9, where x_i, l_i and h_i represent values of an image sample, a low-pass subband coefficient and a high-pass subband coefficient, respectively. The additional coefficients depend on the filter length and the wavelet decomposition depth.

One of the advantages of this method is that it does not need to transmit the shape information and just sends the scaling value s, because the decoder can identify the wavelet coefficients scaled up just by comparing each coefficient with a threshold 2^s. However, with the code stream associated with the most significant s bit-planes, the object cannot be exactly decoded since the decoder cannot distinguish coefficients within the object from coefficients surrounding the object.[58]

6.2. *JPEG2000 scaling-based ROI coding*

In the scaling-based ROI coding adopted in JPEG2000 Part-2,[59] the entire image is transformed and the coefficients associated with the region of interest (within and around the region) are scaled up by a certain number of bit-shifts, as illustrated in Fig. 3(b). Then, the bit-planes of coefficients are encoded from most to least significant plane. Specifying the scaling value can control the difference of image quality between the ROI and the rest of the image. Although JPEG2000 Part-2 specifies scaling-based ROI coding only for rectangular or elliptic areas, the concept of this technique can be easily extended to encode regions of arbitrary shape. In the

scaling-based ROI coding, shape information has to be transmitted to the decoder unlike the Maxshift ROI method. Therefore, in scaling-based ROI coding, the object can be exactly decoded by discarding all of the background. But at the background surrounding the object, the additional coefficients still might cause unwanted effects at an early stage of progressive coding.[58]

6.3. *Object-based coding with shape-adaptive wavelet transform*

Object-based approaches are being studied as a new video coding paradigm,[53,60] in which only the image samples within an object are transformed using a shape-adaptive discrete wavelet transform,[54,61] the object shape information is additionally sent to the decoder, and the resulting coefficients are encoded. Such a coding paradigm may be more easily applied to still images where the object or objects are fixed in shape and position, being non-deformable and motionless. With this technique, the coding can be done to any desired accuracy and will not affect pixels outside the object. Figure 3(c) illustrates that scaling up coefficients only within the ROI does not involve scaling of adjacent background pixels, so that there are no undesirable effects in the background or the ROI.

The object-based coding of an image requires to apply a two-dimensional wavelet transform to a region with arbitrary shape and to efficiently encode the resulting wavelet coefficients. A two-dimensional region with arbitrary shape is comprised of several rows and columns of varying lengths. In a separable two-dimensional wavelet transform, transformation of rows in place is followed by transformation of the resulting columns of coefficients in place. So, for arbitrary-shaped regions, several one-dimensional wavelet transforms of different lengths must be enacted. Figure 5 shows an example of shape-adaptive discrete wavelet transform for a single row of pixels, in which only the hatched samples are transformed and both edges of the segment are extended by the symmetric extension. Different types of extensions, subsampling strategies, and known single coefficient padding, explained in Refs. 54 and 61 are needed to ensure reversibility for any data length (odd or even) and any starting position within a row or column. Note that, in a shape-adaptive wavelet
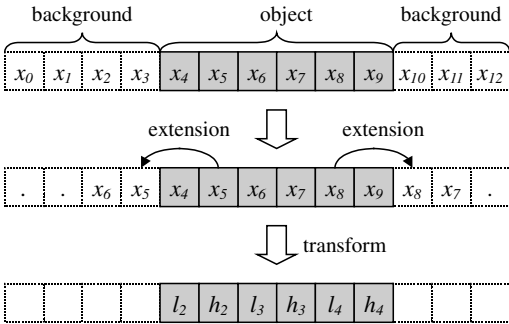

Fig. 5.   One-dimensional shape-adaptive wavelet transform with symmetric extension.

transform, object and background pixels are never mixed together in filtering and that the number of resulting transform coefficients is exactly the same as the number of pixels in the region, unlike the two previous ROI methods.

6.3.1. *Object-Based Set Partitioning In Hierarchical Trees (OBSPIHT)*

The SPIHT (Set Partitioning In Hierarchical Trees) method has been adapted for encoding arbitrarily shaped regions inside an image according to their importance or diagnostic relevance. This object-based adaptation of SPIHT, called OBSPIHT, applies a shape-adaptive wavelet transformation only to those pixels belonging to the region of interest instead of the whole image. Then, OBSPIHT uses a chain code method to transmit the shape information of the ROI. Finally, OBSPIHT progressively encodes only the wavelet coefficients associated to that region.

6.3.1.1. The Set Partitioning In Hierarchical Trees method

The SPIHT is an efficient wavelet-based embedded coding algorithm developed by Said and Pearlman,[30] which is representative of the state-of-the-art in image compression. This coding method is an extension and improvement over the Embedded Zerotree Wavelets (EZW) algorithm developed by Shapiro.[29]

The SPIHT method takes advantage of the spatial self-similarity across subbands inherent of the image wavelet transform, i.e. there are wavelet coefficients in different subbands of the transformed image that represent the same spatial localization in the original image. Studies in the literature have successfully applied SPIHT to lossy compression of medical images.[31,32,36]

When the distortion measure is MSE as in Eq. (3), it is invariant to the wavelet transform, i.e.

$$D_{MSE}(X - \hat{X}) = D_{MSE}(C - \hat{C}) = \frac{1}{N} \sum_{n=1}^{N} (c_n - \hat{c}_n)^2 \qquad (6)$$

where C and $\hat{C}$ represent the image after wavelet transformation and the wavelet transformed image reconstructed by the decoder, respectively. Therefore, to yield the largest distortion reduction at any compression ratio, the wavelet coefficients with larger magnitudes have to be transmitted first, as they have a larger content of information (here, the term information is interpreted to point out how much the distortion can be reduced after decoding that part of the bitstream).

In order to achieve a progressively refined reconstructed image with minimum reconstruction error (embedded method), SPIHT utilizes a bit-plane transmission scheme where wavelet coefficients are partially ordered according to its binary representation (analogous to its magnitude):

$$\lfloor \log_2 c_n \rfloor \geq \lfloor \log_2 c_{n+1} \rfloor \geq \cdots, \quad n = 1, \ldots, N \qquad (7)$$

and transmits first the most significant bits of the most important coefficients. The partial ordering by magnitude of wavelet coefficients is due to sequential comparisons of coefficient magnitudes to a set of decreasing thresholds: given k, a threshold value 2^k is set with the purpose of identifying those wavelet coefficients such that:

$$|c_n| \geq 2^k, \quad k = k_0, \ k_0 - 1, \ k_0 - 2, \ldots \tag{8}$$

Those coefficients with higher magnitude than a given threshold are called significant respect to that threshold value. The initial threshold k_0 is chosen verifying:

$$2^{k_0} \leq |c_n| < 2^{k_0+1}, \quad \forall c_n. \tag{9}$$

With the aim to efficiently determine significant coefficients with respect to a given threshold, SPIHT uses a Set Partitioning Sorting Algorithm that partitions the wavelet transform coefficients into sets that are significant or insignificant with respect to the current threshold. This Sorting Algorithm considers that, given a set of coefficients $\mathcal{U}$, then:

$$\begin{aligned} &\text{if } \max_{c_n \in \mathcal{U}} \{|c_n|\} \geq 2^k : \text{then } \mathcal{U} \text{ is a significant test} \\ &\text{otherwise} : \mathcal{U} \text{ is insignificant set} \end{aligned} \tag{10}$$

In case $\mathcal{U}$ is labelled as insignificant, all coefficients within $\mathcal{U}$ are insignificant, being unnecessary to sort these coefficients. Otherwise, a fixed rule is used to partition $\mathcal{U}$ into new subsets where the test is applied again for determining their significance. If one of the new subsets is labelled as insignificant, no more assessment is done. In the other hand, if the subset is significant, the partition procedure is recursively applied until the significant coefficients are determined.

The specific initial sets and partitioning rules of the SPIHT exploit the spatial self-similarity across subbands inherent in wavelet transformed images, and the fact that coefficients at low-resolution subbands are expected to have higher magnitudes than coefficients at high-resolution subbands. SPIHT considers that wavelet coefficients belonging to similar spatial orientation are organized into descendant trees, spanning wavelet subbands from the lowest spatial frequency (at the root) to the highest frequency subbands (at the leaves). In this tree-structure, called spatial orientation tree, each node corresponds to a wavelet coefficient and its descendants are those of the same spatial orientation in the next subband level. In the initial set grouping, the highest level of the decomposition pyramid is divided into blocks of 2×2 coefficients in which each coefficient, except for the upper left one, represents the root of a hierarchical tree: the upper-right coefficient is the root of the hierarchical tree in the $\mathcal{LH}$ subbands; the lower-left coefficient of the $\mathcal{HL}$ subbands; and the lower-right coefficient of the $\mathcal{HH}$ subbands (see Fig. 6). If a coefficient is insignificant with respect to a given threshold, all its descendants are likely to be insignificant with respect to the same threshold.

Each time the algorithm finds a node c_n whose set of descendants is significant, the tree-based partitioning rule decomposes this set into the four isolated offspring coefficients and all the descendants of the offspring (granddescendants),
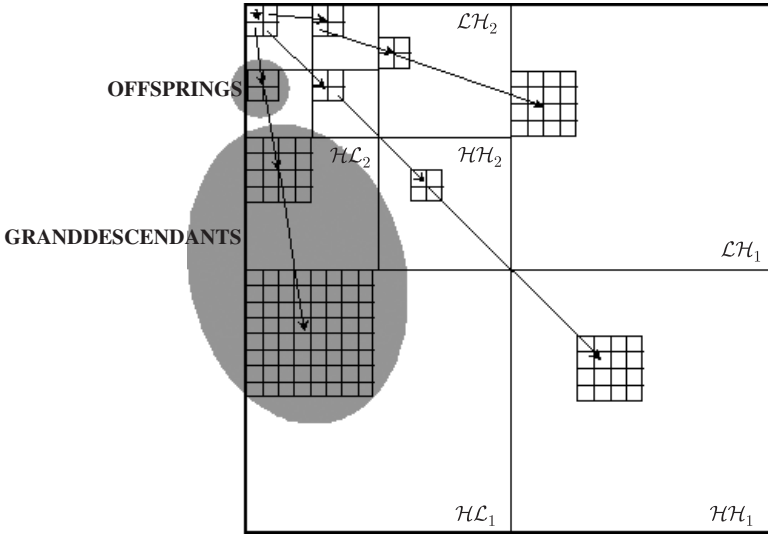
Fig. 6.   Spatial orientation tree in the SPIHT coding algorithm.

as shown in Fig. 6. The magnitude test is then applied to the new five subsets. If the granddescendant subset of coefficient c_n is determined as significant, the tree-based partitioning rule decomposes it into the four subsets of descendants of the offspring of c_n. This process continues until all significant coefficients with respect to the current threshold are identified. A more detailed description of the SPIHT algorithm is found in the original paper.[30]

7. Region-based Compression Methods for Digital Mammography

Compression of digital mammography would benefit from the possibility of determining several regions with different importance within an image before compression. For example, an expert radiologist or a CAD system can determine regions inside the mammogram containing primary radiological signs of breast cancer, such as clusters of microcalcifications and masses, as a preliminary step in a compression scheme for digital mammography. These areas can be compressed at a different ratio than the rest of the breast and, furthermore, the radiological background can be discarded.

Encoding techniques, such as OBSPIHT, permit the compression of multiple regions with precise bit rate control yielding different quality levels within the image. The feature of multiple quality levels is also provided in the general scaling-based method of the ROI coding mode supported in the JPEG2000, but it is not permitted in the Maxshift method.[51] In scaling-based methods, control of bit rate in ROI regions, although supported, but not specified in JPEG2000, would be inordinately cumbersome to implement. The standard JPEG2000 has also avoided

the overhead associated with the encoding of the shape information defining only rectangular and ellipse regions, although principles of the general scaling-based method can be extended to handle regions with arbitrary shape. However, a scaling coefficients procedure, especially in multiple regions, may require downward scaling of the background due to precision limitations. In that case, the least significant bitplanes are lost, causing additional quality degradation of the background.[50]

We present a region-based adaptation of SPIHT to digital mammography.[32]

7.1. *OBSPIHT adapted to digital mammography*

In a typical digitized mammogram much of the area in the image corresponds to the radiological background pixels rather than to tissue pixels. The radiological background does not provide useful information for diagnosis. To optimize compression in digital mammograms it is important to code the least information about background as possible. Detecting the breast border in the image, the OBSPIHT method can be applied considering the breast as the object of interest.

7.1.1. *Mammogram segmentation*

An automatic method is used for detecting the breast border[62]: (1) two cut-off grey levels are established to eliminate artefacts and get an homogenous background; (2) the thresholded image is smoothed; (3) five reference points, (x_1, y_1), (x_2, y_2), (x_3, y_3), (x_4, y_4), and (x_5, y_5), are automatically selected dividing the breast into three regions; (4) the breast border is detected with a tracking algorithm establishing that a point (x_i, y_i), belongs to the breast border if the grey level value, $f(x, y)$, of its nine previous image samples verify the condition:

$$f(x_{i-9}, y_{i-9}) < f(x_{i-8}, y_{i-8}) < \cdots < f(x_{i-3}, y_{i-3}) \leq f(x_{i-2}, y_{i-2})$$
$$\leq f(x_{i-1}, y_{i-1}) \leq f(x_i, y_i). \tag{11}$$

The tracking process is applied along different directions depending on the breast region. In region I the algorithm searches the breast border from left to right; in region II from top to bottom; and finally in region III from right to left (Fig. 7). The detection algorithm is relaxed such that the breast border obtained is always external to the real border.

Calculating the slope of the breast border in the reference points (x_1, y_1) and (x_2, y_2), the detected border is enlarged until it reaches the edge of the digital image. Hence, the breast region where the relevant information for diagnosis is included is completely determined within a closed contour. This contour permits to construct a binary mask determining the breast region within the image (Fig. 8). This mask is used in the next stages of the object-based compression method.

7.1.2. *Contour coding*

In our implementation, a two-link chain coding method presented in Ref. 63 is applied for coding the breast border. This two-link chain coding uses a contour
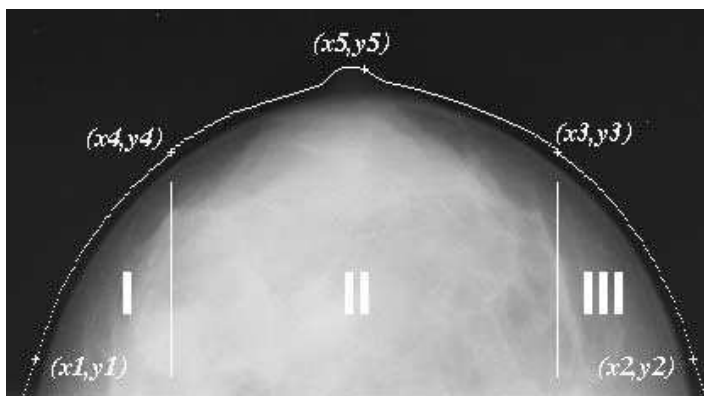
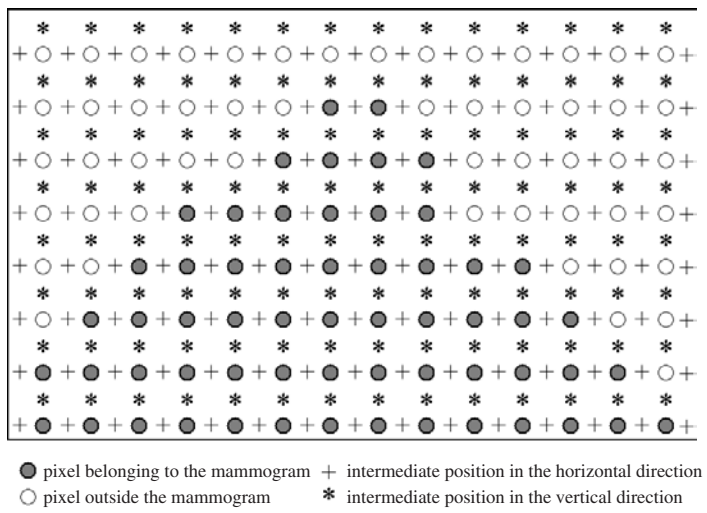Fig. 7. Closed contour determined in the segmentation process. Reference points dividing the breast into three regions are also shown.

Fig. 8. Breast border detection in a digital mammogram and generation of binary mask determining the region to be compressed.

lattice for representing the shape of the breast region and obtaining a more efficient coding of its contour. The contour lattice is formed by the half-pixel positions between two neighbouring pixels in the horizontal (+) and vertical (o) directions of the digital mammogram. The binary mask determining the breast region is represented on the basis of this structure as links between two points on the contour lattice (Fig. 9).

Since the breast has a smooth border, given a point of the contour lattice on the breast border, the number of the next possible links can be limited. This results in

* pixel belonging to the mammogram + intermediate position in the horizontal direction
○ pixel outside the mammogram * intermediate position in the vertical direction

Fig. 9.   Closed contour determined in the segmentation process. Reference points dividing the breast into three regions are also shown.

a reduction of the number of bits necessary to represent the border of the mammogram region.

7.1.3. *Region-based discrete wavelet transform*

The shape-adaptive discrete wavelet transform used to adapt OBSPIHT to digital mammograms, called region-based discrete wavelet transform (RBDWT), was proposed by Barnard.[61]

The binary mask obtained during the segmentation process specifies the region within the image where the RBDWT will be applied, in that case the breast region. Using separable wavelet filters, the two-dimensional RBDWT is achieved applying one-dimensional arbitrary length signal transform by rows and then by columns on the samples inside the image object. Efficient signal extensions allow decomposition up to an arbitrary level for arbitrary length signals with perfect reconstruction property. Moreover, the same number of wavelet coefficients than image samples will represent the signal after transformation.[61] The one-dimensional signal segments inside the arbitrary shape do not all start at an even- or odd-numbered position inside the row or column. Therefore, a proper subsampling strategy in the filtering process has to be chosen to preserve the spatial correlation within the image object. Depending on the parity of the segment length and the parity of the position at which the segment starts in the row or column, there are four classes of signal segments inside the object, which require different signal extension and subsampling.[61]
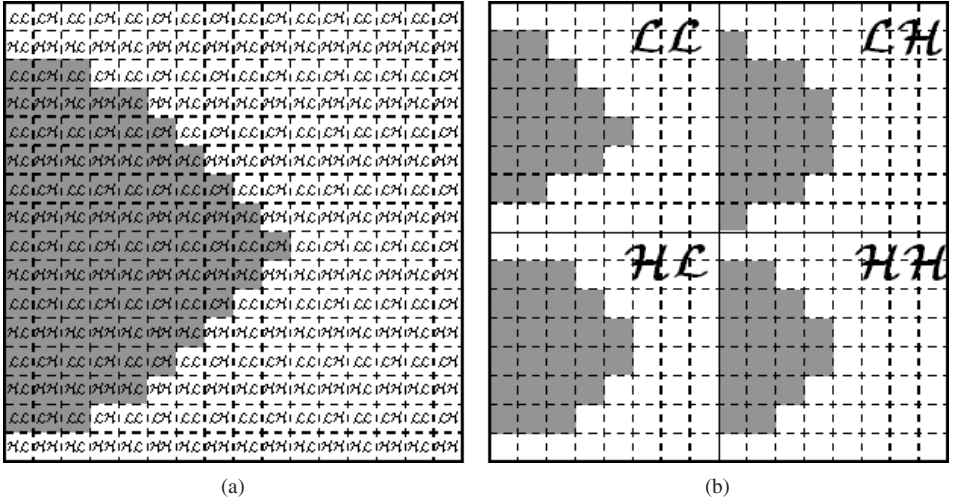
Fig. 10. Decomposition of the binary mask. (a) The original binary mask is divided in blocks of 2×2 pixels, each sample in the block corresponds to a submask; (b) the decomposed binary mask is split into four submasks ($\mathcal{LL}, \mathcal{LH}, \mathcal{HL}$ and $\mathcal{HH}$), analogously to the wavelet subbands.

Once the breast region is decomposed up to the first level, it is necessary to determine the wavelet coefficients in the $\mathcal{LL}$ subband where the RBDWT must be applied again to obtain the next decomposition level. This information is provided by decomposing the binary mask analogously to the wavelet subband decomposition of the image, as it is shown in Fig. 10. The $\mathcal{LL}$ submask specifies the region within the transformed object where the RBDWT will be applied again. The pyramidal decomposition of the binary mask can continue to get the required decomposition level of the breast region.

A 5-level dyadic decomposition of the breast region inside the mammogram was obtained with the two-dimensional RBDWT method using 9-tap/7-tap biorthogonal filters[64] (Fig. 11).

7.1.4. *Coding of the wavelet coefficients*

The binary mask provides the shape information needed for extending the conventional SPIHT algorithm. With the pyramidal shape mask constructed as it has been explained in Sec. 7.1.3, the nodes belonging to the breast region in each subband are known.

Figure 12 shows how the set partitioning produces branches to subsets falling outside the breast region. In the object-based extension, before the coding of the wavelet coefficients, the nodes and child branches outside the image object are pruned. During the coding process, no information about nodes and branches outside the breast region is transmitted.[63]
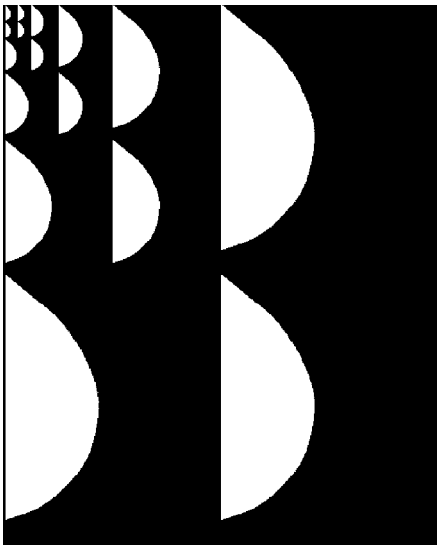
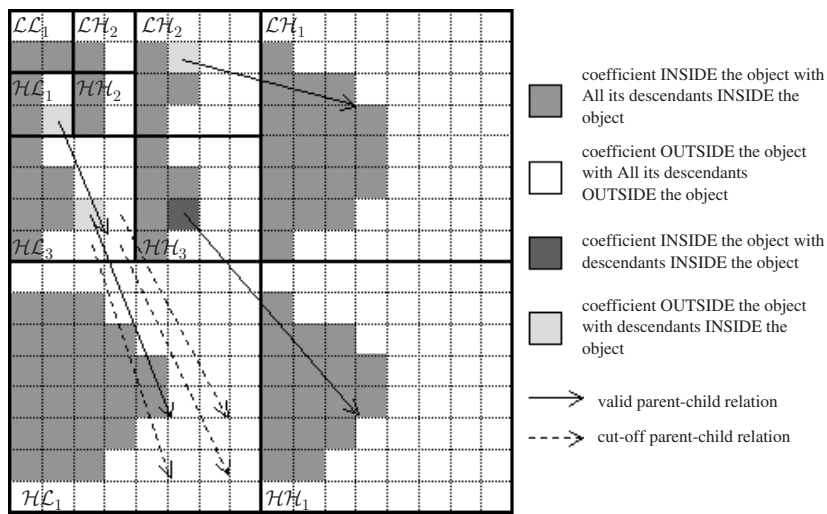Fig. 11. The binary mask of Fig. 8 decomposed up to the five level.

Fig. 12. Set Partitioning rules for the OBSPIHT coding method adapted to digital mammograms.

8. Studies in Region-based Wavelet Compression for Digital Mammography

In Refs. 24, 26, 31, 37, optimization of the compression process is obtained identifying the least rectangular area containing the breast region and compressing it to a different degree than the rest of the image. Recently, the object-based modification of the SPIHT method, OBSPIHT, has been adapted to digital mammography and

presented in Ref. 32. This region-based method compresses only the breast region inside the mammographic image. Quantitative evaluation of OBSPIHT showed an improvement in compression efficiency when using region-based methods compared to original SPIHT or JPEG2000. Qualitative evaluation of OBSPIHT in digital mammography has also been carried out, comparing results with the new standard JPEG2000.[65] Compression ratios up to 80:1 can be obtained with both methods without decreasing the diagnostic accuracy of digital mammograms for detecting important radiological signs.

8.1. *Example of quantitative evaluation*

The quantitative study evaluates the adaptation to digital mammography of the region-based coding method OBSPIHT by means of the PSNR (peak signal-to-noise ratio). The OBSPIHT has been compared to SPIHT and the JPEG2000 when compressing the full digital mammogram and the image with the radiological background previously set to zero. The former quantifies the performance improvement of region-based over full-image techniques; the latter is a comparison to a candidate region-based technique.

8.1.1. *The data set*

Five single-view conventional mammograms containing clusters of microcalcifications and masses, all biopsy proven, were used in this study. Mammograms were digitized with a commercially available laser film digitizer (LUMISCAN 85, Lumisys Inc.; CA, USA) at a resolution of 4,096 horizontal $\times$ 5,120 vertical pixels ($50\,\mu$m/pixel). The optical density range, from 0.03 to 4.1, was digitized to 12 bits precision, which provided 4,096 gray levels per pixel.

8.1.2. *Numerical quality evaluation*

Lossy compression performance of the OBSPIHT was evaluated at different rates (0.1, 0.2, 0.4, 0.6, 0.8, and 1.0 bpp) on the set of five digital mammograms by means of the PSNR. The coding results were compared with those obtained applying the original SPIHT algorithm and the standard JPEG2000[66] on the whole image. OBSPIHT was also compared to SPIHT and JPEG2000 applied to the same set of images with the background previously fixed to a constant value of zero. With these pre-processed images, both algorithms encode the breast region before the background pixels resulting, for this study, in similar implementations to that of OBSPIHT.

The SPIHT algorithm was utilized with the irreversible 9/7 biorthogonal filters[64] and with the reversible S+P filter[67] to examine lossy to lossless coding. Similarly, JPEG2000 was applied using the 5/3 reversible and the 9/7 irreversible filters.[a]

---

[a]JPEG2000 Verification Model Software 8.5, ISO/IEC JTC1/SC29/WG1 N1894, 2000.

Table 1. Average PSNR (dB) at various rates for all mammograms included in the study.

RATE	OBSPIHT	9/7 SPIHT	S+P SPIHT	9/7 JPG2K	5/3 JPG2K	9/7 SPIHTroi	9/7 JPG2Kroi	5/3 JPG2Kroi
0.1	48.22	46.37	46.41	46.91	46.46	47.89	48.33	48.06
0.2	50.64	47.81	47.84	48.39	47.92	50.38	50.97	50.62
0.4	54.77	49.79	49.45	50.08	49.44	54.49	55.35	54.93
0.6	58.80	50.95	50.75	51.27	50.69	58.48	59.58	58.90
0.8	63.16	51.92	51.78	52.38	51.76	62.81	64.04	62.97
1.0	67.57	52.86	52.93	53.44	52.64	67.04	68.36	67.50

In order to compare region-based to full-image coding methods, the PSNR measurements for all methods were calculated within the breast region where the object-based coding method was applied.

8.1.3. *Results and discussion*

On average, the region-based coding method OBSPIHT obtained better results in terms of PSNR than normal SPIHT and JPEG2000 for all bit rates tested (see Table 1). In fact, the average distortion achieved at around 0.6 bpp (26.7:1) with SPIHT or the standard JPEG2000 corresponds to a compression rate of 0.2 bpp (80:1) with OBSPIHT, making clear the improvement obtained with the object-based compression technique. The SPIHT and JPEG2000 methods applied to the images with background set to zero, which we call SPIHTroi and JP2Kroi, proved to be quite competitive in PSNR with the object-based coding method applied to the mammogram region. Actually, on average, the object-based method always introduced less distortion than SPIHTroi, but JP2Kroi was sometimes slightly better than OBSPIHT and SPIHTroi.

Magnified regions of interest containing a cluster of microcalcifications are shown in Fig. 13. Although the reconstructed ROIs are at the same rate, more rice artifacts are visible with general compression methods (SPIHT and JPEG2000) compared to the region-based OBSPIHT and the JP2Kroi methods (Fig. 13).

In this evaluation, OBSPIHT exhibited much higher quality in the breast region at the same compressed file size than full-image compression methods as SPIHT and the standard JPEG2000. Moreover, the SPIHTroi and JP2Kroi, both analogous to scaling-based compression methods applied to the breast region, were competitive in PSNR to OBSPIHT. Although they do not provide the possibility of encoding multiple regions of interest within the image, as the OBSPIHT. These results showed that region-based methods for digital mammography represent an improvement in compression efficiency from full-image methods.

8.2. *Example of qualitative evaluation*

In this study we assessed the effects of the OBSPIHT method on the detection of clusters of microcalcifications and masses on digitized mammograms, comparing
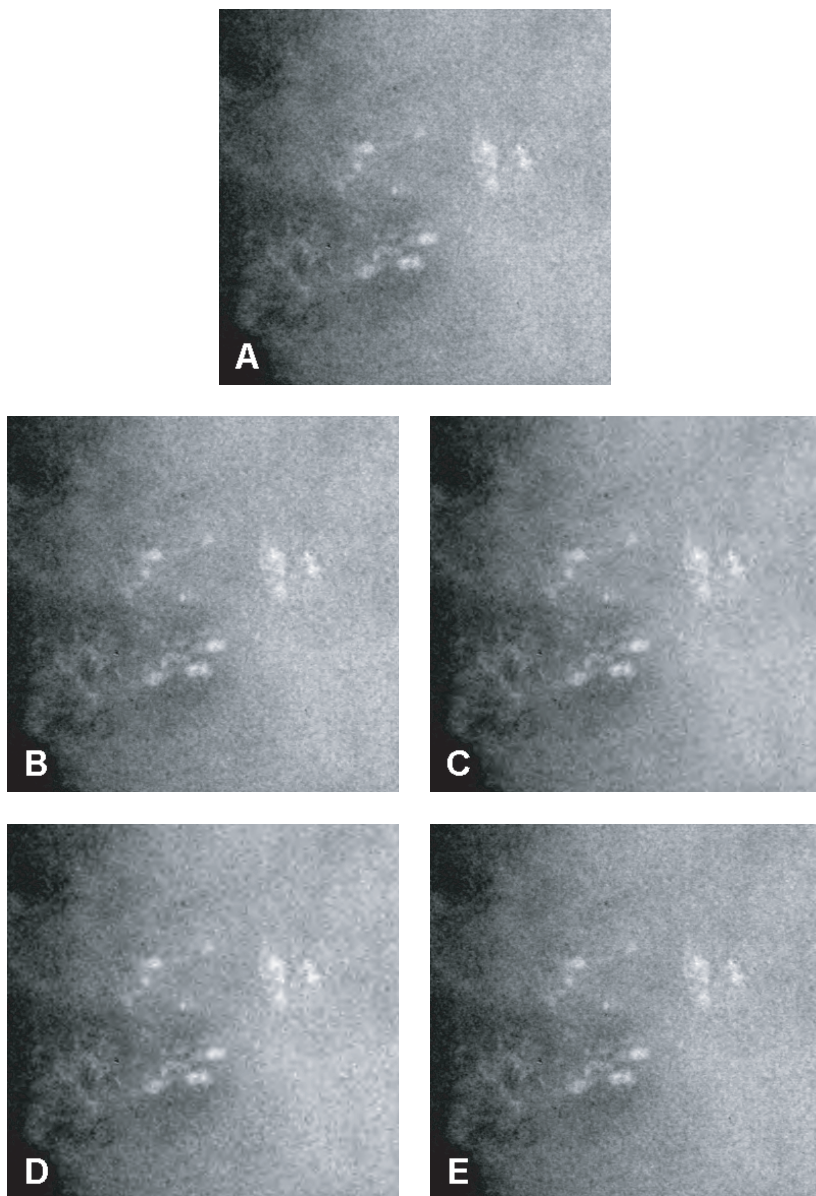
Fig. 13. (a) Originial ROI containing a cluster of microcalcifications; the same ROI is shown: (b) after OBSPIHT compression at 0.4 bpp, (c) after 9/7 SPIHT compression at 0.4 bpp, (d) after 9/7 JPEG2000 compression at 0.4 bpp, (e) after 9/7 JPEG2000 compression at 0.4 bpp of the preprocessed image with background set to zero (JP2Kroi). Rice artefacts are shown with general compression methods (c) and (d).

results with the full-image JPEG2000. The observer performance study took place at the Laboratory for Radiological Image Research at the University of Santiago de Compostela, Spain.

8.2.1. *The data set*

A total of 112 single-view conventional mammograms were used in this study. Of the 112 cases, 28 contained one or two clusters of microcalcifications, 19 mammograms contained one mass, 17 contained both signs and 48 were normal cases. The images presented a total of 54 clusters of microcalcifications (32 determined as malignant) and 36 masses (24 determined as malignant), all biopsy proven. The subtlety of the radiological signs was ranked by an experienced radiologist, not included as a reader in the study, on a four-level scale: 2 of the clusters were rated as obvious to detect, 21 as relatively obvious, 23 as subtle, and 8 as very subtle; 9 of the masses were rated as obvious to detect, 3 as relatively obvious, 16 as subtle and 8 as very subtle.

All images were digitized with a commercially available laser film digitizer (LUMISCAN 85, Lumisys Inc.; CA, USA) at a resolution of 4,096 horizontal $\times$ 5,120 vertical pixels (50 μm/pixel). The optical density range, from 0.03 to 4.1, was digitized to 12 bits precision, which provided 4,096 gray levels per pixel.

Each digitized mammogram was compressed at 0.4 and 0.2 bpp (respectively, 40:1 and 80:1 compression ratios) and then decompressed using the region-based method OBSPIHT and the standard JPEG2000.

8.2.2. *Observer performance study*

After reconstruction, a total set of 560 digital mammograms in five compression levels was obtained: 112 original digitized uncompressed, 112 compressed at 40:1 with JPEG2000, 112 compressed at 80:1 with JPEG2000, 112 compressed at 40:1 with OBSPIHT, and 112 compressed at 80:1 with OBSPIHT. All images were printed onto film with a high resolution DICOM printer (SCOPIX LR5200-P; Agfa-Gevaert, Mortsel, Belgium).

The 560 images were randomly assigned to 16 reading sets, avoiding repetition of the same image from any modality within the same set. Five radiologists with different degrees of experience in mammography, ranging from 10 to 12 years, participated in the study reviewing all images. Images were viewed on a standard light box under conditions of low ambient lighting. Readers could use a magnifying glass. Image sets and images within each set were presented in a different random order for each observer. The observers were told that cases were negative or had at least one biopsied lesion (cluster of microcalcifications, defined as five or more signals within a region of 1 cm^2 of area,[1] or mass). Observers were required to identify the location and extension of any cluster of microcalcifications or mass, and to record a level of confidence for each finding in a four-point FROC (Free-response Receiver

Operating Characteristic) rating scale: level 4, radiological sign definitely present; level 3, probably present; level 2, possibly present; and level 1, questionable. An independent expert mammographer established the truth location (as the distance in millimeters from a reference point) and extension (delimited with the closest contour surrounding it) of each finding, reviewing the patient's radiology and pathology reports. Following the methods of Kallergi *et al.*,[68] when an observer read an image and pointed a detected area approximately no more than four times and at least the 50% of the true area, the case was scored as a true-positive detection.

Since an observer saw each image five times during the FROC experiment (original uncompressed + four compressed versions), a time limit of 20 seconds per film was imposed on the observers to avoid learning effects. An assistant recorded the responses of the observers and removed the images from the light box when the time limit was exceeded. At least one week elapsed between each reading session to further reduce learning effects.

8.2.3. *Statistical data analysis*

Data for the 2,800 observations (5 modalities × 5 readers × 112 cases) were evaluated with the Jackknife Free-response Receiver Operating Characteristic[b] (JAFROC) analysis method.[69] In addition to allowing location of an abnormality to be accounted for the scoring, the method permits multiple responses and multiple lesions per image. This method uses as a figure of merit θ, which is defined as the probability that a lesion is rated higher than a false positive on a normal image. Recent extensive validation studies on simulated data sets have shown that the JAFROC method can yield substantially greater statistical power than the ROC method.[69,70]

8.2.4. *Results and discussion*

Tables 2–4 show the values of the FROC figure of merit (θ) quantifying the performance of the observers in the detection of clusters of microcalcifications, masses, and both radiological signs with the different image modalities tested: uncompressed digitized images, images compressed at 40:1 and 80:1 with OBSPIHT, and images compressed at 40:1 and 80:1 with JPEG2000. JAFROC analysis determined that differences between the five modalities were insignificant at the 5% level, for clusters of microcalcifications (F-stat = 0.2554; p = 0.9026), masses (F-stat = 0.3400; p = 0.8478) and both radiological signs detection tasks (F-stat = 0.1217; p = 0.9747).

Figure 14 shows a region containing a subtle cluster of microcalcifications from an uncompressed image, and the same region after irreversible compression using OBPIHT and JPEG2000 at 40:1 and 80:1. In Fig. 15, an uncompressed region

[b]The JAFROC software used in this study is available for download from the web site http://jafroc.radiology.pitt.edu.

Table 2. FROC figure of merit (θ) values quantifying observer performance in the detection of clusters of microcalcifications for all compression methods and ratios included in the study.

READER	ORIGINAL	JPEG2000 40:1	JPEG2000 80:1	OBSPIHT 40:1	OBSPIHT 80:1
1	0.83	0.80	0.77	0.79	0.79
2	0.85	0.88	0.82	0.83	0.84
3	0.82	0.86	0.91	0.91	0.87
4	0.74	0.72	0.74	0.73	0.77
5	0.75	0.78	0.73	0.78	0.74
Average	0.80	0.81	0.80	0.81	0.80

Table 3. FROC figure of merit (θ) values quantifying observer performance in the detection of masses for all compression methods and ratios included in the study.

READER	ORIGINAL	JPEG2000 40:1	JPEG2000 80:1	OBSPIHT 40:1	OBSPIHT 80:1
1	0.81	0.85	0.82	0.86	0.86
2	0.69	0.67	0.76	0.69	0.71
3	0.85	0.78	0.86	0.81	0.88
4	0.82	0.81	0.78	0.80	0.77
5	0.87	0.85	0.82	0.85	0.786
Average	0.81	0.79	0.81	0.80	0.82

Table 4. FROC figure of merit (θ) values quantifying observer performance in the detection of both radiological signs clusters of microcalcification and masses, for all compression methods and ratios included in the study.

READER	ORIGINAL	JPEG2000 40:1	JPEG2000 80:1	OBSPIHT 40:1	OBSPIHT 80:1
1	0.72	0.73	0.72	0.75	0.73
2	0.72	0.73	0.74	0.72	0.70
3	0.76	0.74	0.82	0.80	0.77
4	0.68	0.68	0.65	0.66	0.66
5	0.74	0.75	0.71	0.73	0.74
Average	0.72	0.73	0.73	0.73	0.72

containing a subtle mass and the corresponding regions after lossy compression with OBSPIHT and JPEG2000 at 40:1 and 80:1 are shown. In both examples, radiological signs are well preserved even at 80:1 compression.

The main conclusion from this analysis is that lossy compression at 80:1 using the OBSPIHT and the JPEG2000 can be used without decreasing the detection accuracy of important mammographic signs. However, as shown in the quantitative evaluation,[32] the region-based method OBSPIHT improves compression performance within the breast region compared to the standard JPEG2000 applied to the
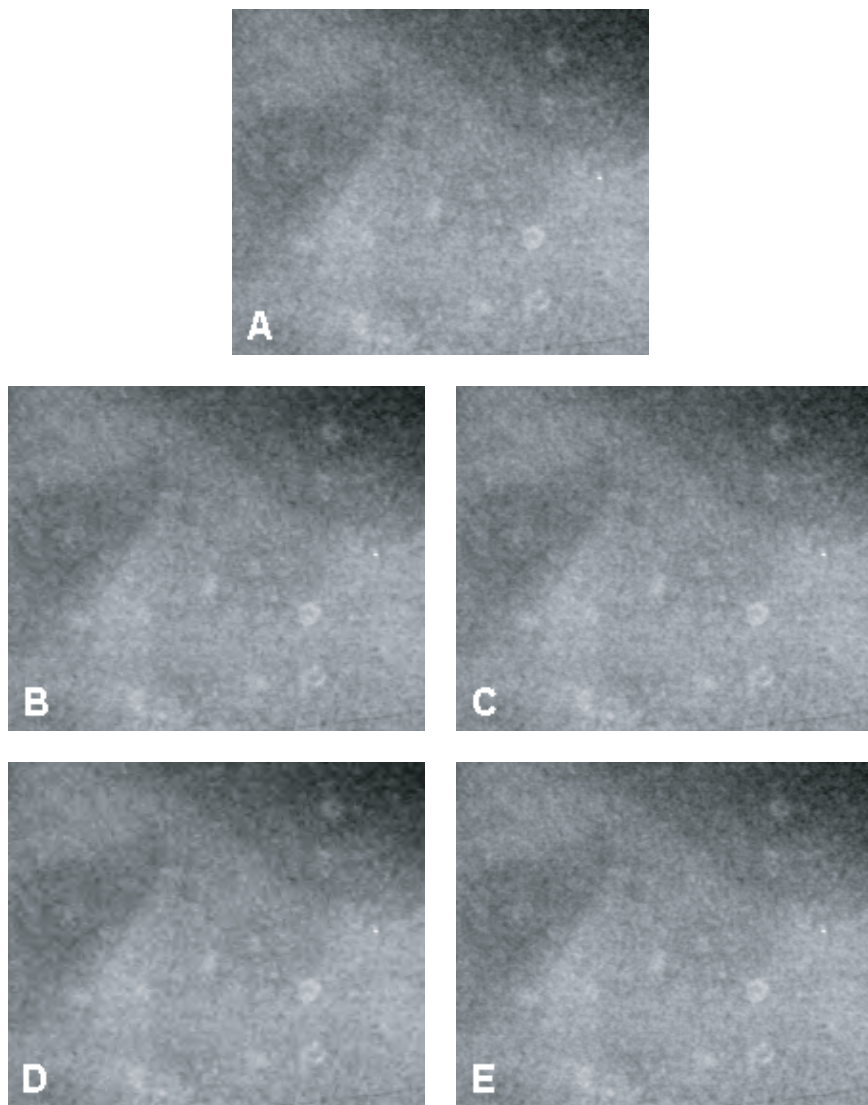
Fig. 14. Region of interest containing a cluster of microcalcification: (a) original region, (b) same region compressed at 40:1 with OBSPIHT, (c) at 40:1 with JPEG2000, (d) at 80:1 with OBSPIHT, and (e) at 80:1 with JPEG2000.

whole image. The compression efficiency of the region-based method OBSPIHT over a full-image method as the standard could be more perceptible when using high-resolution monitors with image visualization tools (such as window/level adjustments, pan and zoom functions, or rotating and flipping images) for displaying the digital mammographic study.
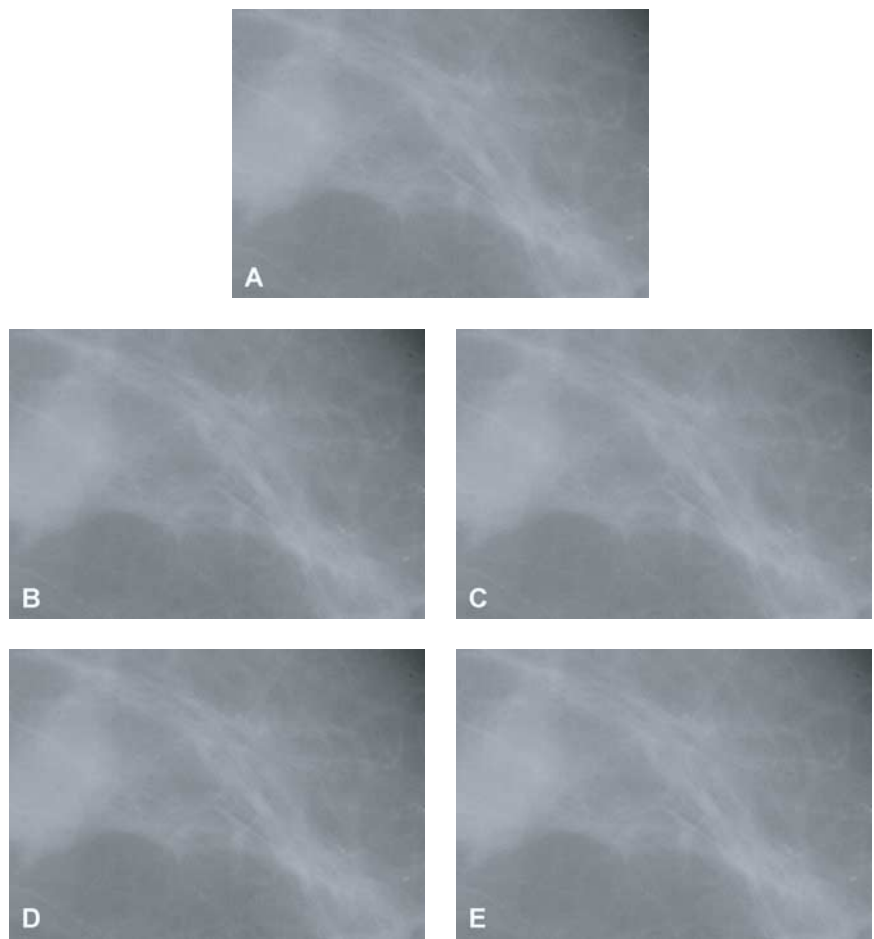
Fig. 15.   Region of interest containing a subtle mass: (a) original region, (b) same region compressed at 40:1 with OBSPIHT, (c) at 40:1 with JPEG2000, (d) at 80:1 with OBSPIHT, and (e) at 80:1 with JPEG2000.

9. Concluding Remarks

The emerging region-based wavelet compression methods are very suitable for medical imaging providing the possibility of compressing several regions at different quality levels according to their diagnostic relevance. We have seen, as a first approximation, how to adapt state-of-the-art ROI compression methods, such as OBSPIHT, to compress only the breast region in a digital mammogram. Application of the region-based method OBSPIHT to several regions within the digital mammogram is straightforward. An expert radiologist or an automatic CAD (computer-aided detection) system can determine these regions beforehand.

The importance of digital mammography in the screening and diagnosis of breast cancer requires to evaluate extensively any compression method that would be

applied to this image modality. The validation for the clinical practice of a region-based compression method for digital mammography should be based on research studies such as the presented in this chapter. Our results showed that detection of important radiological signs is not degraded, even at substantial compression rates. The OBSPIHT compression method could be used in digital technologies such as PACS and telemammography, reducing storage and transmission costs and facilitating digital mammography as a medical imaging modality.

This evaluation, however, does not establish the upper limit where the compression technique can be used so far without losing relevant information. Further assessment of the region-based compression method for its application with computerized detection of radiological signs of breast cancer or evaluating compressed images on high-resolution display workstations should also be carried out.

Acknowledgments

A. Appendix

In this section a brief background on wavelet analysis is introduced. General principles and further details of wavelet theory are explained in Ref. 71.

A.1. *The wavelet transform*

The wavelet transform (WT) decomposes a signal f over wavelet functions obtained as translations u and dilations s of a mother wavelet function ψ of zero average:

$$\mathrm{W}f(u, s) = \int_{-\infty}^{+\infty} f(t) \frac{1}{\sqrt{s}} \psi^* \left(\frac{t - u}{s} \right) \, \mathrm{d}t. \tag{A1}$$

Shifting the translation u and scale s parameters, the wavelet transform, unlike the Fourier Transform, provides a time-frequency representation of the signal. That powerful characteristic led the WT to become a technique used for non-stationary signals in several applications, including biomedicine.[16] Additionally, sampling the translation and scale parameters as $u = 2^j n$ and $s = 2^j$, it is possible to construct

a discrete orthogonal wavelet basis holding the signal details necessary to obtain a finer resolution representation of the signal, which is related with a multiresolution representation.

In a multiresolution approximation, approaches of a signal at different resolutions are determined with a discrete orthogonal basis obtained as dilations and translations of a scaling function ϕ. This multiresolution representation is completely specified by a discrete conjugate mirror filter h. It has been demonstrated that an orthogonal wavelet basis is constructed with a mother wavelet ψ, which is derived from ϕ and h, providing the detailed information lost when passing the signal to a coarser resolution representation. The orthogonal wavelet is designed with a conjugate mirror filter g given by:

$$g[n] = (-1)^{1-n}\, h[1-n]. \tag{A2}$$

A signal f at a scale 2^j is then represented in a coarse resolution as:

$$f = \sum_{n=-\infty}^{+\infty} a_{j+1}[\mathrm{n}]\phi_{j+1,n} + \sum_{n=-\infty}^{+\infty} d_{j+1}[\mathrm{n}]\psi_{j+1,n} \tag{A3}$$

where $a_{j+1}[n] = \langle f, \phi_{j+1,n}\rangle$ and $d_{j+1}[n] = \langle f, \psi_{j+1,n}\rangle$ are the approximation coefficients and the wavelet coefficients of f at a coarse resolution, respectively.

The fast implementation of the discrete wavelet transform is computed with a filter bank, which decomposes the signal with these conjugate mirror filters h and g, respectively low and high pass filters, and subsamples the output by 2 (Fig. A1). Extension to multiple dimensions is easily obtained with separable wavelet filters, which extracts signal details at different scales and orientations, applying conjugate mirror filters along each dimension. A schematic diagram of the wavelet decomposition of an image is shown in Fig. A2.

In a wavelet-based lossy compression method, the original image is first transformed and once in the wavelet domain, coefficients are quantized (represented with a less number of bits, which incurs in a loss of information) and entropy coded (coded with the minimum number of bits required), obtaining a compressed file. The decompression procedure inverts all this steps in reverse order, obtaining the
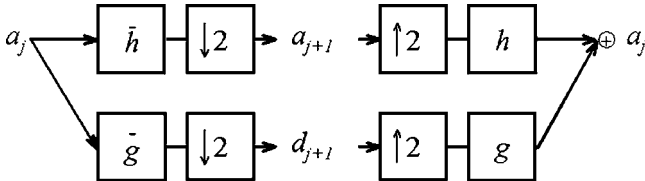

Fig. A1. Fast implementation of the one-dimensional wavelet transform: decomposition of a_i is computed with a cascade of filtering followed by a factor 2 downsampling ($\downarrow$2); reconstruction of a_i is done by inserting zeros between samples of a_{i+1} and d_{i+1}, filtering and adding up the output.

First decomposition level
N/2 x N/2

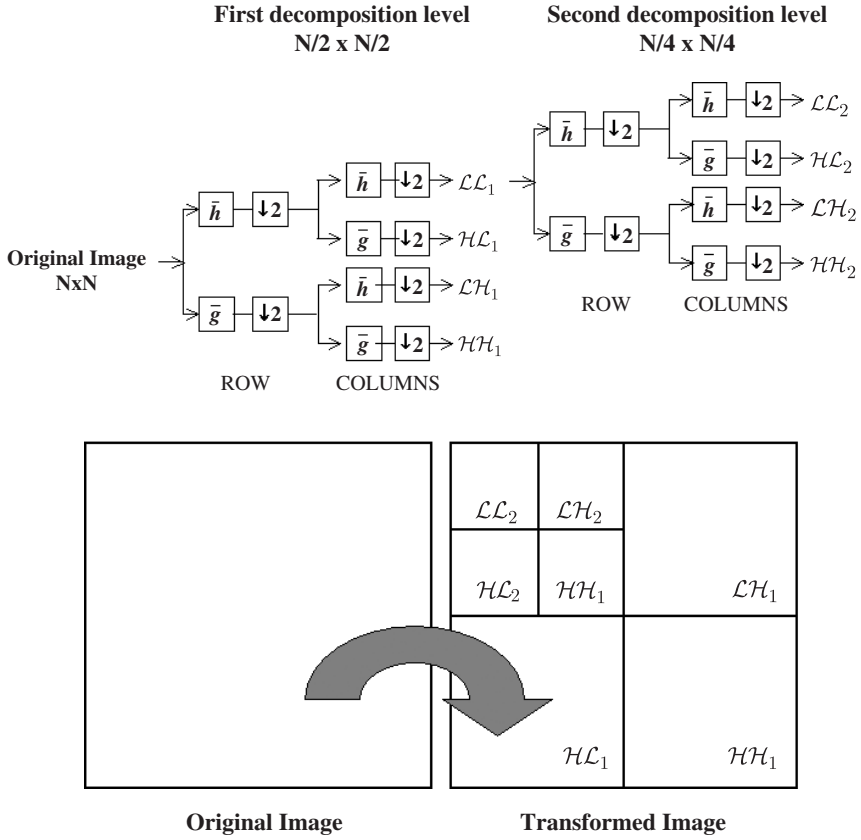**Second decomposition level**
N/4 x N/4

Fig. A2.   Schematic diagram of the decomposition two-dimensional dyadic wavelet transform. In the original image, first each row is filtered and subsampled by 2, then, each column is filtered and subsampled by 2. Four subimages are obtained, called wavelet subbands, referred to as HL, LH, HH (high frequency subbands), and LL (low frequency subband). The LL subband is again filtered and subsampled to obtain four more subimages. This process can be repeated until the desired decomposition level.

reconstructed image, which is not exactly the original image due to the quantization step. Further performance improvement is obtained with embedded techniques, which combine in the compression process the quantization and coding strategies that add in characteristics of the wavelet decomposition, such as the SPIHT algorithm and the JPEG2000. These coding methods utilize a bit-plane coding technique to produce a bitstream that can be truncated at any point (equivalent to stopping the compression process at a desired compression rate), whereupon the image can be reconstructed with the minimum distortion.

Numerical studies have shown that the 9/7 biorthogonal filters provide the best distortion rate performance for wavelet-based lossy compression of images.[64,71,72]

References

1. D. B. Kopans, *Breast Imaging* (Lippincott-Raven Publishers, Philadelphia, PA, 1998).
2. *Premarket Applications for Digital Mammography Systems; Final Guidance for Industry and FDA*, Center for Devices and Radiological Health of the Food and Drug Agency. US Department of Health and Human Services (2001).
3. J. J. James, *Clinical Radiology* **59** (2004) 1–10.
4. O. Jarlman, L. Samuelsson and M. Braw, *Acta Radiologica* **32** (1991) 110–113.
5. D. S. Brettle, S. C. Ward, G. J. Parkin, A. R. Cowen and H. J. Sumsion, *British Journal of Radiology* **67** (1994) 464–468.
6. A. R. Cowen, G. J. Parkin and P. Hawkridge, *European Radiology* **7** (1997) 918–930.
7. L. A. Venta, R. E. Hendrick, Y. T. Adler, P. DeLeon, P. M. Mengoni, A. M. Scharl, C. E. Comstock, L. Hansen, N. Kay, A. Coveler and G. Cutter, *Am. J. Roentgenology* **176** (2001) 1241–1248.
8. E. A. Berns, R. E. Hendrick and G. R. Cutter, *Medical Physics* **29** (2002) 830–834.
9. S. Suryanarayanan, A. Karellas, S. Vedantham, H. Ved, S. P. Baker and C. J. D'Orsi, *Radiology* **225** (2002) 801–807.
10. E. D. Pisano, E. B. Cole, E. O. Kistner, K. E. Muller, B. M. Hemminger, M. L. Brown, R. E. Johnston, C. M. Kuzmiak, M. P. Braeuning, R. I. Freimanis, M. S. Soo, J. A. Baker and R. Walsh, *Radiology* **223** (2002) 483–488.
11. S. Obenauer, S. Luftner-Nagel, D. von Heyden, U. Munzel, F. Baum and E. Grabbe, *European Radiology* **12** (2002) 1697–1702.
12. R. E. Bird, T. W. Wallace and B. C. Yankaskas, *Radiology* **184** (1992) 613–617.
13. H. P. Chan, L. T. Niklason, D. M. Ikeda, K. L. Lam and D. D. Adler, *Medical Physics* **21** (1994) 1203–1211.
14. M. D. Harpen, *Medical Physics* **25** (1998) 1985–1993.
15. D. F. Schomer, A. A. Elekes, J. D. Hazle, J. C. Huffman, S. K. Thompson, C. K. Chui and W. A. Murphy, *Radiographics* **18** (1998) 469–481.
16. M. Unser and A. Aldroubi, *Proc. IEEE* **84** (1996) 626–638.
17. P. Hansgen, P. E. Undrill and M. J. Cree, *Computer Methods and Programs in Biomedicine* **56** (1998) 1–10.
18. P. Saipetch, B. K. T. Ho, R. Panwar, M. Ma and J. Wei, *IEEE Eng. Med. Biol. Magazine* **14** (1995) 587–593.
19. S. K. Thompson, J. D. Hazle, D. F. Schomer, A. A. Elekes, D. A. Johnston, J. Huffman and C. K. Chui, *Medical Physics* **27** (2000) 276–288.
20. B. Erickson, A. Manduca, P. Palisson, K. Persons, F. Earnest, V. Savcenko and N. Hangiandreou, *Radiology* **206** (1998) 599–607.
21. A. Munteanu, J. Cornelis and P. Cristea, *IEEE Trans. Med. Imaging* **18** (1999) 272–281.
22. A. Baskurt, I. E. Magnin and R. Goutte, *Optical Engineering* **31** (1992) 1922–1928.
23. H.-P. Chan, S.-C. B. Lo, L. T. Niklason, D. M. Ikeda and K. L. Lam, *Medical Physics* **23** (1996) 1325–1336.
24. W. F. Good, J. H. Sumkin, M. Ganott, L. Hardesty, B. Holbert, C. M. Johns and A. H. Klym, *Am. J. Roentgenol* **175** (2000) 1573–1576.
25. G. R. Kuduvalli and R. M. Rangayyan, *IEEE Trans. Med. Imaging* **11** (1992) 430–445.
26. B. Zheng, J. H. Sumkin, W. F. Good, G. S. Maitz, Y. H. Chang and D. Gur, *Acad. Radiol.* **7** (2000) 595–602.
27. W. B. Richardson, *IEEE Eng. Med. Biol. Magazine* **14** (1995) 551–560.
28. Z. S. Yang, M. Kallergi, R. A. Devore, B. J. Lucier, W. Qian, R. A. Clark and L. P. Clark, *IEEE Eng. Med. Biol. Magazine* **14** (1995) 570–577.

29. J. M. Shapiro, *IEEE Trans. Signal Processing* **41** (1993) 3445–3462.
30. A. Said and W. Pearlman, *IEEE Trans. Circuits Syst. Video Techn.* **6** (1996) 243–250.
31. S. Perlmutter, P. Cosman, R. Gray, R. Olshen, D. Ikeda, C. Adams, Betts BJ, M. Williams, K. Perlmutter, J. Li, A. Aiyer, L. Fajardo, R. Birdwell and B. Daniel, *Signal Process* **59** (1997) 189–210.
32. M. Penedo, W. A. Pearlman, P. G. Tahoces, M. Souto and J. J. Vidal, *IEEE Trans. Med. Imaging* **22** (2003) 1288–1296.
33. M.-M. Sung, H.-J. Kim, E.-K. Kim, J.-Y. Kwak, J.-K. Yoo and H.-S. Yoo, *IEEE Trans. Nucl. Sci.* **49** (2002) 827–832.
34. S. Suryanarayanan, A. Karellas, S. Vedantham, S. M. Waldrop and C. J. D'Orsi, *J. Digital Imaging* **17** (2004) 64–70.
35. H. MacMahon, K. Doi, S. Sanada, S. Montner, M. Giger, C. Metz, N. Nakamori, F. Yin, X. Xu and H. Yonekawa, *Radiology* **178** (1991) 175–179.
36. V. Savcenko, B. J. Erickson, P. M. Palisson, K. R. Persons, A. Manduca, T. E. Hartman, G. F. Harms and L. R. Brown, *Radiology* **206** (1998) 609–616.
37. O. Kocsis, L. Costaridou, L. Varaki, E. Likaki, C. Kalogeropoulou, S. Skiadopoulos and G. Panayiotakis, *European Radiology* **13** (2003) 2390–2396.
38. *Guidance for the Submission of Premarket Notifications for Medical Image Management Devices*, Center for Devices and Radiological Health of the Food and Drug Agency. US Department of Health and Human Services (2000).
39. *ACR Technical Standard for Teleradiology*, American College of Radiology (2002).
40. C. J. Vyborny and M. L. Giger, *American J. Roentgenology* **162** (1994) 699–708.
41. C. E. Metz, *Seminars in Nuclear Medicine* **8** (1978) 283–298.
42. N. A. Obuchowski, M. L. Lieber and K. A. Powell, *Academic Radiology* **7** (2000) 516–525.
43. R. G. Swensson, *Medical Physics* **23** (1996) 1709–1725.
44. K. O. HajianTilaki, J. A. Hanley, L. Joseph and J. P. Collet, *Academic Radiology* **4** (1997) 222–229.
45. C. M. Rutter, *Academic Radiology* **7** (2000) 413–419.
46. D. Chakraborty, *Medical Physics* **16** (1989) 561–568.
47. D. Chakraborty and L. Winter, *Radiology* **174** (1990) 873–881.
48. A. Jarvi, J. Lehtinen and O. Nevalainen, *Signal Processing-Image Communication* **14** (1999) 683–696.
49. E. Atsumi and N. Farvadin, *Proc. IEEE Int. Image Processing Conf.* (1998) 87–91.
50. C. Christopoulos, J. Askëlof and M. Larsson, *IEEE Signal Process Lett.* **7** (2000) 247–249.
51. M. Rabbani and R. Joshi, *Signal Process: Image Comm.* **17** (2002) 3–48.
52. N. D. Black, R. J. Millar, M. Kunt, M. Reid and F. Ziliani, *Advances in Imaging and Electron Physics*, **112** (2000) 1–54.
53. G. Minami, Z. X. Xiong, A. Wang and S. Mehrotra, *IEEE Trans. Circuits Syst. Video Technol.* **11** (2001) 1063–1068.
54. S. Li and W. Li, *IEEE Trans. Circuits Syst. Video Techn.* **10** (2000) 725–743.
55. H. Freeman, *IRE Trans. Electron. Comput.* **10** (1961) 260–268.
56. *JPEG2000 Image Coding System*, ISO/IEC JTC1/SC29/WG10 15444-1 (2000).
57. D. Taubman and M. W. Marcellin, *JPEG2000: Image Fundamentals, Standards and Practice* (Kluwer Academic Publishers, 2002).
58. I. Ueno and W. A. Pearlman, *Proc. SPIE 5022, SPIE/IS&T Electronic Imaging* (2003) 1048–1055.
59. *JPEG2000 Image Coding System: Extensions*, ISO/IEC JTC1/SC29/WG10 15444-2 (2000).

60. Z. Lu and W. Pearlman, in Picture Coding Symposium (PCS-2001), 413–416 (2001).
61. H. J. Barnard, *Image and Video Coding Using a Wavelet Decomposition*, Dept. Electrical Engineering, Delft Univ. Technol., Delft, The Netherlands (1994).
62. A. J. Mendez, P. G. Tahoces, M. J. Lado, M. Souto, J. L. Correa and J. J. Vidal, *Computer Methods & Programs in Biomedicine* **49** (1996) 253–262.
63. Z. Lu and W. Pearlman, in Proc. SPIE (2002).
64. M. Antonini, M. Barlaud, P. Mathieu and I. Daubechies, *IEEE Trans. Image Processing* **1** (1992) 205–220.
65. M. Penedo, M. Souto, P. G. Tahoces, J. M. Carreira, J. Villalón, G. Porto, C. Seoane, J. J. Vidal, K. S. Berbaum, D. P. Chakraborty and L. L. Fajardo, Radiology (in press).
66. J. Askëlof, M. Carlander and C. Christopoulos, *Signal Process: Image Comm.* **17** (2000) 105–111.
67. A. Said and W. Pearlman, *IEEE Trans. Image Processing* **5** (1996) 1303–1310.
68. M. Kallergi, G. M. Carney and J. Gaviria, *Medical Physics* **26** (1999) 267–275.
69. D. P. Chakraborty and K. S. Berbaum, *Medical Physics* **31** *in press* (2004).
70. D. Chakraborty, *Acad Radiol* **9** (2002) 147–156.
71. S. G. Mallat, *A Wavelet Tour of Signal Processing* (Academic Press, San Diego, 1999).
72. J. Villasenor, B. Belzer and J. Liao, *IEEE Trans. Image Processing* **4** (1995) 1053–1060.

CHAPTER 7

TECHNIQUES IN SEGMENTING 3D IMAGES WITH ANISOTROPIC SPATIAL RESOLUTION AND FOR TRACKING TEMPORAL IMAGE SEQUENCES AND THEIR APPLICATION

XOSÉ M. PARDO* and DAVID L. VILARIÑO†

Dept. Electrónica e Computación
Universidade de Santiago de Compostela
15782 Santiago de Compostela, Spain
*pardo@dec.usc.es
†dlv@dec.usc.es

Three-Dimensional segmentation and tracking of anatomical organs is widely used for diagnosis of diseases and planning of surgical interventions. However, these are difficult tasks due to the complexity and variability of the anatomic shapes, sampling artifacts and noise. Moreover, 3D medical image modalities provide voxel datasets in which the sampling is normally uniform along each axis though the voxels themselves tend to have anisotropic resolution. Where 3D image data sets have no isotropic spatial resolution, features cannot be extracted in 3D, and therefore a slice-to-slice framework is more adequate.

Deformable models have been developed for segmentation and tracking. In this work two active contour approaches are reviewed and applications include automatic contouring of organs and tracking of organ motion. The first active contour is devoted to tackle segmentation tasks where the desired features of the object of interest could slightly change from slice to slice, while the second active contour is dedicated to tracking tasks where the features of the object of interest do not change from frame to frame, but the shapes do. The segmentation approach aims to learn and update the best discriminants from one slice to the next. The tracking approach, however, assumes well established goal features and intends to reach them as fast as possible.

Keywords: Segmentation; tracking of anatomical organs.

1. Introduction

An important property of 3D medical images is that they are discrete samples of an object at a finite number of voxels in three dimensions. A medical image might be made of a single slice, a series of parallel slices with uniform spacing, a series of slices with varying spacing and/or orientation, or volumes. But multislice images can also be made up of non-parallel slices. For example, the majority of ultrasound images are acquired with a free-hand transducer. If the transducer is moved in a controlled way or tracked, the relative positions of the frames can be recorded, and a three-dimensional dataset obtained.

The process of forming a discrete image from a continuous object is called sampling. In some modalities of medical image, data are acquired from an entire

volume, but usually, a stack of individual slices is acquired. The gap between the centre of adjacent samples in a given direction is the sampling interval, which is often anisotropic. For instance, the sampling interval in the through-slice direction in tomographic images is often larger than in the slice plane.

A 3D image typically has a large number of voxels and is very compute intensive for processing such as segmentation and tracking. While the raw image can be readily displayed as 2D slices, 3D analysis and visualization requires explicitly defined object boundaries, especially when creating 3D surface models. For example, to create a 3D rendering of a human brain from a MR image, the brain needs to be identified first within the image and then its boundary marked and used for 3D rendering. Image segmentation identifies the attributes of voxels and defines the boundaries for voxels that belong to the same group. Additionally, measurements and quantitative analysis for parameters such as area, perimeter, volume and length, can be computed when object boundaries are defined.

In medical imaging, the subject being imaged has three spatial dimensions, but it also changes with time. For instance, clinicians use ultrasound imaging to generate real-time movies of left ventricular motion to assess if a heart is functioning well. To automatically derive quantitative measures of heart function is a challenge because cardiac ultrasound image sequences are noisy and good images difficult to acquire. Tracking is the task of estimating the position of an object along a temporal sequence. In medical image analysis, techniques for segmenting of anisotropic spatial sequences and for tracking of temporal sequences, usually, come close. In that context, the goal of segmentation is to find out the boundaries of one or more target objects along the spatial sequence of slice images, meanwhile tracking is aimed to follow the movement of one or more objects from the same point of view. Where spatial proximity can be assumed in a 3D anisotropic sequence, the appearance of two adjacent slice images is similar, and the same happens along a temporal sequence if the steps in time are not too long.

Because of the importance of identifying objects from an image, there have been extensive research efforts on image segmentation and tracking for the past decades. A number of image segmentation/tracking methods have been developed using fully automatic or semi-automatic approaches for medical imaging and other applications.

Before introducing a brief revision of segmentation/tracking techniques we report a final reason for performing a slice to slice processing rather than dealing with the whole bulk of data.

1.1. *Scene versus object based interpolation*

Usually, 3D medical image modalities provide voxel datasets in which the sampling is normally uniform along each axis ($256 \times 256 \times 124$ for a typical MR image volume), though the voxels themselves tend to have anisotropic resolution (e.g. $0.8\,\text{mm} \times 0.9\,\text{mm} \times 2\,\text{mm}$). The third dimension is often treated differently from

the first two, so a 3D image is commonly treated as a sequence of 2D images. In general, it is safe to say that high image quality can often be achieved only at the expense of patient throughput, and a difficult decision is the selection of the most appropriate tradeoff between quality and throughput.

Algorithms for volume analysis are often based on filtering techniques which presume isotropy. Interpolation is required whenever the acquired image data are not at the same level of discretization as the level that is desired. For an appropriate medical treatment it is necessary that the data at requested locations are reconstructed as precisely as possible taking into account not only the characteristics of a 3D signal conveying the data but also the topological properties of the explored structures. Interpolation algorithms can broadly be divided into two categories: image-based and shape-based.

The simplest image based interpolation method involves linearly interpolating the grey values in the slices to fill in the grey values in the missing slices. However, artifacts (diffuse contours and grey level discontinuities) are produced when the contour locations on two adjacent slices shift considerably. To reduce these artifacts, higher order functions have been proposed to perform interpolation. Several approaches are based on the matching between feature points on different slices, and then the intermediate slices are generated from the interpolation of locations and values of matched points.[19] Figure 1 illustrates an image based interpolation.

From the beginning shape-based algorithms required a description of objects to be interpolated. The first approaches were based on contour interpolation of the 2D segmented regions. Raya and Udupa[47] proposed an interpolation of distance fields computed from binary slices. The method consists of first segmenting the given image data into a binary image, converting the binary image into a distance map from the slice boundary, and then interpolating between each pair of adjacent distance maps. The zero-valued voxels define the object boundaries. Turk and O'Brien[60] proposed a method which combines the generation of distance functions (contour description) and the interpolation between these two functions into a single step. Given two shapes in the plane, a set of boundary and normal constraints are generated for each shape, and instead of using each set of constraints separately to create the two different 2D implicit functions, the method embeds all the constraints in 3D and invokes 3D variational interpolation to create a single scalar-valued function in 3D.

Although image-based filters have received a lot of attention in the visualization community, there is repeated evidence in the literature of the superior performance of shape-based over scene-based interpolation techniques.

Figure 2 illustrates a shape-based interpolation of a femur in a anisotropic 3D CT image sequence. The figure in the top shows the voxels of bone regions after eliminating other tissue voxels. It can be noted the anisotropic sampling. The figure at the bottom shows the geometric reconstruction of the femur after shape based interpolation of the stack of femur contours along the slices.
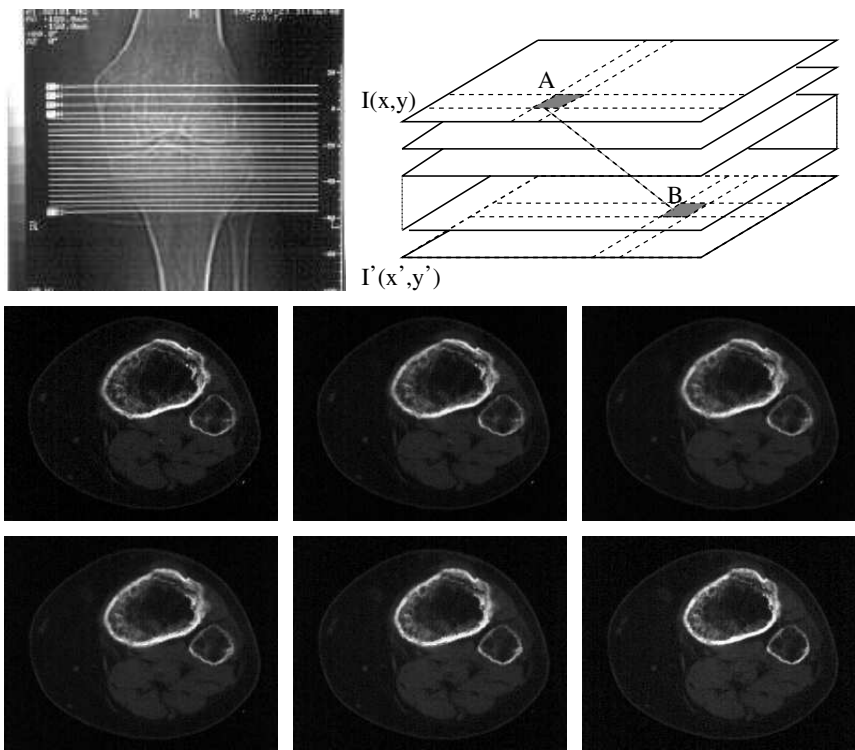
Fig. 1.   Image based interpolation: (Upside) slices location and the correspondence between two feature points on two slices between which several slices are interpolated to reach isotropic resolution. (Bottom) Original slice-i, four intermediate interpolated slices and original slice-$i-1$.
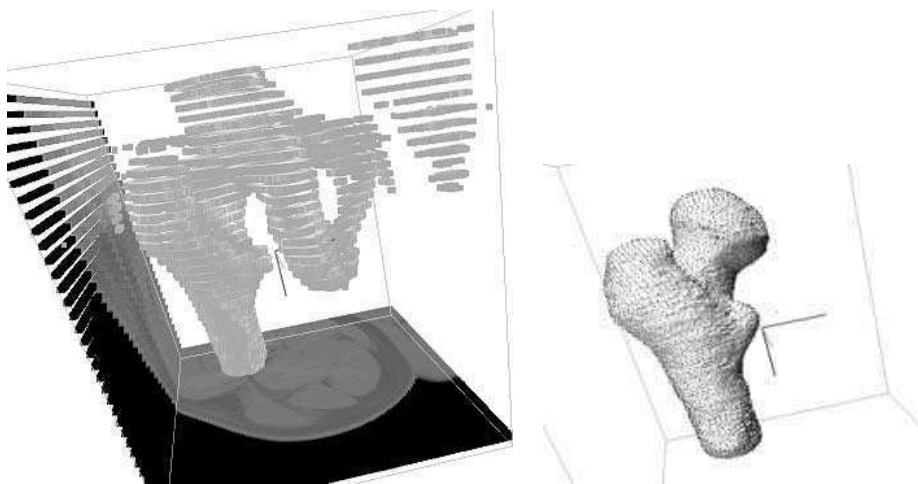

Fig. 2.   Shape based interpolation: (Upside) segmentation. (Bottom) Interpolation between contours.

1.2. *Segmentation and tracking*

Three-Dimensional segmentation of anatomical organs is widely used for diagnosis of diseases and planning of surgical interventions. However, the segmentation of medical images is a difficult task due to the complexity and variability of the anatomic shapes, sampling artifacts and noise.

Despite the capabilities of the different medical imaging modalities, it is difficult to develop automated or semiautomated methods to extract object surfaces from image data. Edge-based methods allow delineation of regions through the location of their contours, typically as points of high gradient. However, this approach has two major difficulties. First, the threshold value to achieve effective edge detection varies from slice to slice, leading to significant errors. Second, it has a high sensibility to noise, particularly between regions with small contrast, and needs gap closing methods to link edges belonging to the same boundary.

Image segmentation by thresholding provides binary images composed by regions that identify objects segregated from the background. Greyscale thresholding works well when images have uniform regions and contrasting background. Its major drawback is the lack of local adaptability, voxels of the same tissue may exhibit different intensities between slices and even within the same slice. It is generally necessary to postprocess the image using mathematical morphology,[2] and/or combining thresholding with region growing.[70]

Region-based methods yield parts of the image which meet the requirements of a given uniformity criterion. One of the most widely used techniques in this category is region-growing. In its simplest form, the method starts with a small region and then examines its neighbourhood in order to decide whether or not they have similar features. If they do, then they are grouped together to form a new bigger region. It has two drawbacks: first, post-processing is usually required to cope with infra- and over-segmentation, generally applying domain knowledge;[13,54] and second, it also suffers from bad location of boundaries.

Regardless what segmentation method is used, automated methods often cannot generate perfect object boundaries. The difficulty of segmentation is an aspect of global/local duality. Consideration of local features can lead to over-segmentation, and erroneous fusions (infra-segmentation) can be obtained using only global features. To tackle the problem of obtaining a tradeoff between infra/over-segmentation, (local) edge information and (global) homogeneity information must be combined during image segmentation.[9,38,45,68] Moreover, a meaningful segmentation must take into account high level descriptions of the objects.[8]

Region-based segmentation can be improved by introducing knowledge about the domain. As a matter of example, we describe a rule-based approach for region growing whose scheme is sketched in Fig. 3.[41] The system was developed for the segmentation of CT image sequences and has two main blocks: low level (extracts low level features) and high level (performs region growing by using domain knowledge). The system has also two memories: data base (input, output, intermediate
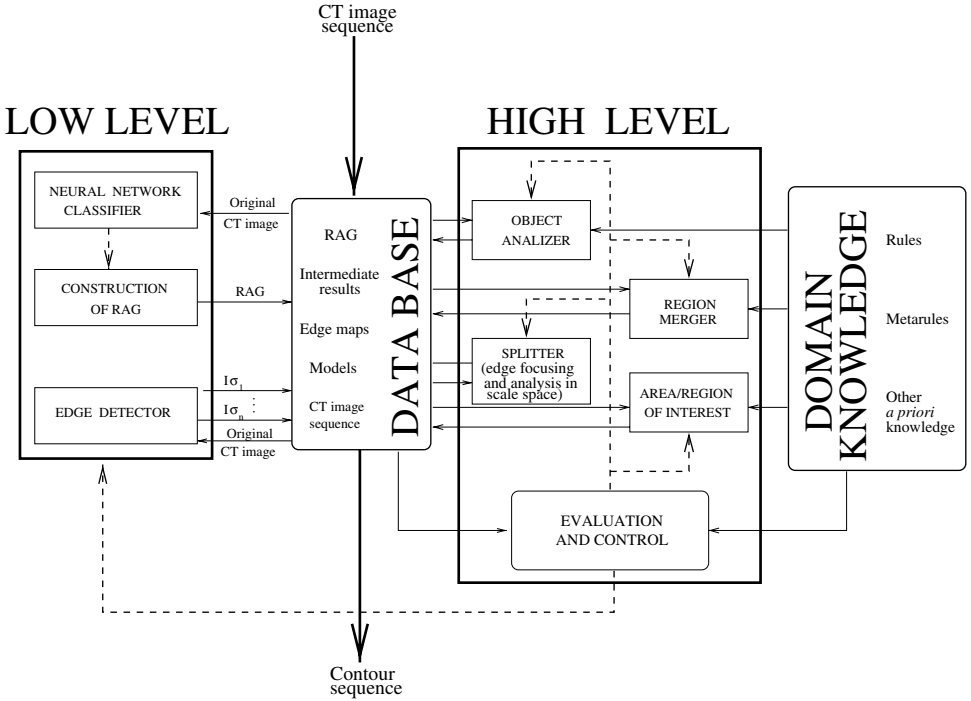
Fig. 3.   Global scheme of a knowledge-based region growing system.

results, etc.), knowledge domain (rules, strategies, etc.). The system begins to carry out a low level presegmentation by means of a neural network classifier. Then, regions and their features are organized in a region adjacency graph (RAG), where the nodes contain region features and the weighted arcs represent the size of the shared contours. After the presegmentation, the high level stage works directly over the RAG. It begins with the selection of growing kernels of bone areas and completes the recognition of the bone structures by applying rules for the merging and the splitting of regions. The growing process is driven by a set of fuzzy rules that embodies criteria as compactness, closure, spatial closeness and distance to a model. Given the 3D continuity of the bone structure, shape and location of bone in each slice are similar to those in the previous slice. So, the bone contour delineated in the previous slice is considered as a model for the current slice. The process starts by defining an area of interest (AOI) by means of the dilation of the model mask. Regions whose centroid are out of this area are discarded. From the remaining regions, the most dense ones are chosen as the growing kernels (the higher the intensity, the higher the probability of belonging to the bone). All these tasks are carried out by the block *Area/Region Of Interest*. The bone growing is performed by the combined action of two blocks *Region Merger* and *Object Analizer*, which decide on the merging of regions to bone kernels, and the

incorporation of new bone kernels respectively. The process iterates until a stop criterion, derived from the measurement of matching between objects and models, is achieved in the *Evaluation And Control* block. If different structures appear overlaid, then a splitting process, guided by edge information (low level), has to be carried out. Edge information is obtained combining the outputs of an edge operator with several scales $(I_{\sigma_1}, \ldots, I_{\sigma_n})$ and taking into account the shape and location of the model (constraints imposed by the knowledge on the domain). The system has a special sub-block for edge focusing and scale space analysis.

Deformable models approaches provide a simplest and more effective alternative to rule-based systems. Deformable models stand out from other segmentation techniques due to their ability to interpret sparse features and link them to obtain object contours. An exhaustive summary of the application of deformable models to medical images can be found in the work by McInerney and Terzopoulos.[31]

Deformable models have been developed for segmentation and tracking. In this paper two active contour approaches will be reviewed and applications will include automatic contouring of organs and tracking of organ motion. The first active contour technique is devoted to tackle segmentation tasks where the desired features of the object of interest could slightly change from slice to slice, while the second active contour technique is dedicated to tracking tasks where the features of the object of interest does not change from frame to frame, but the shape does. The segmentation approach aimed to learn and update the best discriminants from one slice to the next. The tracking approach, however, assumes well established goal features and intends to reach them as fast as possible.

2. Segmentation by Active Contours

A deformable model is an energy minimization method, where the energy functional is defined in terms of intrinsic shape attributes (internal energy) and desired image features (external potential).[23] The external potential originates forces that attract the model to specific image features while the internal energy causes stress forces that try to maintain model continuity and smoothness. When a minimum is reached the geometric deformation finishes. Therefore, the model evolves from its initial location to approach the nearest energy minimum, which maybe does not correspond to the desired object surface.

Although the classical edge-oriented deformable contour models have demonstrated high efficiency in the segmentation of biomedical structures, they are not free of limitations. Since the classical snakes are based on gradient information, they share some limitations of edge detectors. Ronfard[48] proposed a new deformable model which uses no edge information, but local computations around contour neighbourhoods (region information). Statistical models of object and background regions are used to push the model towards border points. Region information can provide clues to guide the evolution where gradient information vanishes (homogeneous areas). This method was proposed for the case of step-edges but cannot

distinguish boundaries between objects of similar features. Ivins and Porril proposed a related approach where a snake is driven by a pressure force that is a function of the statistical characteristics of image data.[21] The model expands until it finds pixels that lie outside user-defined limits relative to the statistical characteristics of a seed region; when these limits are crossed the pressure force is reversed to make the model contract.

Region features present the advantage of being less sensitive to noise. Their inconveniences reside in that the change of the pixel characteristics that suggest the presence of a contour remains diluted in the global character of the region features. This frequently provokes the undesired merging of regions. This problem can be alleviated with gradient/edge information. Chakraborty *et al.*[6] proposed a deformable boundary finding approach that integrates region and edge information for the segmentation of homogeneous structures surrounded by a single background. This approach must know *a priori* which regions belong to the object of interest. This idea was also extended to 3D.[56]

A unifying framework which generalizes the deformable model, region growing, and prior matching approaches was proposed by Zhu *et al.*[71] The method combines local and global information, and local and global optimization. It can be applied both for obtaining global segmentation of images and finding individual regions.

Several researchers augmented snake-like models with prior information about typical mean shapes and normal variations.[12,55] A number of researchers have incorporated knowledge of object shape using deformable shape templates. We can find in the literature statistical-based approaches, which preferably attract the curve towards image features consistent with their trained shapes[4,6,12,34,55,67] and/or grey level appearance. Turk and Pentland[59] have used principal component analysis to segment face images in terms of a set of basis functions. Paragios and Deriche[39] have proposed geodesic active contours based on a supervised statistical modeling, which combine *a priori* knowledge about the desired boundary, and the exterior and interior region properties. Yezzi *et al.*[65] have presented active contours with a natural use of global and local information to segment regions distinguishable by a given set of statistics. Some authors consider the full appearance (shape and grey level) of the object.[25,26]

The model-based segmentation could be stated as a two (no necessarily disjoint) phases: registration and free form deformation. On one hand, registration describes a transformation with far less degrees of freedom than the free form deformations. Therefore, their ability to represent shape variations is less important than the free form deformation. On the other hand, because of their restricted degrees of freedom, they tend to be more robust than free form deformations. So, the first is better for describing global shape and location, and the second is better in detecting fine details.

The integration of the selection of features, and their scales, in a statistical framework is the most convenient approach to relate the segmentation technique to the specific segmentation task. Our approach combines the selection of the best

scale, for the fine local fitting, with a multiscale approach to cope with cases where there is a big change in contour location from one slice to the next slice.

When deformable models are applied to 3D medical data, two kinds of segmentation are possible: $2\frac{1}{2}D$ (slice by slice) and 3D. Initialization is simpler for $2\frac{1}{2}D$ segmentation, since it is applied just to the first slice. Each slice is initialized with the result of the previous one. When processing slice sequences, the contour in the previous slice can be used as the initial deformation state in the present slice. The construction of the object surface is usually carried out by stacking the contours obtained from the 2D segmentation of each CT slice and connecting adjacent contours using polygons.

Manual edition of initial 3D surfaces is very laborious, and the automatic or semiautomatic initialization is usually more complex than the 2D counterpart. In deformable model literature we can find different approaches to cope with the initialization problem in 3D. Some approaches are based on the manual or semiautomatic selection of anchor points. Among them is the imposing of interactive constraints in the form of springs and volcanoes,[23] and the method of Neuenschwander *et al.*,[35] that allows the user to fix a set of seed points and their normal vectors, which cannot be changed during deformation.

Some authors propose multiple initialization through several seed models. During the deformation process, several initial models will merge and the superfluous ones should be removed.[58] Important decisions have to do with: the number and location of initial seed models, the stopping criteria of the seed growing process, and choosing one result among the final models.

Fully automatic initialization can be achieved by matching the object in the image with a prototype, as done by Bajcsy and Kovacic[3] who used brain atlases in their specific-purpose initialization techniques. There are several problems with the deformable atlas approach: the technique is sensitive to initial positioning of the atlas, and the presence of neighbouring features may also cause matching problems. One solution is to use image preprocessing in conjunction with the deformable atlas.[51]

Moreover, the length of the objects to be segmented can vary significantly, so the use of 3D *a priori* models can be very complex.

As a counterweight to the low robustness due to the usual low accuracy in initialization, deformable surface models have the power to ensure smoothness and coherence in 3D shapes. However, where 3D image data sets have no isotropic spatial resolution, features cannot be extracted in 3D, and therefore a slice-to-slice framework is more adequate.

2.1. *Enhancement of desired feature-points*

The success of any active contour is based on the discrimination and spatial continuity of the image features used to guide the curve. Obtaining correct result of the segmentation is intricately tied to the type of the sought features and the criteria used for discriminating between the extracted features.[22]

Usually, image features are not selected in accordance with the specific organ and image modality; thus too many feature points that do not belong to the organ of interest are enhanced, while other important ones go unnoticed. Several authors proposed the combination of different image features depending on the application domain and the objects to be segmented.[4,67] However, the best way of selecting and integrating different features remains an open problem. The feature space must be capable of representing all image features of interest. To achieve such a broad description, we propose employing a bank of Gaussian derivative filters. Numerous physiological measurements support the theory that receptive field profiles can be modelled by Gaussian filters of various widths and their partial derivatives.[17,24,66] From a mathematical point of view, it has been proven that the complete family of scale-space Gaussian kernels defines a sufficient Cartesian family for complete determination of local image structure.[16]

A very important issue concerns the construction of an irreducible set of differential invariants in a fixed N-jet (the set of partial derivatives up to Nth order). This irreducible set should provide optimal tools to measure the amount of *featureness* at a certain position for any generic image configuration.[36] We consider the problem of constructing an optimal subspace of the complete feature space, where a certain part of the object boundary better discriminates from others. For example, it is well-known that to discriminate texture patterns only second Gaussian derivatives are needed.[29]

Another important issue which must be addressed is the selection of filter scales. In general, an optimal single scale does not exist. One alternative lies in searching for the best local scale, but it depends on the imaging modality and the specific image feature.[15,57] Multiscale analysis is usually used when there exists no clear way to choose the right scale. The decisions of how many scales should be used and how to combine them are the main issues of this approach.[42]

Knowledge propagation is essential in expanding single slice processing into volume-based efforts, both in time and accuracy. Essentially, knowledge gained from processing previous slices, can be used as a guideline for decisions in next slices.

We represent the snake curve by means of a cubic B-spline, and therefore each snaxel (snake control point) influences four curve segments. For each snaxel, we define a *contour patch* as the central part of its four curve segments of influence, avoiding the overlap between adjacent contour patches. Each snaxel moves in accordance with the energy accumulated along its contour patch.

The statistical discriminant snake (STD-snake) has a triple objective: to obtain a general representation of any contour class in terms of a set of features, to determine the optimum discriminant for each desired contour part and to deform towards the target object boundary.

To make the snake discriminant, each snaxel must be able to distinguish between its corresponding contour target and other structures. The discriminant snake should learn the more relevant features for each specific segmentation problem,
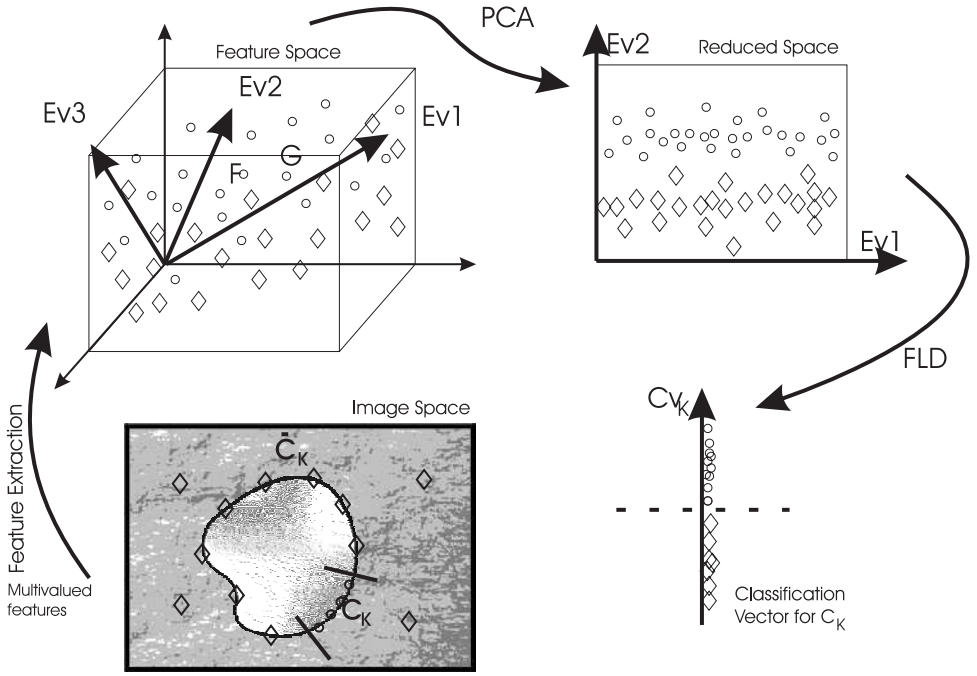
Fig. 4. From feature space to classification vector for each contour part.

and generalize the classical approach that only uses gradient information. For each desired contour part, a set of samples is obtained and then a supervised learning aimed at the maximization of the between-class scattering of image features, is accomplished. We define a parametric potential that depends on the features of each part of the target boundary. By parametric we mean that each snaxel sees the image features in a different way, because each one is aimed to fit different object boundary parts.

Figure 4 illustrates the learning process carried out for each contour part. The features of the target contour part and the complementary class are extracted from the training contour and background. Then, Fisher linear discriminant analysis (FLDA) gives the optimal discriminant. During the optimization process each snaxel will evolve guided by its discriminant. The following sections describe in detail these steps.

2.2. *Image features learning*

We use a bank of Gaussian filters to extract the image features. Since the directional derivative operator is steerable, each filter $G^d(x, y, \sigma, \phi)$ of degree d and orientation ϕ can be synthesized by interpolating $d + 1$ basis functions $\left\{G^d(x, y, \sigma, \theta_k)\right\}_{k=0}^{d}$.[18] We define $\mathcal{G}$ as the $\mathcal{D}$-jet filter bank with Σ scales: $\sigma \in$

$\{2^0, 2^1, \ldots, 2^{\Sigma-1}\}$. The size of the bank of filters is a function of $\mathcal{D}$ and Σ:

$$\mathcal{N} = \dim(\mathcal{G}) = \Sigma \sum_{d=0}^{\mathcal{D}} (d+1) = \frac{\Sigma}{2}(\mathcal{D}^2 + 3\mathcal{D} + 2). \tag{1}$$

Our bank of filters contains derivatives up to degree three, because variance of higher-order filters can be expected to approach that of image noise, and tend to be highly correlated to the outputs of lower order filters.[46] The impulse responses of the Gaussian filters are given by:

$$G^0(x, y, \sigma) = \frac{1}{2\pi\sigma^2} \exp\left\{-\frac{1}{2}\left[\frac{x^2 + y^2}{\sigma^2}\right]\right\},$$

$$G^1(x, y, \sigma, \theta_1) = R\left(\left(-\frac{x}{\sigma^2}\right) G^0(x, y, \sigma)\right)^{\theta_1},$$

$$G^2(x, y, \sigma, \theta_2) = R\left(\left(-\frac{1}{\sigma^2} + \frac{x^2}{\sigma^4}\right) G^0(x, y, \sigma)\right)^{\theta_2},$$

$$G^3(x, y, \sigma, \theta_3) = R\left(\left(\frac{3x}{\sigma^4} - \frac{x^3}{\sigma^6}\right) G^0(x, y, \sigma)\right)^{\theta_3},$$

where the superscript in $G^d(\cdot)$ represents the derivative degree, and $R(\cdot)^{\theta_d}$ represents a rotation of an angle $\theta_d \in \{\frac{k\pi}{d+1}|k = 0, 1, \ldots, d\}$.

The convolution of $\mathcal{G}$ with the image $\mathbf{I}$ allows the generalized description of edge, ridge, valley or whatever image intensity features. We define a multivalued feature $\mathcal{P}$ as the result of this convolution:

$$\mathcal{P} : \mathbf{R}^2 \to \mathbf{R}^{\mathcal{N}}$$

$$\mathbf{I} \to \mathcal{G} * \mathbf{I}.$$

The image potential will refer to what each snaxel interprets on this multivalued feature.

Each snaxel will give different relevance to each component of the feature vector $\mathcal{P}$. In that sense, we shall talk about variable (parametric or locally defined) external energy. To determine the relative importance of each component of the feature vector, a supervised learning is performed. The algorithm starts with a set of N sample feature vectors on object contour and non object contour $\{s_1, s_2, \ldots, s_N\}$, $s_i \in \mathbf{R}^{\mathcal{N}}$. For each contour patch k, each multivalued feature s_j is assigned to one of two classes $\{C_k, \bar{C}_k\}$, contour k and the complementary class respectively.

Then a FLDA provides the optimal linear discriminant for each patch k that maps the original $\mathcal{N}$-dimensional space into a 1-dimensional space. The FLDA is a recognized method to analyze the separation ability between classes $\{C_k, \bar{C}_k\}$ in the feature space and it can demonstrate the separation ability from the original feature data directly if the underlying distributions of the observations are multivariate

normal.[14] The Fisher discriminant function is a scalar function:

$$F(\mathcal{P}) = W^T \cdot \mathcal{P} + \alpha,$$

where $\mathcal{P} \in \mathbf{R}^{\mathcal{N}}$ is the multivalued feature, $W \in \mathbf{R}^{\mathcal{N}}$ is a vector that projects the features of each pixel into the reduced feature space $\mathbf{R}$, and α is a constant that determines the border between classes.

The optimal projection W_{C_k} for each contour patch C_k is defined as the one which maximizes the ratio of the between-class to the within-class scatter of the projected samples:

$$W_{C_k} = \arg \max_W \frac{|W^T S_{B_k} W|}{|W^T S_{W_k} W|}$$

where

$$S_{B_k} = N_{C_k}(\mu_{C_k} - \mu)(\mu_{C_k} - \mu)^T + N_{\bar{C}_k}(\mu_{\bar{C}_k} - \mu)(\mu_{\bar{C}_k} - \mu)^T$$
$$S_{W_k} = \sum_{s_j \in C_K} (\mu_{C_k} - s_j)(\mu_{C_k} - s_j)^T + \sum_{s_l \in \bar{C}_k} (\mu_{\bar{C}_k} - s_l)(\mu_{\bar{C}_k} - s_l)^T$$

are the between-class S_{B_k} and within-class S_{W_k} scatter matrices; being N_{C_k} and $N_{\bar{C}_k}$ the number of samples in class C_k and $\bar{C}_k$ respectively; μ_{C_k} and $\mu_{\bar{C}_k}$ are their mean feature vectors, and μ is the mean vector over all samples. W_{C_k} is the S_{W_k}-generalized eigenvector of S_{B_k} with the largest corresponding eigenvalue.

If S_{W_k} is nonsingular we can obtain a conventional eigenvalue problem by writing

$$W_{C_k} = S_{\mathrm{Wk}}^{-1}(\mu_{C_k} - \mu_{\bar{C}_k}).$$

Before computing W_{C_k}, it is necessary to prevent the existence of a singular S_{Wk}. To this purpose, principal component analysis (PCA) is firstly performed on all training samples. FLDA was proven to be equivalent to PCA plus FLDA for small sample size problems.[64] In these problems, the number of training samples is less than the dimension of the feature vector, thereby the within-class scatter matrix is singular. As a result, we obtain W_{pca} as a matrix of dimension $d_{\mathcal{G}} \times m$ formed by the m most important eigenvectors of the scatter matrix S_T of the full set of samples. PCA is then followed by FLDA, to obtain the best discriminant, $\mathcal{V}_{C_k}$, for each contour part C_k:

$$W_{C_k}^T = W_{\mathrm{k\_flda}}^T W_{\mathrm{pca}}^T$$
$$W_{\mathrm{k\_flda}} = \arg \max_W \frac{|W^T W_{\mathrm{pca}}^T S_{B_k} W_{\mathrm{pca}} W|}{|W^T W_{\mathrm{pca}}^T S_{W_k} W_{\mathrm{pca}} W|}.$$

Finally, we can write the optimal projection as:

$$W_{C_k} = W_{\mathrm{pca}} \left(W_{\mathrm{pca}}^T S_{W_k} W_{\mathrm{pca}}\right)^{-1} W_{\mathrm{pca}}^T(\mu_{C_k} - \mu_{\bar{C}_k}), \tag{2}$$

where W_{C_k} maximizes the ratio between-class/within-class discrimination, and μ_{C_k} is the mean feature vector of the patch k.

The discriminant for each patch k consists of the classifier vector W_{C_k} and the class centre α_{C_k} of the samples of patch k in the feature space:

$$\alpha_{C_k} = W_{C_k}^T \cdot \mu_{C_k}.$$

We define a component of the local external energy of the snake in terms of the distance D_{C_k} from the image features $\mathcal{P}$ in the current location of the snake to the desired contour features:

$$D_{C_k} = \left(\alpha_{C_k} - W_{C_k}^T \mathcal{P}\right)^a. \tag{3}$$

In a previous paper[43] we have used a value $a = 2$ which defines a squared Euclidean distance, however it is more interesting to define an odd function because the sign of the score $\alpha_{C_k} - W_{C_k}^T \mathcal{P}$ is important: high positive value classifies a data into the complementary class, however a high negative value classifies a data into the same class, but no exactly the learned feature, and a value around zero represents a good matching. Therefore, it is convenient to penalize more strongly high positive scores than the symmetric negative ones. The function must have the minimum in 0, and this behavior can be achieved with values in $(1.5 + n, 2.5 + n)$ with n integer. Figure 5 represents the behaviour of the function $f(x) = x^a$ parameterized by a. In our experiments we use the value $a = 1.6$.

We give an explicit expression for the external energy and forces in Sec. 2.4.

Each snaxel j learns different contour types (W_{C_j}, α_{C_j}) in slice $i - 1$, and therefore each snaxel interprets in different way the image features $\mathcal{P}$ in slice-i. This behaviour of our method is illustrated in the synthetic example of Fig. 6. Figure 6(a) contains a synthetic image of two circles on a white background. Figures 6(b), (c) correspond to the initial contour and final fitting of the snake guided by the classifiers (one per contour patch) in the multivalued feature space. Figs. 6(d), (e) contain the components of the multivalued feature for two scales and derivatives up to degree three. Figure 6(f) shows the image potentials as are seen by each snaxel.

In a slice to slice segmentation it can be assumed that corresponding features in two adjacent slices will be close in the cross-sectional plane. So, we add to the expression of the local external energy the knowledge about the distance from the location of the centre of the learned contour, as follows:

$$E_{C_k} = D_{C_k} + b(1 - e^{-d_k^2}), \tag{4}$$

where b is a normalization constant and d_k^2 is a measure of the distance from an image point to the location of the centre of the learned contour patch.

A real example in the field of the medical images is showed in Fig. 7, where alternating colours on the training contour represent different contour patches. Figures 7(b–e) represent the distances (in feature space) to contour classes placed on the west, east, north and south sides of the training contour. As is desirable, the minimum of each distance map is around each learning contour class, and there is a maximum in the diametrically opposed part of the contour.
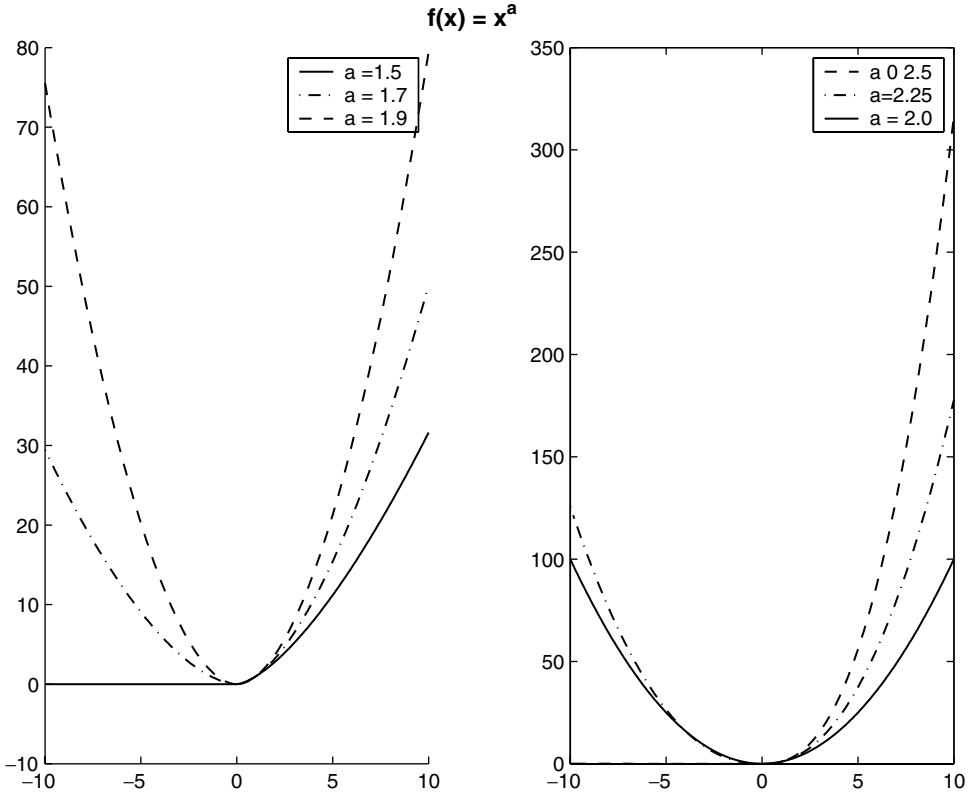
Fig. 5.   Distance measure as a function of parameter a.

2.3. *Continuity and regularization of classification vectors*

In general, features smoothly vary along the boundary of an object in a medical image. If the contour is densely sampled, smooth variation of the classifiers can be expected to occur along the learning contour. Figure 8 illustrates this hypothesis. It contains the representation of components of the classification vector along the internal contour of a coronary vessel in an intravascular ultrasound (IVUS) image. As can be seen, the features and the components of the classification vectors (W_{ij}) vary smoothly along the sequence of contour patches.

If the inter-slice resolution is close to the intra-slice resolution, the similarity between adjacent slices will be high, which will allow us to learn the contour features in a slice and use it to segment an adjacent slice. Figure 9 shows the continuity of the classifiers between adjacent slices in a CT image sequence of femur.

We assume that features vary smoothly along the object surface (inter- and intra-slices) to improve the robustness of the learning process. Due to possible mislocation of snaxels at the beginning of the learning step, some spurious contour configurations could be learned. We introduce a regularizing process which increases
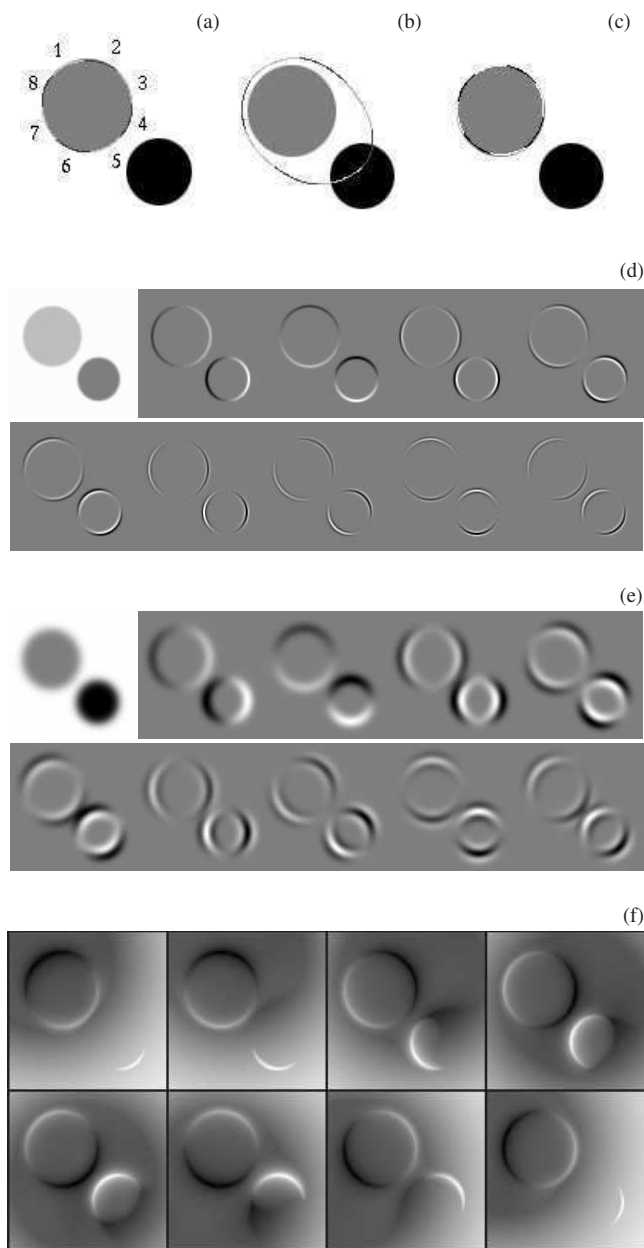
Fig. 6.   Training contour and patch numbering (a); initial contour (b); final contour (c); filter responses for scales $\sigma = 1, 4$ (d,e); distance map to the different contour classes of the top left circle, top left map correspond to patch 1 and bottom right map correspond to patch 8 (f).
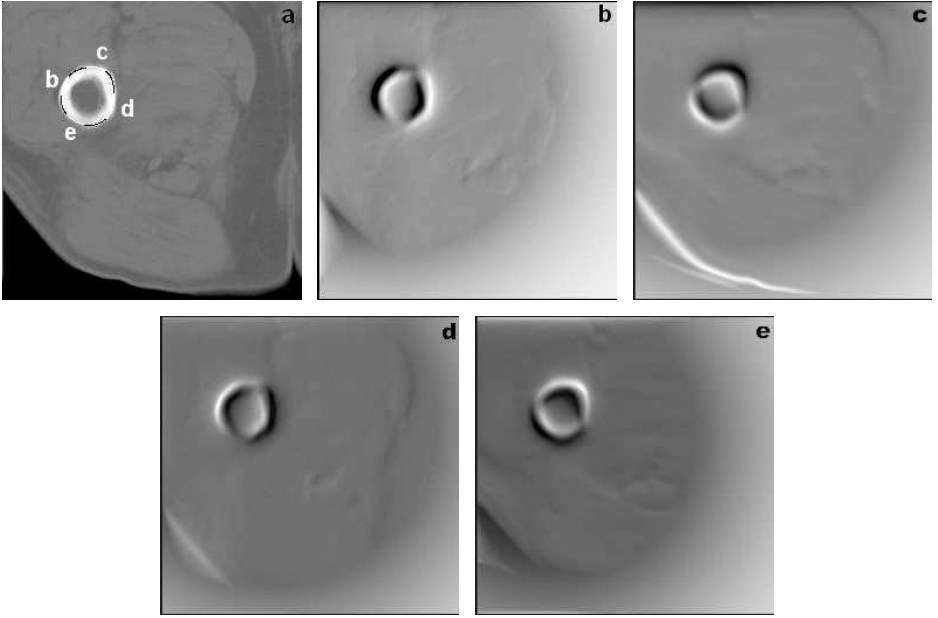
Fig. 7. Selected patches (b, c, d, e) on the training contour (a); and distance maps generated by the discriminants of each learned contour patch (b–e).

the continuity of the classification vector, and makes the learning step more robust. In this process, the classifier of each contour patch p is regularized taking into account its direct neighbour classifiers:

$$W'_{C_p} = \zeta_p W_{C_p} + \frac{1 - \zeta_p}{2}(W_{C_{p-1}} + W_{C_{p+1}}), \qquad (5)$$

$$\alpha'_{C_p} = \zeta_p \alpha_{C_p} + \frac{1 - \zeta_p}{2}(\alpha_{C_{p-1}} + \alpha_{C_{p+1}}),$$

where the superscripts indicate new values, and $\zeta_p \in [0, 1]$ is a regularizing parameter. From the regularized classifiers, we define a continuous classification function as follows:

$$W(s) = \sum_{p=0}^{N-1} W'_{C_p} \varphi_p(s)$$

where φ is considered in terms of finite-element nodal shape functions. For linear functions:

$$\varphi_p(s) = \begin{cases} |u(s) - u_p|, & \text{if} \quad u \in [u_{p-1}, u_{p+1}] \\ 0, & \text{otherwise} \end{cases},$$
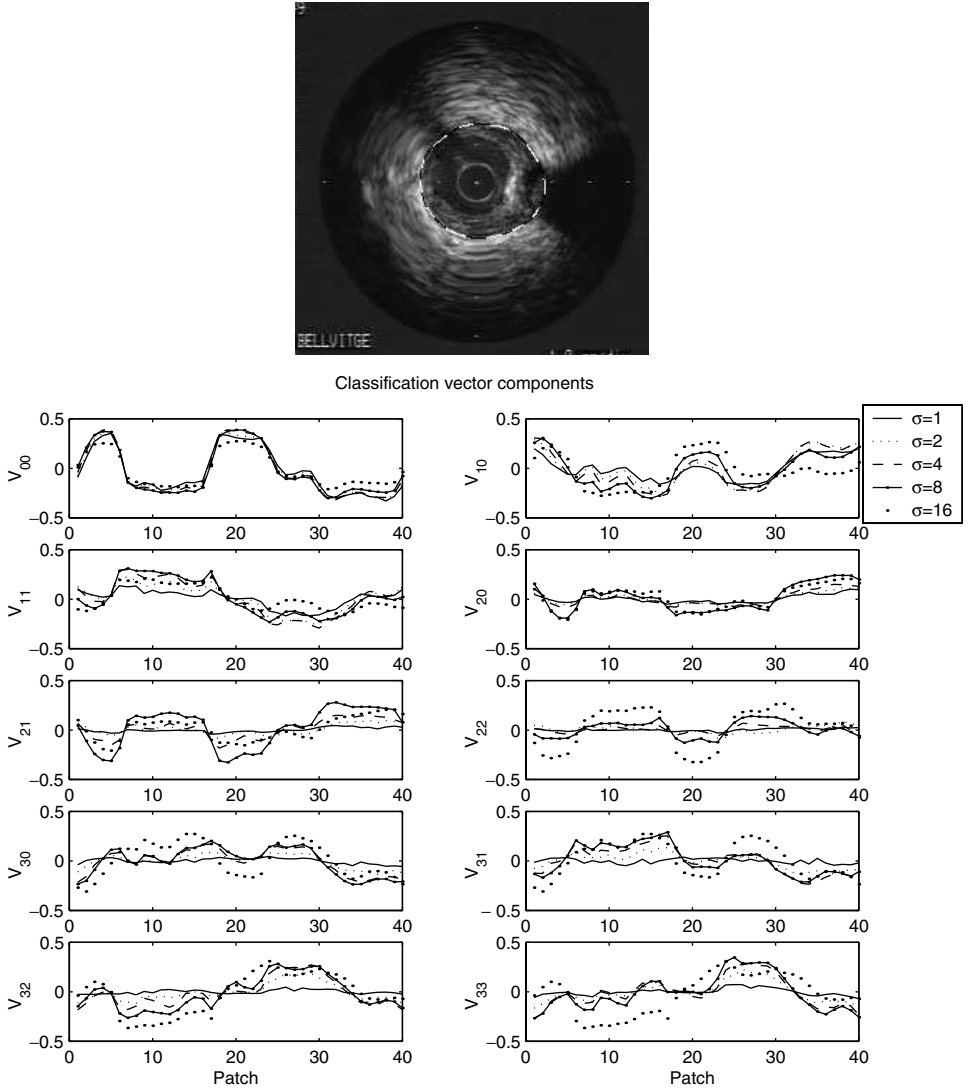
Fig. 8. Components of the classifier along the patches of the learning contour. In W_{ij}, $i \in [0, 1, 2, 3]$ indexes derivative degree and j indexes orientation ($\theta_i \in \{\frac{j\pi}{i+1} \mid j = 0, 1, \ldots, i\}$), and each line corresponds to a different scale.

where $u(s)$ represents the contour curve and u_p are the nodal points assuming uniform knots ($|u_j - u_{j+1}| = 1$) of the same curve.

To correct erroneous vectors without seriously affecting their neighbours, we compute parameter ζ_p as follows:

$$\zeta_p = \max\{W_{C_p}^T W_{C_{p-1}}, W_{C_p}^T W_{C_{p+1}}\}. \tag{6}$$
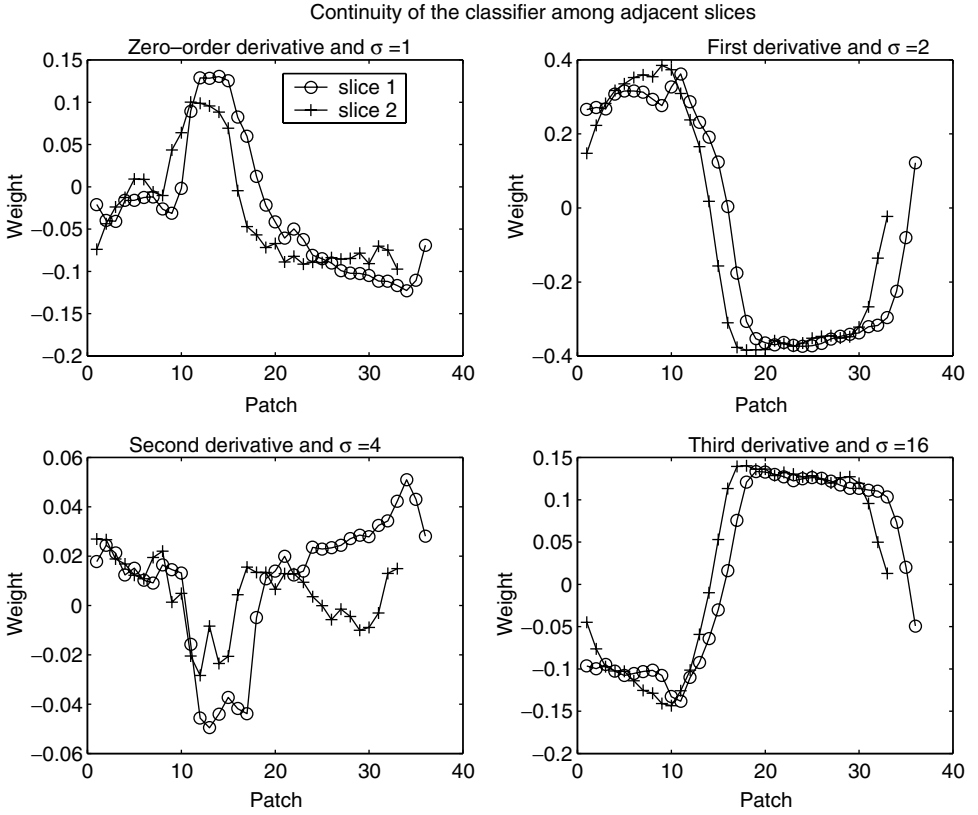
Fig. 9.   Continuity of the classifiers among adjacent slices for different derivatives and scales.

Provided that classifiers are normalized, ζ_p is the cosine of the angle between both vectors:

$$W_{C_p}^T \cdot W_{C_{p\pm1}} = |W_{C_p}^T| \cdot |W_{C_{p\pm1}}^T| \cos \angle_{W_{C_p}, W_{C_{p\pm1}}}$$

$$|W_{C_p}^T| = |W_{C_{p\pm1}}^T| = 1 \Rightarrow W_{C_p}^T \cdot W_{C_{p\pm1}} = \cos \angle_{W_{C_p}, W_{C_{p\pm1}}}$$

where $\angle_{.,.}$ denotes the angle between vectors. If ζ_p takes values close to 1 it means that classifier W_{C_p} has at least one similar neighbour classifier, so it is not greatly smoothed. This prevents that contours with real sharp changes smooth severely, and only erroneous vectors are significantly corrected.

In Fig. 10 one can see the 0, 1st and 3rd derivatives with respect to the coordinate x in the classifiers of snake patches ($p = 0, \ldots, 45$). Note that different weights form a continuous function of the internal parameter of the snake curve. In patches 35 and 37 smoothing of the weights is observed after applying Eqs. (5) and (6).
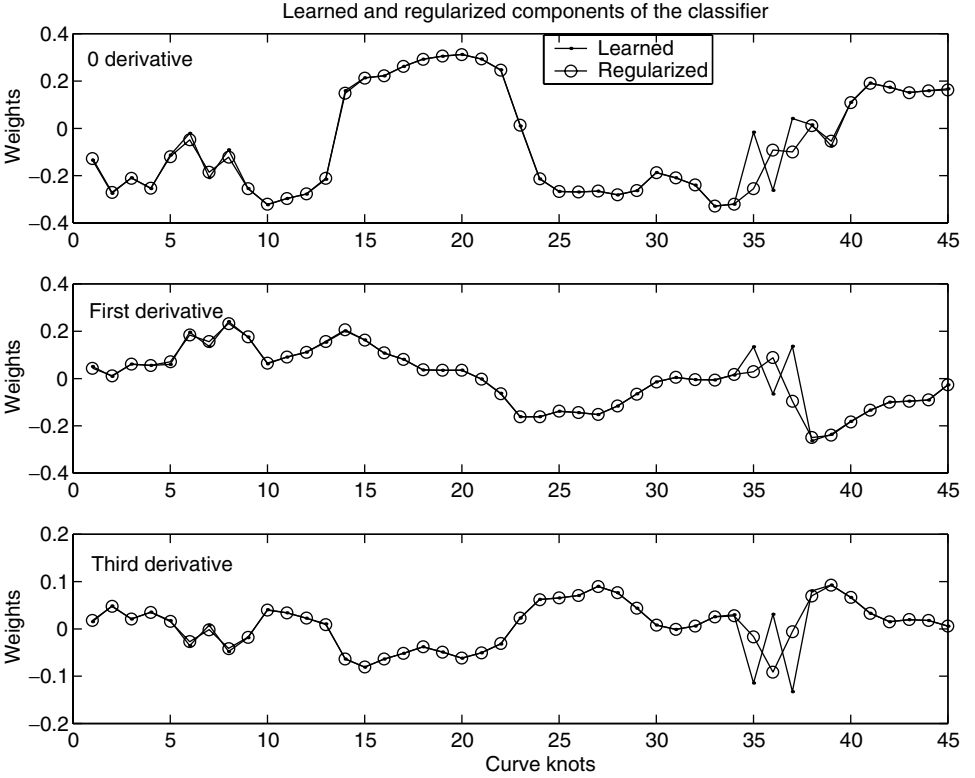
Fig. 10.   Learned and regularized components of the parametric classifier.

2.4. *STD-Snake implementation issues*

Let $C(s) = \sum_i V_i B_i(s)$ be a B-spline representation of the curve, where V_i are the control points and s is the curve parameter. The total energy can be written[32]:

$$
\begin{aligned}
E_{\text{total}}(C) &= \int_0^1 \left(E_{\text{int}}(C(s)) + E_{\text{ext}}(C(s)) \right) \mathrm{d}s \\
&= \int_0^1 \left(\omega_1 \left| \frac{\partial C(s)}{\partial s} \right|^2 + \omega_2 \left| \frac{\partial^2 C(s)}{\partial s^2} \right|^2 + E_{\text{ext}}(C(s)) \right) \mathrm{d}s \\
&= \sum_j \left[\omega_1 \left(\sum_i V_i \frac{\partial B_i(s_j)}{\partial s} \right)^2 + \omega_2 \left(\sum_i V_i \frac{\partial^2 B_i(s_j)}{\partial s^2} \right)^2 + E_{\text{ext}}(C(s_j)) \right].
\end{aligned}
$$

We define the external energy as a function of image features and the training contour classes. For a point in the patch k, the external energy is:

$$
E_{\text{ext}}(C(s_j)) = \left(\alpha_{C_k} - W_{C_k}^T \mathcal{P}(C(s_j)) \right)^a + b \left(1 - e^{-d_k^2} \right). \tag{7}
$$

We are looking for control points V_i, $i = 0, \ldots, N$ that minimize the total energy, i.e. that satisfy:

$$\frac{\partial E_{\text{total}}}{\partial V_l} = 0, \quad \forall l \in \{0, \ldots, N\}.$$

These equations can be written in a matrix form similar to the classical snake:

$$A_b V + G = 0, \tag{8}$$

where the stiffness matrix A_b for B-spline snakes is still a banded matrix and G plays a role equivalent to forces in the classical snake. The i-th element of G is:

$$G_i = \sum_j B_i(s_j) \nabla \left(\left(\alpha_{C_k} - W_{C_k}^T \mathcal{P}(C(s_j)) \right)^a + b(1 - e^{-d_k^2}) \right). \tag{9}$$

The solution of the segmentation problem is found in an iterative and incremental way:

$$V_{t+1} = (A_b + \gamma I)^{-1} (\gamma V_t + G(C_t)),$$

where γ is the damping parameter that determines the convergence rate and is equal to 1 in our implementation. The method is incremental because it begins with the highest scale components of the classification vector, incorporates a new scale after convergence, and finishes with all the scales. The weights corresponding to the scales that are not being considered are set to zero. In each incremental step the algorithm iterates until all the control points stop evolution or oscillate.

As the number of control points increases the delineations of convoluted shapes improve. However, when the number of control points is increased the size of the patches decreases and discrimination power decreases. We found that a distance of about 5 pixels between adjacent control points worked well in all tested examples. To maintain these distances we augmented the algorithm with a reparameterization capability during the deformation process.

3. Evaluation of the Segmentation Approach

In order to validate the method we have tested STD-snake on organs from different imaging modalities. First, we are going to show some properties of our approach with respect to classical implementations, and afterwards we will show a comparison to expert tracings.

3.1. *Comparison to classical snakes*

We are going to compare the method with classical snakes to both prove the advantages of our method and to justify the choices made. We will begin by showing the advantages of using multiple scales and a continuous classifier, and will finish by demonstrating the ability of our method to describe and learning different image features.

3.1.1. *Integration of multiple scales*

The use of multiple scales has two objectives: first, the greater the scale, the wider the spatial effect of the potential. Hence high scales allow large displacement of the snake contour to reach the desired object. Second, multiple scales greatly reduce the number of local minima.

Figure 11 shows an example that illustrates this property of the multiscale vector. We have computed the distance map generated by the classification vectors over all the image. Figure 11(b) shows the computed distance map using only the lowest scale in the classification vector, and Fig. 11(c) contains the distance map computed by the multiscale classification vector. As can be seen in Fig. 11(d), (e) the approximate gaussian shape of the histogram has lesser variance and higher mean distance in Fig. 11(e) than in Fig. 11(d). This means that the number of convex components in the distance map diminishes and there are less potential minima in Fig. 11(c) (multiscale) than in Fig. 11(b) (single scale). In all the examples shown in this work, we used scale parameters with values $\sigma = 1, 2, 4, 8$.

3.1.2. *Continuous classifier*

The definition of a classifier that continuously varies along the contour has several advantages. One of them is the capability of filtering noise in the learned classifiers without excessive smoothing.

In Fig. 12 the regularized classifier is compared to simple average. The contour of a rectangular object against a background is learned in a synthetic binary image. The contour presents sharp changes in features (the filter weights change significantly in corners) which is in harmony with a square wave shape of the first and third derivatives. Simple constant average ($\zeta_p = 0.5, \forall p$) blurs the sharp changes, while the new regularizing process preserves them (the learned and regularized values agree). Figures 12(c), (d) show the final fitting for both cases and Fig. 12(e) shows two samples of the smoothed and the regularized components of the classifier along the contour.

The continuous classifier also allows us to add new control points to fit to concave shapes. After convergence, the algorithm tests the distance between adjacent control points; when this distance is greater than two times the initial distance a new control point is introduced. If new control points were added the optimization would be done again. The smoothed continuous classifier allows the snake to track changes in shape and texture features along the image sequence.

Figure 13 illustrates this ability: Fig. 13(a) contains a binary synthetic image where the object contour is learned, Fig. 13(b) shows the initial curve and Fig. 13(c) contains the final curve. If few contour patches are considered, erroneous fitting is achieved (Fig. 13(c)), but relying on the continuity of the classification vector, more control points can be dynamically introduced to allow a correct fitting (Fig. 13(d)). Figs. 13(e–g) show the inverse case (deformation in the opposite direction).

The classical snake has the well known difficulty in segmenting concave objects due to the curve is equally attracted by different contour configurations (Fig. 14).
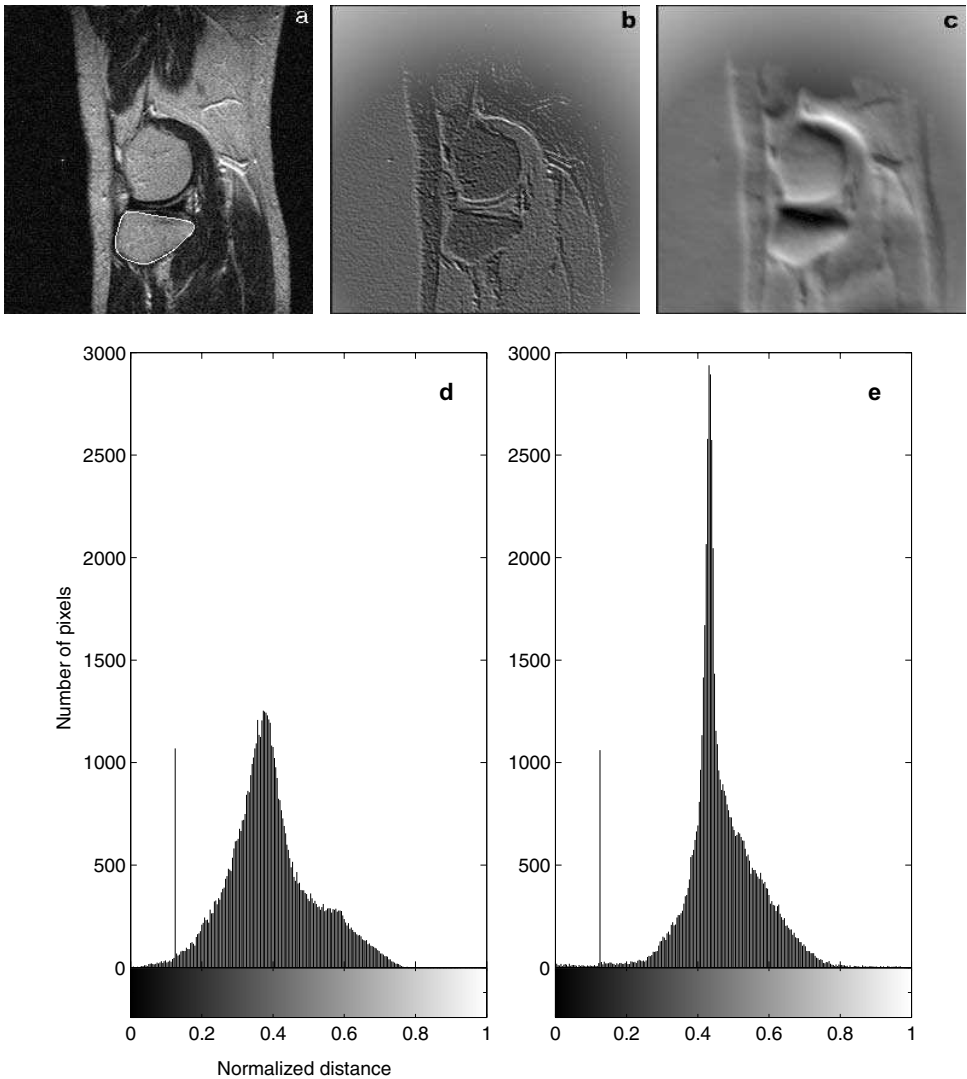
Fig. 11. Training contour: MR Image of the knee joint (a); distance map images generated by the classifier of the contour patch placed in the middle top of the bone contour in (a) considering: only the smallest scale (b), all the scales used in learning phase (c); histograms for the distance map: for single scale (d), and all scales (e).

As a consequence, the curve falls in a zone of equal distance to different contour configurations and stabilizes far from the desired contour. Figure 14(d) shows the final fitting reached with classical snake where alternating colours of contour patches allow us to see that contour patches near to concavities increase in size because they are attracted by the ends of concavities.

Several methods have been proposed in the literature to cope with this problem. A well illustrated example about the behavior of traditional snakes in the
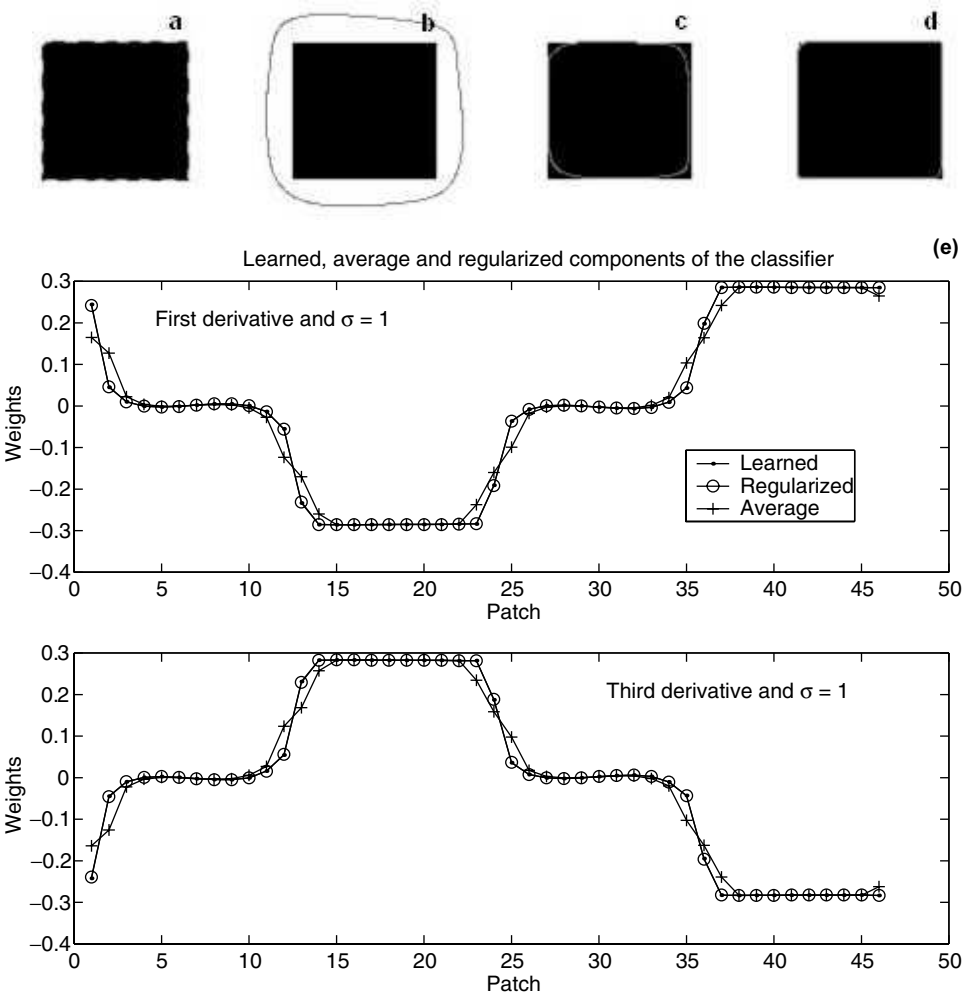
Fig. 12.   Learned contour (a); initial contour (b); final contour using smoothed (average) classifiers (c); final contour using regularized classifiers (d); average ($\zeta_p = 0.5$) and regularized (preserving real sharp changes) components of the continuous classifier for the rectangle contour (e).

presence of concavities and revision of previous proposed solutions can be found in Neuenschwander et al.,[35] and Xu and Prince.[63] We think that our approach is more general in the sense that it handles non-convex shapes in a natural way, and has the ability to adapt to concave contours due to the selectivity of external energy and the smoothness of the discriminant function (classifier) along the contour.

The discriminant snake allows the tracking of contours with some abrupt changes in shape or texture through slices. Figure 15 illustrates this property in a CT femur slice. Figure 15(a) shows the learning contour where alternating colours indicate different contour patches. Between this slice and the next one, a sharp change exists around the pointed patch; Fig. 15(e) shows the profiles normal to the pointed patch in the two adjacent images. However, the potential minimum for the
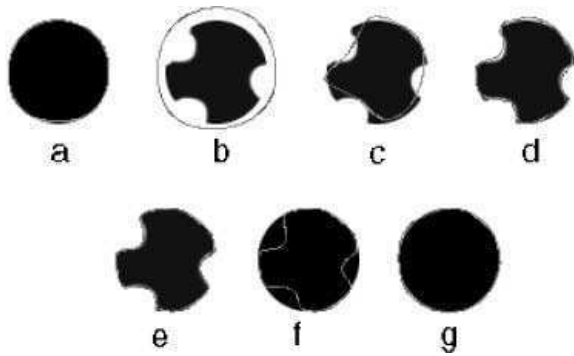
Fig. 13.   Training contour (a); initial contour (b); intermediate contour (c); final contour after interpolation (d); training contour (e); initial contour (f); final contour (g).
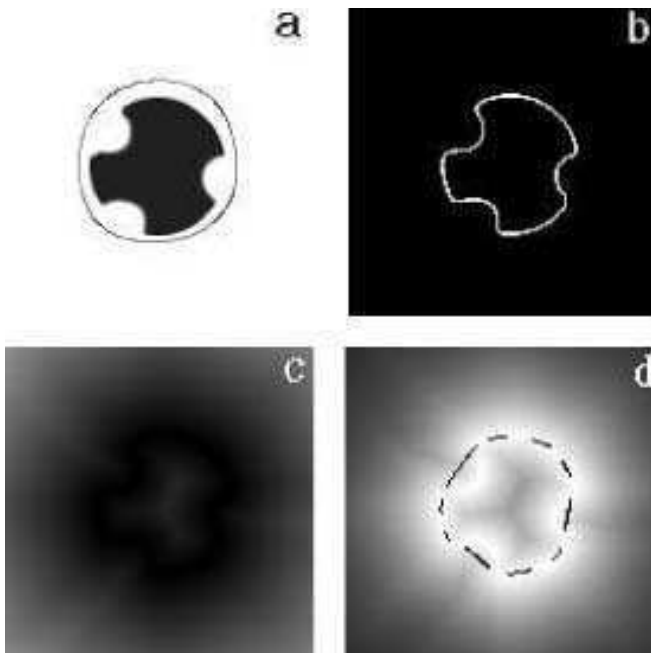
Fig. 14.   Initial contour (a); edges (b); classical distance map (c); fitting of the classical snake in black over the inverted potential (d).

corresponding patch in the next slice is in the correct place as Fig. 15(b) shows. Figure 15(c) contains the translation of learning contour in the previous slice to the next slice, and Fig. 15(d) shows the correct final delineation.

3.1.3. *Generalized description of features*

Our snake does not have a predefined goal, but it learns the desired contour features for each specific segmentation problem. Moreover, each snake patch searches for its
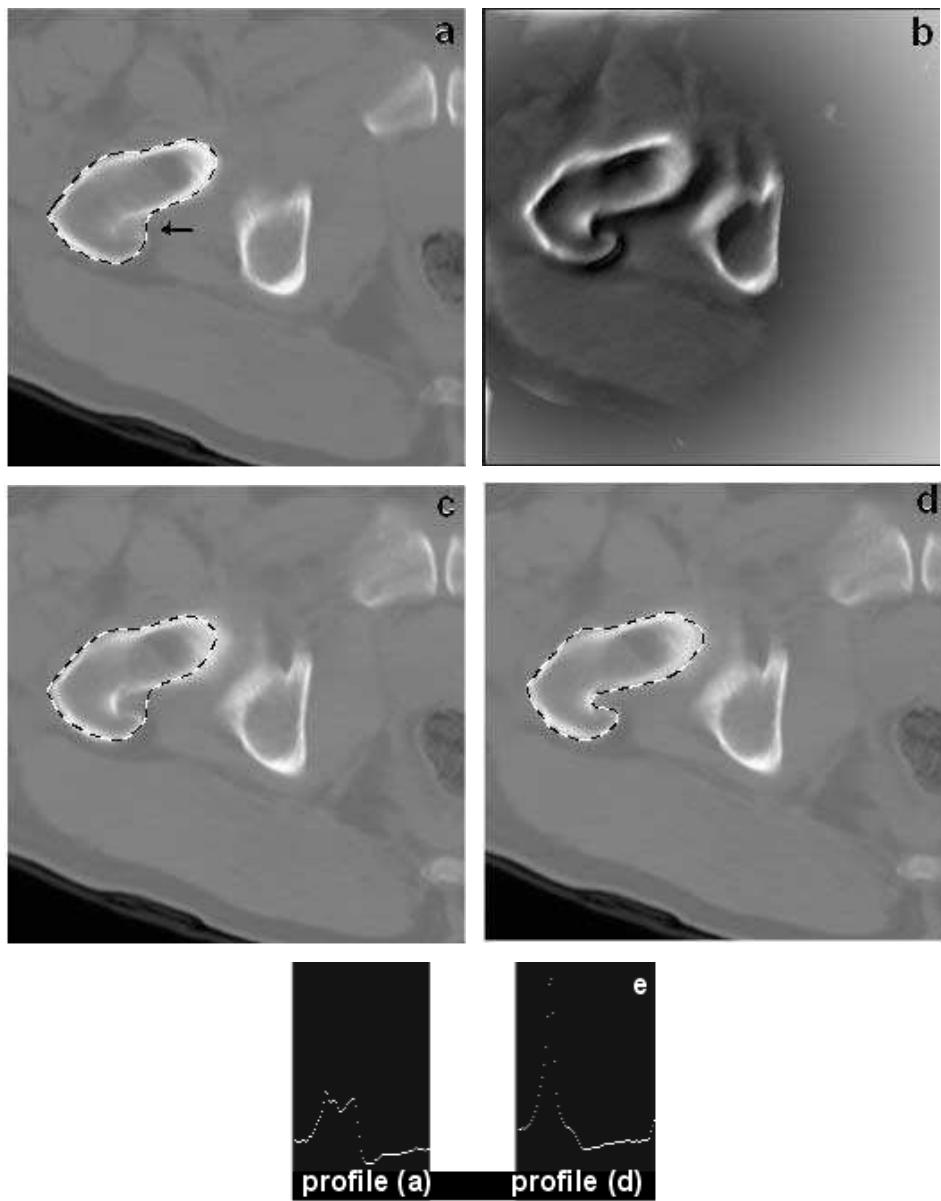
Fig. 15. Training contour (a); energy image generated by the classifier of the pointed patch in the following slice bf (b); initial contour (c) and final fitting (d). Corresponding profiles normal to the pointed patch (e).

specific contour class, which avoids ambiguities in the event of there existing several neighbouring objects.

Figure 16(a–d) illustrates the second capability. Figure 16(a) contains the original image and Fig. 16(b) shows the contour used for a supervised learning of
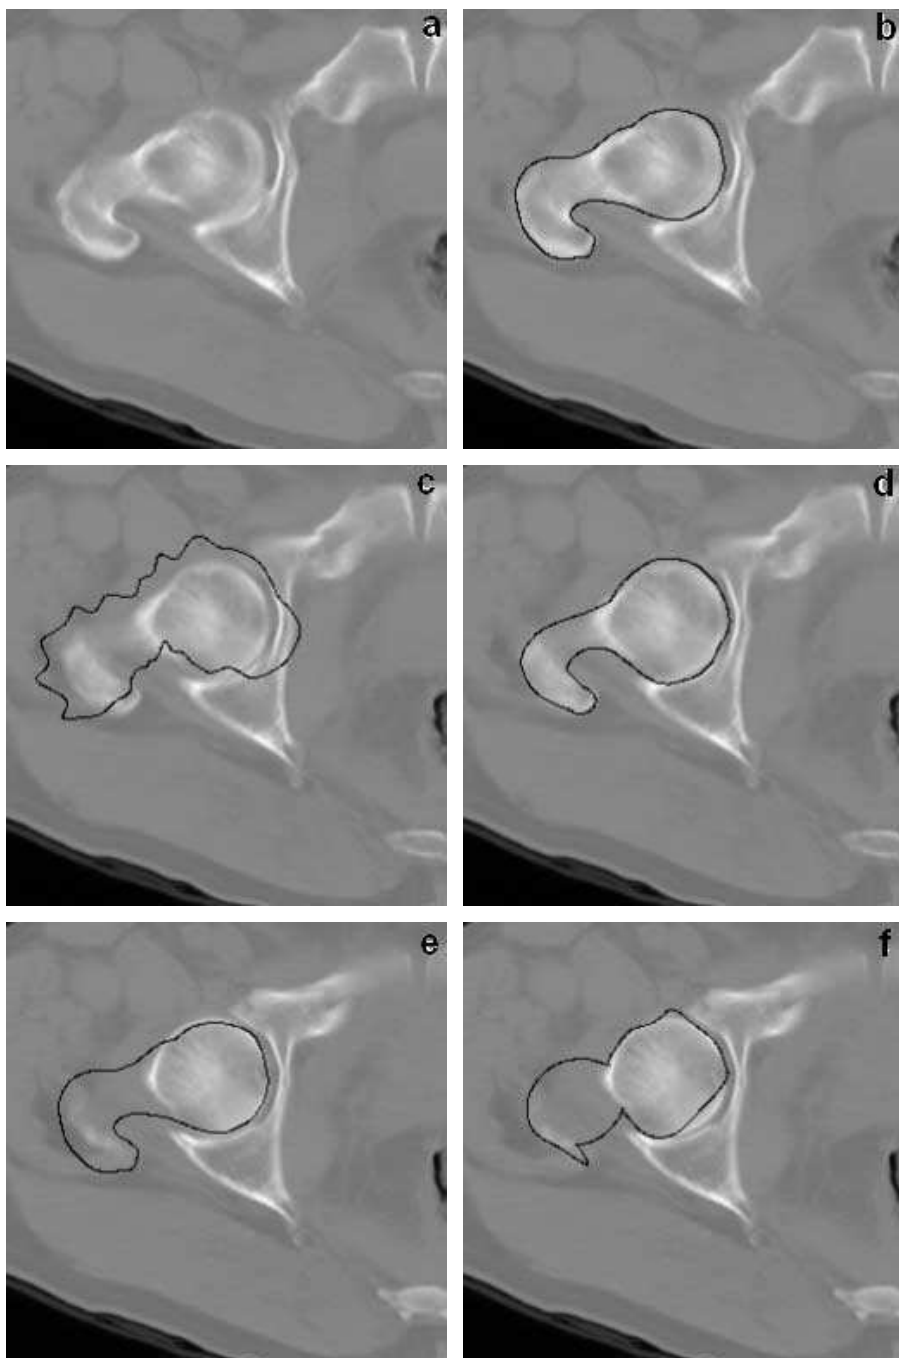
Fig. 16. Original image (a); learned contour (b); initial contour and final contour for a moderately distant slice (c, d); initial contour and final contour for a very distant slice (e, f).

different parts of the external cortical bone of femur in an CT image. To show the robustness of our proposal we consider another slice (non adjacent to that) which has different contour configuration (texture and contrast). We use the classification vector in the previous slice, but with the initial contour deformed, and placed between two bones, see Fig. 16(c). Figure 16(d) represents the final contour fitted to the bone of interest by our method. Although the initial contour is not close to the final contour and the characteristics of the contour are different, the snake found the contour more similar to the learned one, and this provided a good result. Different parts of the cortical bone have different filter responses, hence they belong to different classes and the snake moves towards the desired contour even in the case of nearby objects.

If the texture features and the geometry changes dramatically then the segmentation fails. Figure 16(e) shows the learned contour in Fig. 16(b) over the slice to be segmented. The final result shown in Fig. 16(f) is incorrect because the learned contour and the actual contour of interest are very different. In a normal medical image sequence the distances between adjacent slices are short so we can hope not to find such strong changes.

Now, we are going to show the ability to learn different features by segmenting medical images (from several modalities) where the desired object contours have different characteristics.

Figure 17 shows the segmentations of a set of slices where a considerable change in contour shape takes place. The full sequence is made up of 60 slices with a distance between slices of 5 mm. This is a typical case where the 3D image has not the same resolution in all the spatial dimensions, and therefore, a multislice segmentation approach is the best option. The supervised learning takes place in the first slice and converged active contours and updated classifiers are propagated from slice to slice. The snake is capable of following texture and shape changes. Note that the new approach naturally fits convoluted shapes. Moreover, the new external energy eliminates ambiguities between internal-external contours and contours of nearby structures. Figure 2 shows the 3D reconstruction of the femur.

We also apply our method to MR images. Figure 18 shows three samples of the segmentations of a sequence of 15 MR images of the distal femur. In these images contour has mostly the appearance of an intensity valley, unlike CT images where the contour features are mainly edges. Figure 2 shows the 3D reconstruction obtained from the segmentation.

The next imaging modality that we consider is the IVUS in two different cases: abdominal aorta (AA) and coronary vessel. In these images the contours have less contrast and the signal to noise ratio is higher. Figure 19 illustrates the segmentation of a sequence of IVUS images of the abdominal aorta of a pig. This figure contains the first image with the initial training contour, and other intermediate ones of a sequence made up of 100 slices. Classifiers are learned in the first slice and the deformed contour is carried to the next slice. Afterwards, the classifiers are again updated and transferred together with the final contour to the next slice. Although
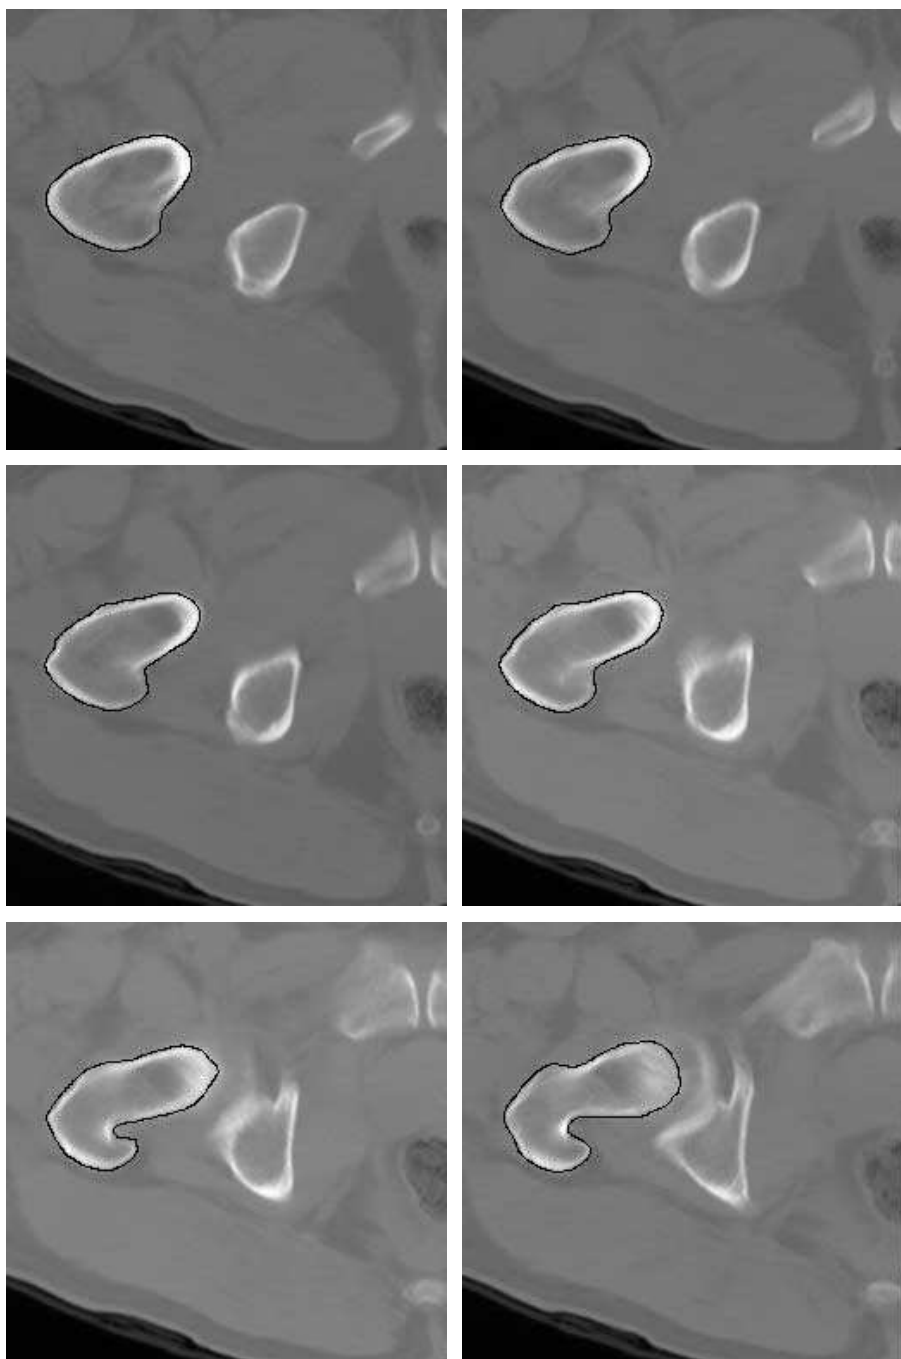
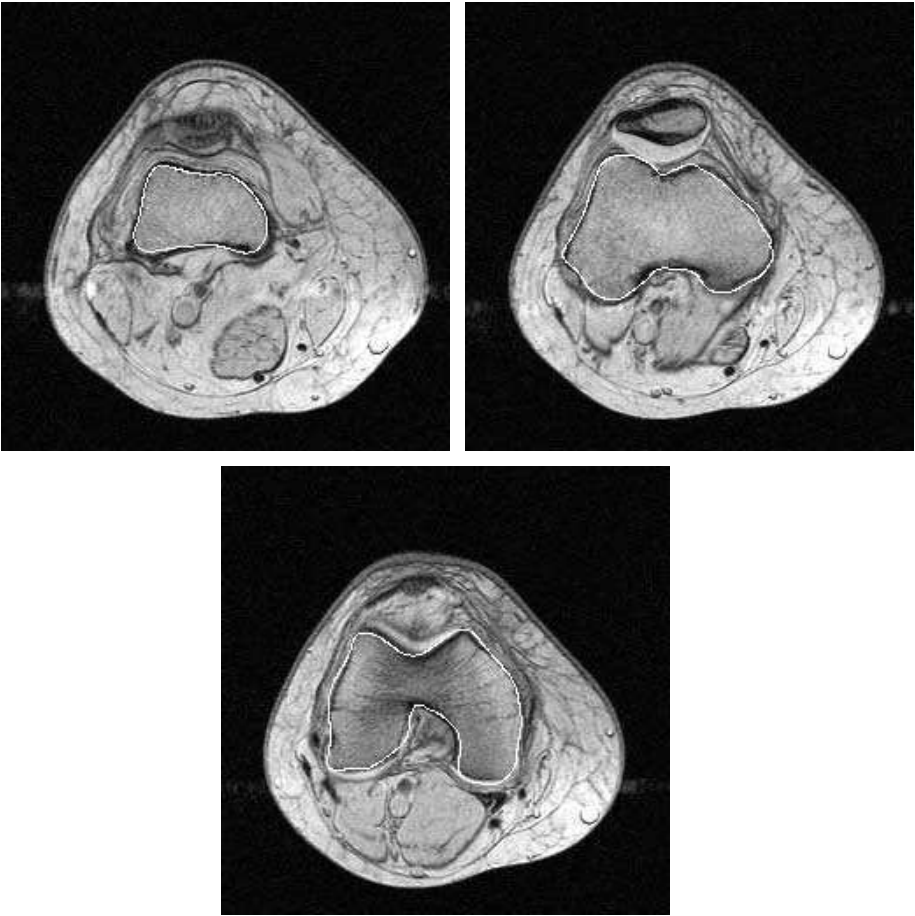Fig. 17.   CT images of slices 35–40.

Fig. 18.   MR images of the distal femur.

there are several different contours nearby, the snake fits the learned one. The reconstruction of the 3D surface derived from the segmentation is also shown.

4. Tracking by Active Contours

The snake model was the first of a plethora of mathematical formulations that have been proposed mainly for both segmenting and tracking problems. Among them, level-set techniques[5,28] are the more suitable method to track moving objects. While the snakes are physically motivated and represent the contours explicitly as parameterized curves in a Lagrangian formulation, level-set approaches are implicit, intrinsic, parameter free and can manage topological changes in the moving target. Nevertheless almost all of them have in common high computational requirements which might limit their use in tracking tasks needing fast time response. This inconvenience is alleviated by the development of new strategies for the numerical
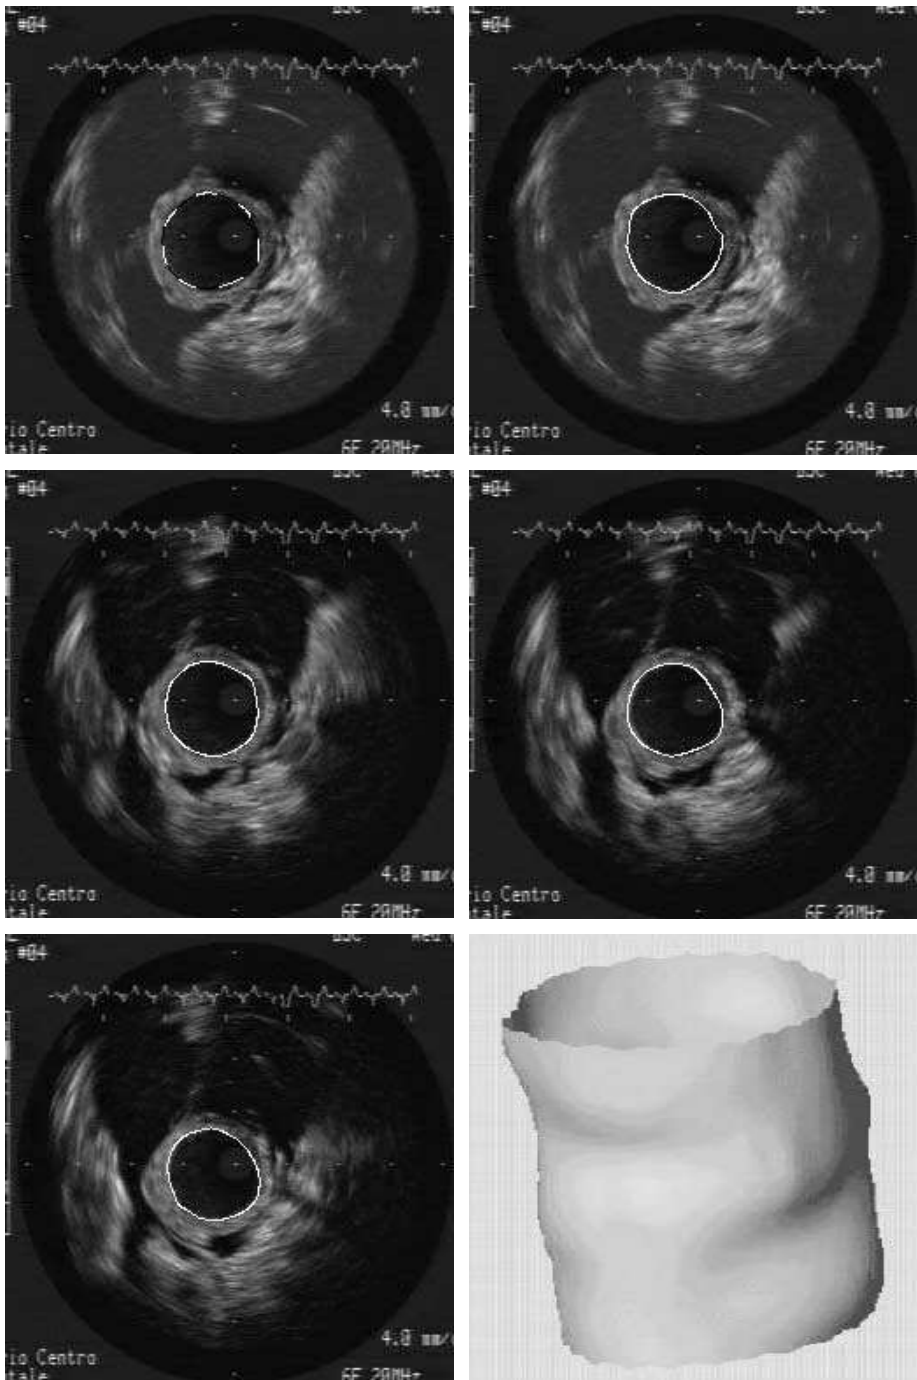
Fig. 19. Initial training contour, several slice segmentation examples, and the reconstructed internal surface of the abdominal aorta of a pig.

simulation of the equations which govern the dynamic of the curve evolution which usually lead towards a compromise between processing speed and flexibility in the contour evolution.[1,4,40,52,72]

The so-called pixel-level snakes (PLS) appeared originally intended to resolve the high computational cost inherent to the classical active contour techniques, mainly, in tracking temporal sequences. They are based on a pixel-level discretization of the contours and on a massively parallel computation on every contour cell which lead to a high speed processing without penalizing the efficiency of the contour location.[61] They represent a topographic and iterative active contour technique where the contours are explicitly represented and evolve guided by local information. In the context of PLS, the active contours are represented as sets of 8-connected activated pixels, i.e. black pixels in a binary image called contour image. This binary image has the same dimensions as the original image containing the objects to be segmented and tracked. The contour evolution consists on an iterative process of activation and deactivation of the contour pixels based on binary and local morphological operations extended along the four cardinal directions. The goal after each cycle (four iterations, one for each cardinal direction) is to obtain new well-defined contours slightly shifted and/or deformed in order to come closer and fit themselves to the boundaries of the objects of interest.

PLS can be conceptually defined based on three different modules which interact dynamically as depicts Fig. 20:

(1) A module responsible for extracting the information to guide the contour evolution. This includes the computation of the internal potential, pressure (balloon) potential, and the combination with the external potential derived from image features.
(2) A module dedicated to the contour evolution. This performs an iterative pixel-to-pixel shift of the contours driven by the guiding information.
(3) A module undertaken to handle the possible collisions between contours and to change or preserve the topology of the contour as required.

In the following the working of these building blocks are described in detail.

4.1. *Guiding information extraction*

Guiding forces compel the contours to move towards the desired image features. These forces derive from potential any $P(x, y)$ defined in the image space:

$$\mathbf{F} = -\nabla P(x, y) = \left(-\frac{\partial P}{\partial x}, -\frac{\partial P}{\partial y} \right) = (F_x, F_y). \tag{10}$$

PLS is one example of the general technique of matching a deformable model to an image using energy minimization. Each pixel on the starting contour, subjected to some curvature constraints, moves into alignment with the nearest salient feature in an image. The direction of the movement coincides with the direction of the decreasing potential. In a pixel-level iterative technique only the sign of both
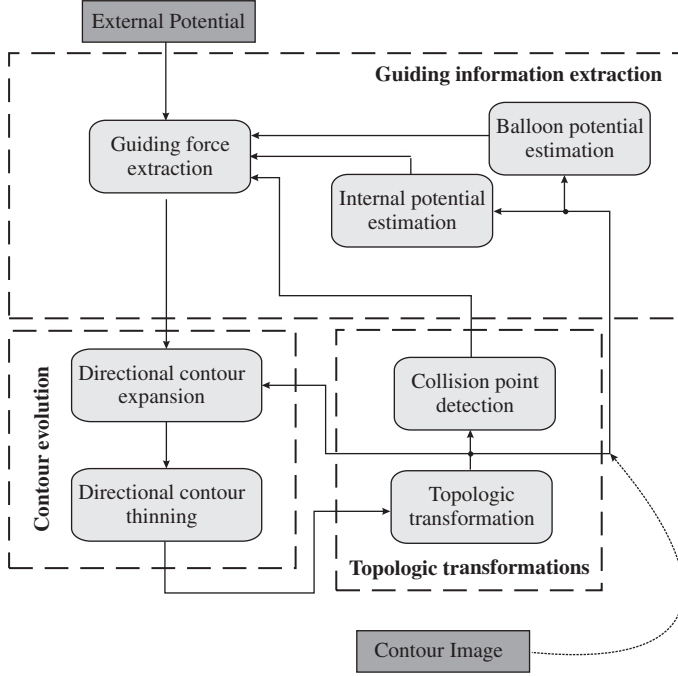
Fig. 20.   Flow diagram of the PLS algorithm showing the main modules and their iterations. The contour image is externally provided only for the first iteration (dashed line).

components of the guiding forces is actually needed. Therefore, given a direction of processing the guiding information should be the sign of the component of the guiding forces along that direction. These operations are performed in the so-called guiding force extraction module (GFE). Therefore:

$$OUT_{GFE}(NORTH) = -OUT_{GFE}(SOUTH) = sgn(F_y),$$
$$OUT_{GFE}(EAST) = -OUT_{GFE}(WEST) = sgn(F_x),$$

The derivation of the components of the guiding forces from the potential fields is approached by a directional gradient estimation based on Sobel-type operators (D_Gr). The sign of these forces are subsequently extracted by a thresholding operation (Fig. 21).

Therefore, the output of this GFE module will represent a binary map with activated pixels in those locations where the potential is decreasing along the direction under study. Figure 22 illustrates the operations in the GFE module by means of an example of contour evolution based on external potential. The sign of force components indicates the way of movement of each pixel on the deformable model along the given direction.

The external potential should be defined in such a way that the boundaries of the object of interest coincide with the minima in the image potential. The exact feature that defines an object boundary is strongly dependent on the particular application
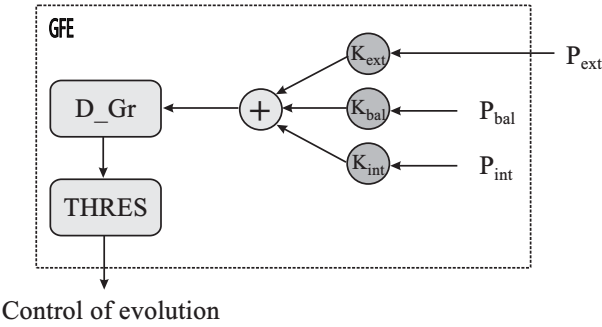
Fig. 21. Flow diagram for the guiding force extraction from the potential fields. P_{ext}, P_{int} and P_{bal} makes reference to external, internal and balloon potential respectively. The constants k_{ext}, k_{int} and k_{bal} weight the influence of each term.
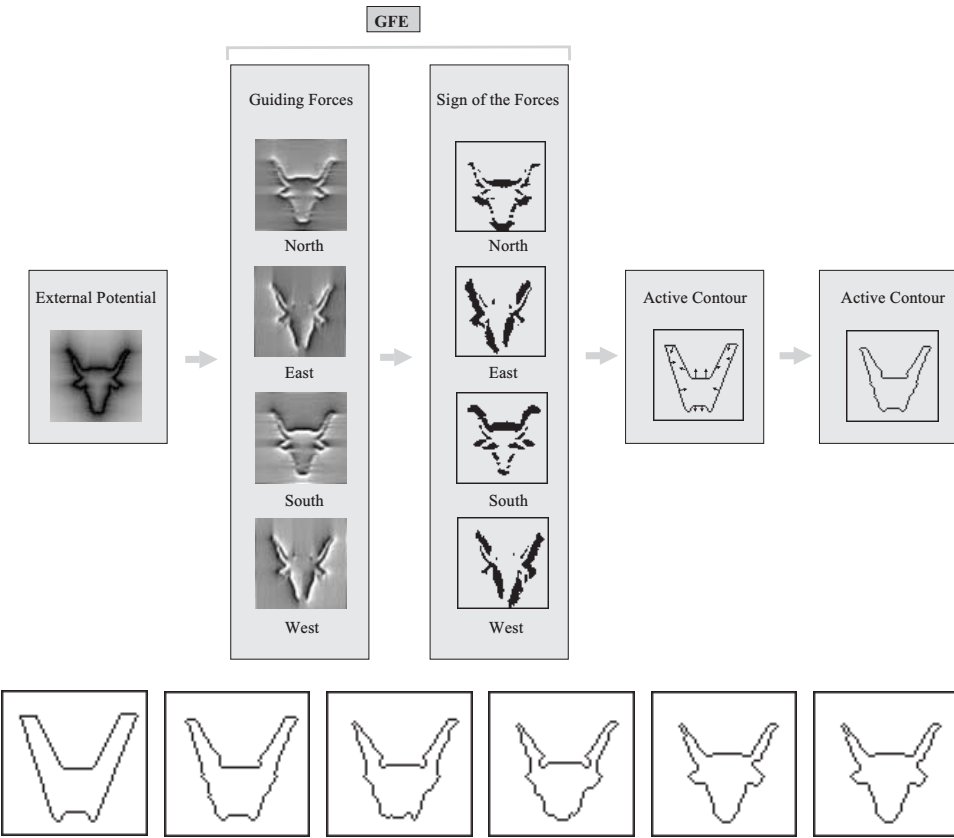
Fig. 22. (Up): Guiding force extraction (GFE) from an external potential. Lower potential is represented by lower intensity. By means of directional gradients the component of the guiding forces for each direction is obtained. (Bottom): Several snapshots of the contour evolution.

and represents an *a priori* knowledge that incorporates the PLS algorithm. Typical potential terms are[11]:

$$P_l = \pm G_\sigma * I(x, y),$$

where I is the intensity image, G_σ represents a Gaussian filter with a scale parameter σ. This potential attracts the curve to line features described by high or low intensity points after convolution with a Gaussian smoothing to mitigate the undesired effect of noise;

$$P_g = -|\nabla G_\sigma * I(x, y)|,$$

where ∇G_σ represents a Gaussian derivative filter with a scale parameter σ, and attracts the curve to intensity edges; and

$$P_e = -e^{-d(u(s))^2},$$

where $d(u(s))$ is the distance to the closest boundary point, which pushes the contour to the edge points of an edge map. An edge map is an image of contours, usually defined by zero-crossings in a Laplacian filtered image.

In addition, PLS incorporates internal regularizing constraints and balloon potential. The formers ensure that the model remains smooth limiting the amount that they can bend. The equivalent term of the internal potential in classical approaches depends on the tension and the flexion of the contour and it may be measured as a function of distances among adjacent points according to the considered discretization. This approach cannot be directly included in the PLS formulation because the contour is not defined as a predetermined number of discretization points but as a set of black pixels of a binary image, and its evolution is based on the activation and deactivation of pixels belonging to this contour image. However the desired smoothing effect can be obtained by assigning higher potential values to those pixels in the contour image situated out of the contour cavities, with respect to those situated inside. One way to perform this internal potential estimation (IPE) is by means of a recursive low-pass filtering or diffusion operation acting on the contour image.[61] Therefore, directional gradient operation acting on diffused contour image will originate positive internal forces which push the contour to reduce the local curvature and therefore to smooth the contour shape. This approach recalls the thin-plate energy in parametric deformable model strategies or the regularizing term in the implicit formulations. This idea is illustrated in Fig. 23, where one contour is guided by only this kind of internal forces. A directional gradient operation will originate forces proportional to the local curvature which guide the contour evolution in order to regularize its shape. It is well known that a planar closed contour whose evolution relies only on the local curvature will adopt a circular shape and finally will collapse. This behaviour is observed with the proposed internal potential estimation which demonstrates the curvature dependence of our approach.

When not subject to image forces a deformable model shrink until collapsing to a point due to the unbalancing effect of internal forces. Also, it is feasible that the
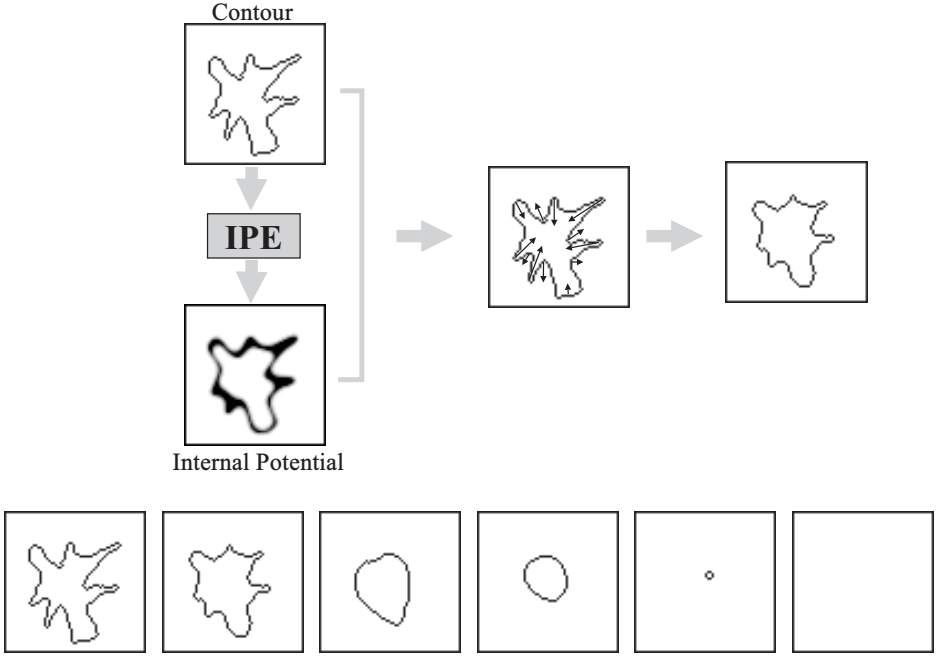
Fig. 23.   (Upside) Generation of the internal potential for the PLS. Lower potential is represented by lower intensity. (Bottom) Several snapshots of the contour evolution guided by the only internal forces.

deformable model remains trapped by sparse weak features far away the desired object boundary. Proceeding towards the avoiding of these problems, the so-called *balloon forces* have been introduced.[11] The incorporation of these forces can help the deformable model to trespass spurious isolated weak image edges and counteract their tendency to collapsing to point in the absence of image forces. The balloon forces can be expressed as follows:

$$\mathbf{F} = k_1 \mathbf{n}(s), \tag{11}$$

where $\mathbf{n}(s)$ is the unit vector normal to the curve $u(s)$ and k_1 is a constant which controls the strength of inflation or deflation, depending on the sign of k_1. The addition of a pressure term reduces the sensitivity of active contours to initialization and spurious features. However, in each run of the method, it has to be decided whether an inflationary or deflationary force is required.

Both implicit and pixel-level active contours can easily incorporate this kind of guiding terms. The implicit models implement them as geometric-independent advection terms.[37] The pixel-level snakes can effectively inflate or deflate the deformable model by the definition of a new potential field:

$$P_{\mathrm{bal}}(x, y) = s_{\mathrm{bal}} * \xi(x, y), \tag{12}$$
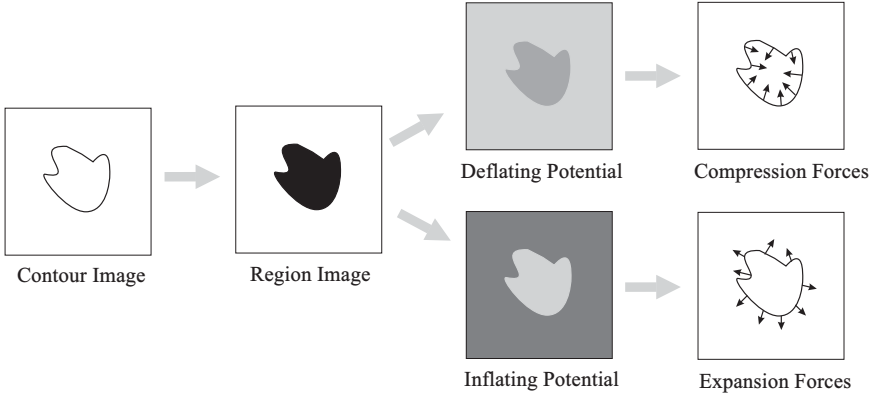
Fig. 24. Generation of inflating/deflating potentials for the PLS. Lower potential is represented by lower intensity.

where $\xi(x, y) > 0$ if (x, y) belongs to the set of locations enclosed by an active contour. Otherwise, $\xi(x, y) < 0$. $s_{\text{bal}} = +1$ defines an inflating potential and $s_{\text{bal}} = -1$ a deflating potential. Therefore in the discrete image space the balloon forces $\mathbf{F}_{\text{bal}} = -\nabla P_{\text{bal}}(x, y)$ can be approached by finite differences. The implementation of the balloon potential is mainly supported by means of a weighted hole filling operation as illustrates Fig. 24. A directional gradient operation originates forces guiding the contour evolution outwards (inflation) or inwards (deflation).

The global potential from which the guiding forces of the contour evolution are extracted will be a combination of the previously described potentials:

$$P(x, y) = k_{\text{int}} P_{\text{int}}(x, y) + k_{\text{ext}} P_{\text{ext}}(x, y) + k_{\text{bal}} P_{\text{bal}}(x, y), \qquad (13)$$

where the parameters k_{ext}, k_{int} and k_{bal} weigh the contribution of the external potential and the *a priori* knowledge on shape and feature strengths of the objects to be tracked. There is no exact rule to determine the influence of each kind of potential field. Nevertheless the search of the parameter values can be based on some general hints. Usually the external potential should have the higher influence in the global potential because it contains the information of the image under processing. However, a low value for k_{int} can produce protuberances in the contour shape, which increase the internal potential and thus compensate the initial weak influence of this kind of potential. Therefore a low weight for the internal potential is usually preferred. Finally the influence of the balloon potential is often required to guide the contour evolution when the external potential is too weak. Therefore the k_{bal} value should be clearly lower than k_{ext}. On the other hand its relation with the internal potential depends on the nature of the balloon potential. The internal potential enforces the deflating potential and counteracts the inflating potential.

4.2. *Contour evolution*

The contour evolution is based on iterative operations of expansion and thinning along the four cardinal directions conducted by the binary map from GFE. For each direction of processing, the contours are shifted one pixel along it by the directional contour expansion (DCE) module depicted in Fig. 25. The resulting binary image is combined with the output of the GFE module by a logical AND which keeps activated only those shifted contour pixels whose locations coincide with activated pixels in the GFE output. Finally, the combination of the result with the original contour image by a logical OR will produce a binary image with the contours expanded, or dilated, only in those locations where the guiding forces are positive along the current direction of processing.

The expanded contour image is processed in the directional contour thinning (DCT) module composed by the following operations (see Fig. 25): After DCT(1) those contour pixels which are frontier along the processing direction are deactivated. This leads to thinned contours but with possible breakpoints. Steps DCT(2) and DCT(3) act on areas with opposite slope in order to reactivate those pixels (deactivated in DCT(1)) needed to restore the continuity of the contour.[61] The outcomes of these three operations are combined by logical OR operations which results in a binary image containing well-defined contours slightly shifted and/or deformed based on the guiding information. DCT implements the morphological hit and miss operations which obey the rules portrayed in Fig. 26 only for the North
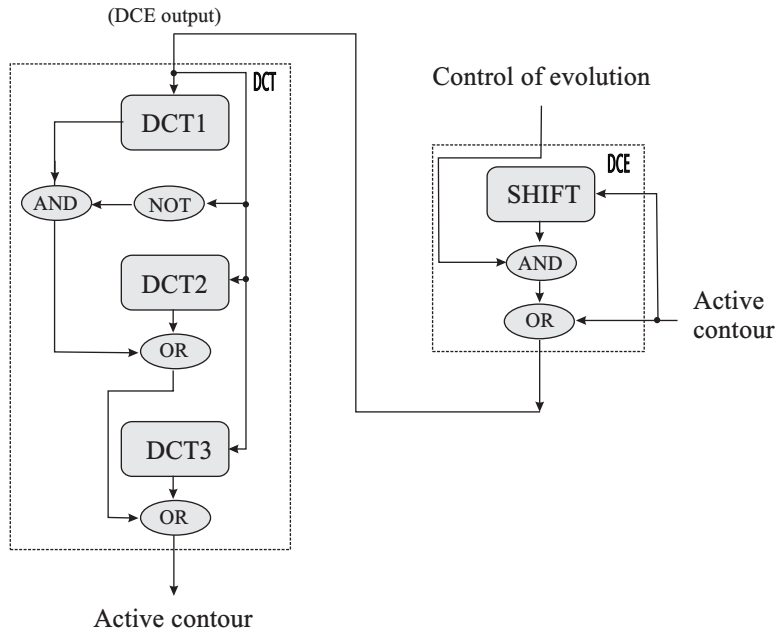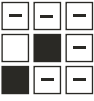

Fig. 25.   Flow diagram containing the operations to be performed in DCE and DCT modules.

Operation	Pattern	Yes	No
DCT1		□	■
DCT2		■	□
DCT3		■	□

■ = ON   □ = OFF   ⊟ = don't care

Fig. 26. Patterns and local rules for the morphological operations of the directional contour thinning along the North direction. The first column labels the operation. The second column refers to the pattern to be satisfied.

direction. Given a pixel, if the configuration of its neighbouring cells coincides with the pattern then the new pixel value should be that indicated in the third column. Otherwise it should be that in the four column. For the processing along the other three cardinal directions the patterns should be suitably rotated.

4.3. *Topologic transformations*

Segmentation and tracking by means of active contours are highly effective for matching boundaries provided the initial deformable models are close to the objective. However, sometimes the number of objects in the scene does not coincide with the number of initial contours into the contour image. This can be because of a bad initiation process due to insufficient information about the domain, or the topology of the same object changes from one frame to the next (branching), or two objects overlap. In these situations the collision between different contours (or different parts of the same contour) may occur. Pixel-level snakes, due to their characteristics of evolution and their nonparametric nature, can handle these changes of topology by simple inspections of the contour map. This capability notably increases the number of tasks where the PLS can be applied. Topologic transformations can be approached by PLS in different ways. One of the most efficient is based on implicit deformable models where contours are implicitly defined as wavefronts of propagating waves.[62] In this framework, when two deformable models (or two parts of the same deformable model) collide, the collision points no longer belong to the set of contour pixels. Therefore, changes in topology are handled implicitly by updating
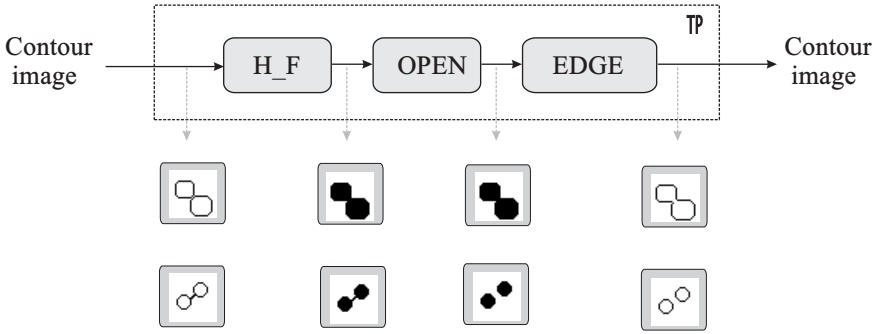
Fig. 27.   Flow diagram of the operations to manage changes in contour topologies.

the borders of the regions enclosed by the deformable model. The transformation
is performed by the three operations indicated in Fig. 27. The regions are obtained
from the contours by means of a hole filling operation (H_F) followed by a one-step
morphological opening (OPEN). In the last step the region boundaries are obtained
by a binary edge detection (EDGE) recovering the deformable model.

Figure 28 illustrates the capabilities of our method to manage topologic trans-
formations. In this case study the target are the blood vessel contours. Retinal
blood vessel morphology can be an important indicator for many diseases such as
diabetes, hypertension and arteriosclerosis, and the measurement of geometrical
changes in veins and arteries can be applied to a variety of clinical studies.[69] The
figure contains several snapshots of the evolution of deformable models along blood
vessels from a retinal angiography image. The contour evolution was guided by the
combination of the intensity image ($k_{ext} = 0.7$) and inflating potential ($k_{bal} = 0.3$).
As it can be seen the delineation of blood vessels entails the evolution along narrow
cavities and several changes of topology.

Based on the commented strategy topologic transformations can be performed.
Nevertheless, there are some practical applications where the contour topology
should be preserved.[20] An example is illustrated in Fig. 29, showing the results
of segmenting a sequence of computed tomography images of proximal tibia and
fibula. For each slice image, the initial contours correspond with the outcomes of
the segmentation of the previous slice image. The aim of the segmentation is to
delineate the outer side of cortical bones as a previous step to the 3D reconstruc-
tion of the knee. The width and the high density of the cortical bone and also
the clear separation between tibia and fibula make the segmentation in the distal
part non complex. As we approach the proximal part of the tibia and fibula, the
thickness of the cortical bone decreases and the proximity between the bones is
reduced. When knee injuries are considered problems of bone loss, cortical bone
narrowness, malformations, etc., can be found. This translates into badly defined
external contours of the bones, little separation between the internal and external
contour of the cortical bone and the proximity between tibia and fibula, making
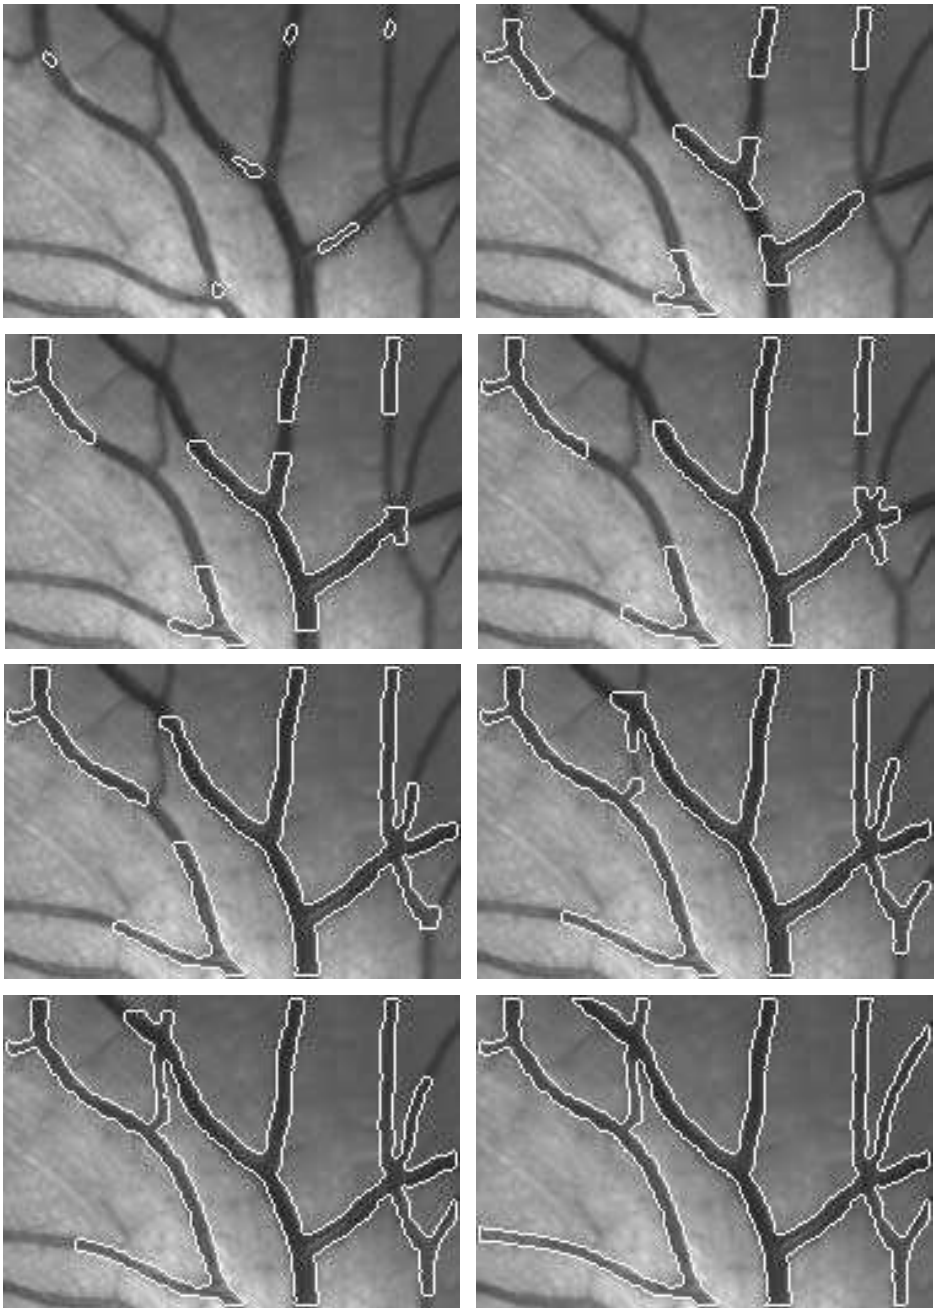
Fig. 28. Example of segmentation of blood vessel in an angiogram of retina from several initial contours. The sequence goes from left to right and top to bottom.
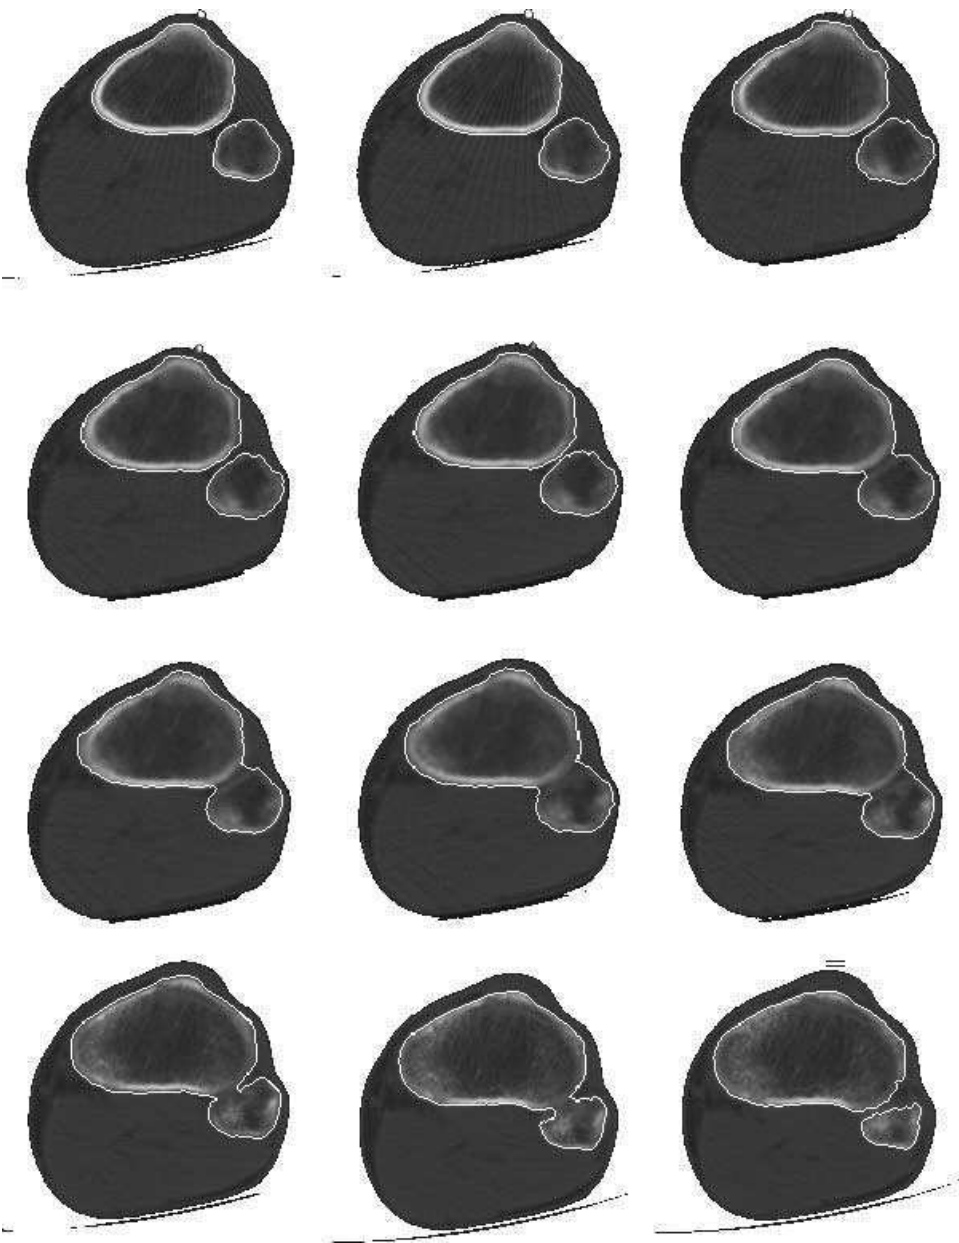
Fig. 29.   Example of segmentation on a CT image sequence corresponding to the proximal part of the tibia and the fibula. The topologic transformations impede to correctly define both structures.

the segmentation process particularly difficult. Taking this into account we have approached the segmentation of tibia and fibula simultaneously. As it can be seen in Fig. 29 a topologic merging of the two evolving contours impedes to distinguish the tibia and the fibula frontiers. This situation changes in higher slices where both bones becomes again separated provoking a new topologic transformation.

In order to avoid undesired changes in the contour topology those locations where contour collision may appear should be pre-estimated. This can be performed easily with PLS because the contours evolve as the effect of activation and deactivation of pixels in the contour image. Thus, the contours move pixel to pixel which allow to foresee the contour locations and shapes in the next iteration. The collision point detection (CPD) is carried out mainly by a simple pattern recognition which takes as input the binary contour image and returns a binary image with white pixels in those locations where a collision between contours can appear in the next iteration. Therefore, by the projection of this binary map onto the output of the GFE module, the pixel activation can be avoided on those conflictive locations and consequently the contour collision will be prevented. Figure 30 shows the flow diagram with the operations to be performed in CPD. These are described as morphological hit and miss operations. Figure 31 shows the patterns and rules to be fulfilled in the processing along North direction. Given a pixel, if the configuration of its neighbouring cells coincides with the pattern then the new pixel value should be that indicated in the third column. Otherwise it should be that in
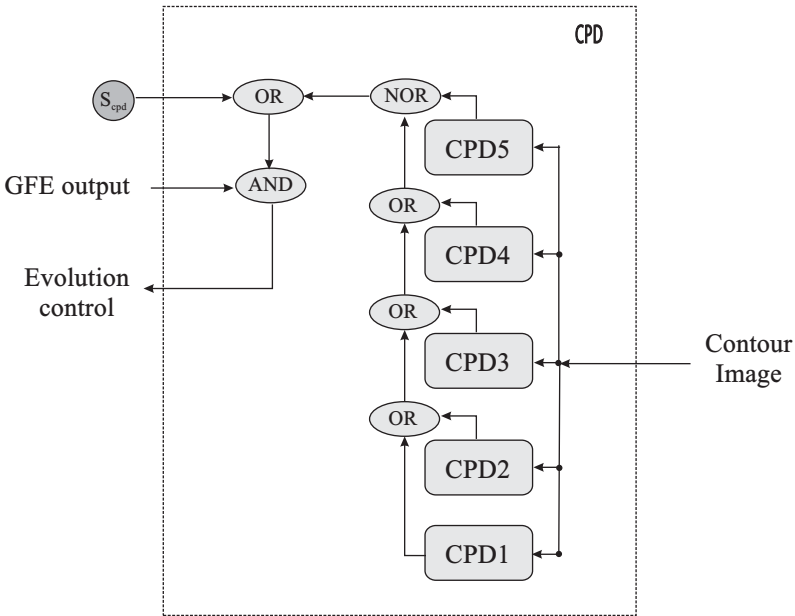

Fig. 30. Flow diagram of the operations to preserve changes in the contour topology. The parameter s_{cpd} is a switch which enables (logical '1') or disables (logical '0') the CPD action.

Fig. 31.   Patterns and local rules for the morphological operations of the collision point detection along the North direction. The first column labels the operation. The second column refers to the pattern to be satisfied.

the four column. For the processing along the other three cardinal directions the patterns should be suitably rotated.

In a physical context, the effect of this operation is equivalent to the generation of an infinite potential barrier between contour pieces to avoid the collision. Figure 32 shows the same example as Fig. 29 but preserving the contour topology by the action of the CPD module.

4.4. *PLS implementation issues*

PLS algorithm consists of simple local dynamic convolutions, morphological hit and miss and simple arithmetic and logical operations. Therefore it meets the requirements to be implemented onto simple instruction multiple data (SIMD) processor arrays, particularly on those architectures based on the cellular neural network (CNN) concept[10] like the CNN Universal Machine (CNNUM).[49] CNN constitutes
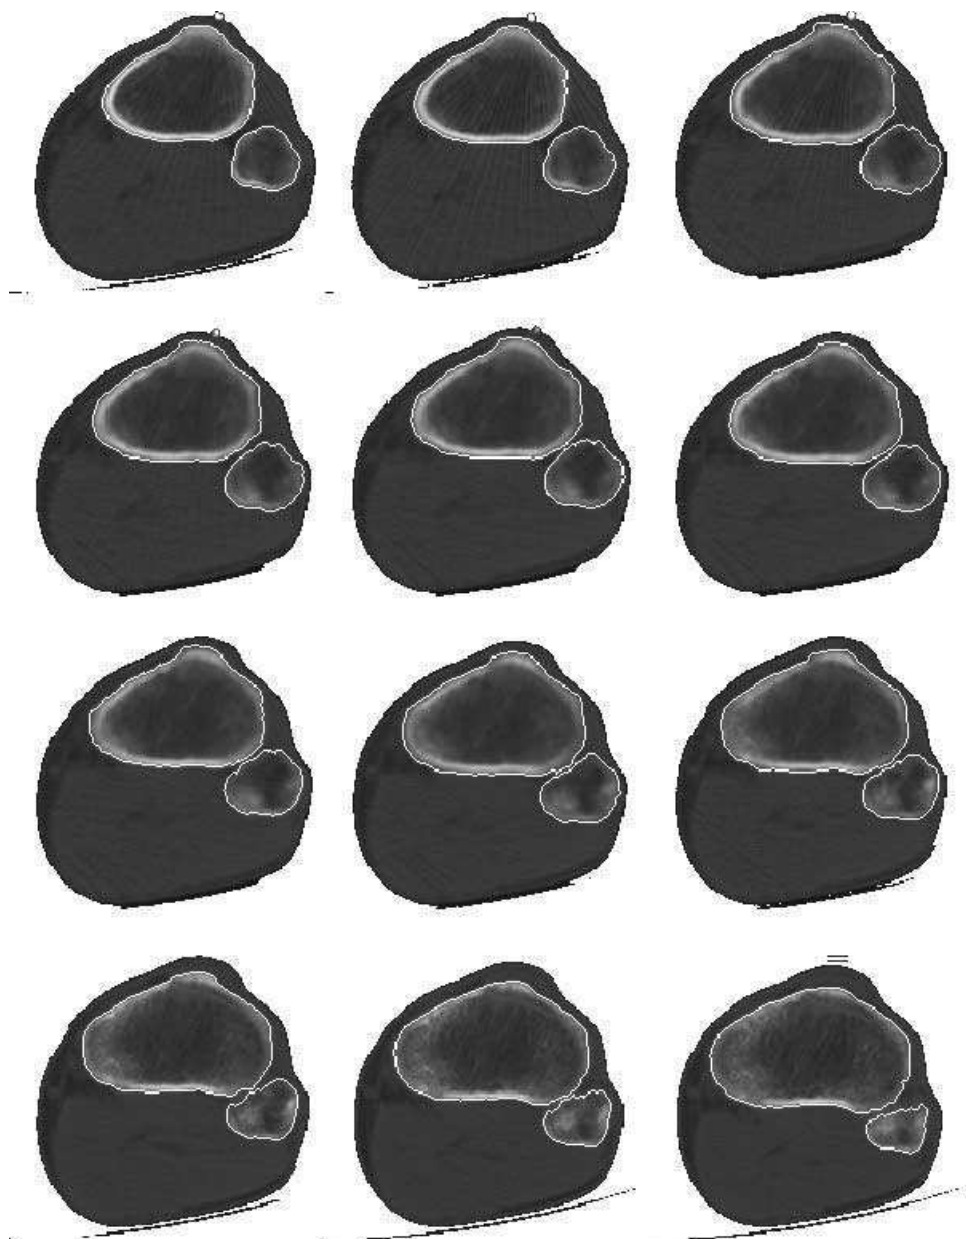
Fig. 32. Example of segmentation of tibia and fibula when avoiding topological merging of contours.

Table 1. Execution time of all major processing modules in the PLS algorithm extracted from the implementation on the ACE4K.

Module	Iteration	Cycle
GFE	$250\,\mu s$	$1000\,\mu s$
DCE	$60\,\mu s$	$240\,\mu s$
DCT	$160\,\mu s$	$640\,\mu s$
CPD	$300\,\mu s$	$1200\,\mu s$
TP	-	$760\,\mu s$
BPE	-	$750\,\mu s$
IPE	-	$40\,\mu s$
Maximum		$3820\,\mu s$

a class of recurrent locally coupled array of identical and simple processing elements. Since the coupling between processing elements is exclusively local, the hardware complexity does not increase with the array size. This has made the hardware implementation of a considerable number of processing elements into a chip feasible.[27,44]

The projection of the PLS-algorithm onto CNN structures permits to exploit the characteristics of massively parallel processing of this kind of architectures. This fact guarantees a fast computation making the pixel-level snakes a valid tool for those applications needing fast time response like the segmentation and tracking of moving objects.

The described PLS-algorithm has been implemented and tested on a 64×64 CNNUM chip (ACE4K[27]).[62] In Table 1, the execution times for each module of the algorithm in Fig. 20 are gathered.

Note that the estimation of the internal and balloon potentials (IPE and BPE) are actually required only once per cycle (one cycle represents four iterations, one for each cardinal direction). Furthermore the topologic transformations can be checked once per cycle or even less since the changes of topology are correctly performed in any time after the collision between contours. On the other hand, the action of modules TP and CPD is never required at the same time. The maximum time in Table 1 makes reference to the more expensive case in processing time, i.e. that one including CPD but not TP. Therefore the implemented algorithm requires less than 4 ms to complete one cycle running on the 64×64 CNNUM chip. We have observed that in real time applications like video object segmentation and tracking less than ten iterations per frame are usually needed. Therefore, even with the full version of the algorithm the speed processing of 25 $frame/s$ is feasible.

In order to illustrate the capabilities of PLS in segmentation and tracking of moving structures we have approached the automatic tracing of the boundaries of the human left ventricle (LV) from sequences of ultrasound (US) echocardiography images. The aim of the operation is to segment and track the contour of the left ventricle from frames as a previous step of the determination of the ventricular volume and the wall motion in different stages of the cardiac cycle. The tracing of LV boundaries from 2D US images sequences covering the entire cardiac cycle
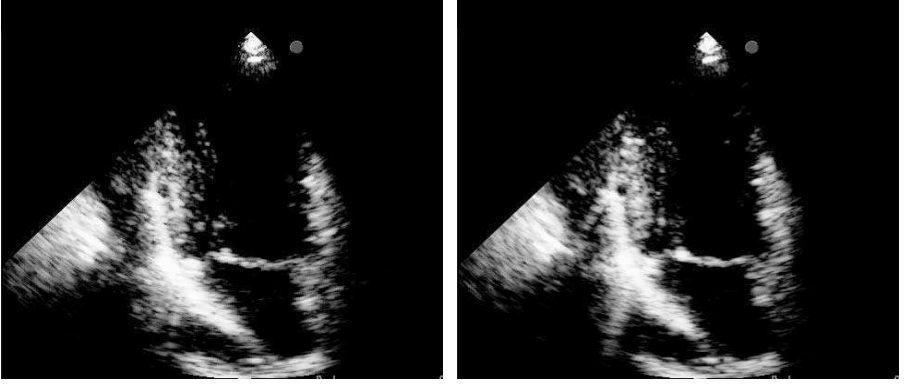
Fig. 33.   Two consecutive frames in a 2D US image sequence.

allows quantitative interpretation of LV contraction dynamics. Furthermore, the 3D motion estimation and reconstruction of cardiac chambers can be efficiently approached from the emergent 3D rotational echocardiography.[50] Following, we will show segmentation results from both 2D a 3D US image data.

The proximity between the boundaries of the heart chambers in two consecutive frames depends on the frame rate of the US image acquisition process and the cardiac cycle of the patient. Nevertheless, usually they are close enough to make suitable the use of active contour techniques in this kind of application.[7,33] Figure 33 shows two consecutive frames corresponding two a 2D sequence with 15 frames per cardiac cycle.

The proximity between boundaries of interest in two consecutive frames allows the control of the active contour based on only local information. The contour evolution is guided by a combination of external potential ($k_{ext} = 0.5$) from the US images together with internal potential ($k_{int} = 0.025$) to keep smooth the contour shape. In order to implement the complete system in the ACE4K chip, the external potential is derived from an inverted version of the current frame filtered by a local diffusion which reinforces the evolution towards high intensities, combined with the result of a diffused edge detection onto this image which encourages the contour evolution towards edges (Fig. 34).

Therefore by using as initial contour the result from the previous processed frame it is possible to fit it to the boundaries of the chamber in the current frame (Fig. 35).

This approach is valid for intermediate frames but not for the first frame of the sequences. For the first frame an inflating potential ($k_{bal} = -0.2$) is considered to put one initial seed situated into the left ventricle, close to the boundaries of interest (Fig. 36).

Figures 37 and 38 show experimental results from the on chip implementation of the PLS algorithm, corresponding to the tracking of the human left ventricle from 2D and 3D sequences of US echocardiography images.
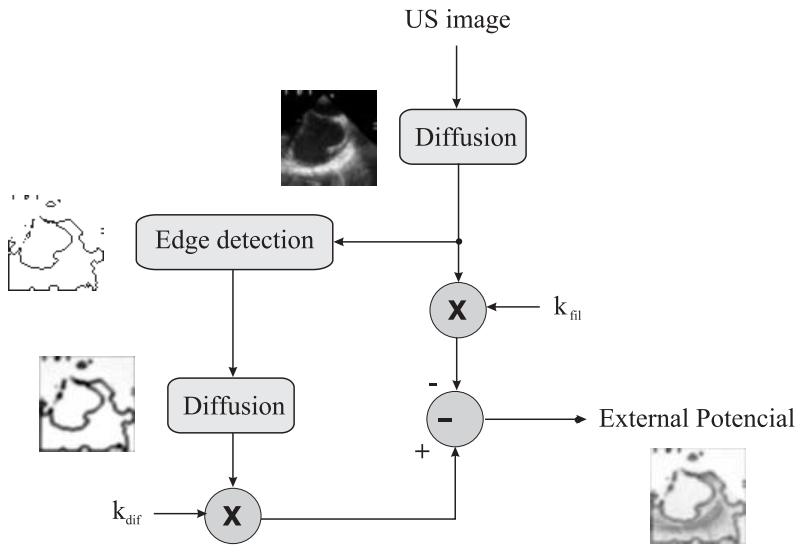
US image

Fig. 34.   External potential estimate for the contour location in US image of the human left ventricle. k_{fil} and k_{dif} are constants which weigh the influence of the filtered image and the image resulting from the diffused edge detection.
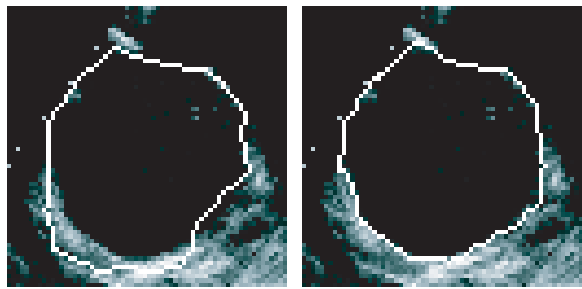
Fig. 35.   Processing on an intermediate frame. The initial contour on the left (from the previous frame) and the final contour on the right.

To validate the left ventricular boundary tracking the experimental results have been compared with the contours delineated by cardiology specialists. Though the validation is subjective, this seems to be the best method to give a qualitative and quantitative assessment of the algorithm performance. Figure 39 shows the experimental results from the on-chip implementation together with the manually delineated contours from a US image sequence. Figure 40 shows the area comparison (Hamming metric) along with a correlation analysis (linear regression) of the hand-traced and the computed contours. This comparison allows to observe that the dynamics of the left ventricle is captured and that there is a close match between the contour areas which is at the same order as the intervariability of the human observers.[62]
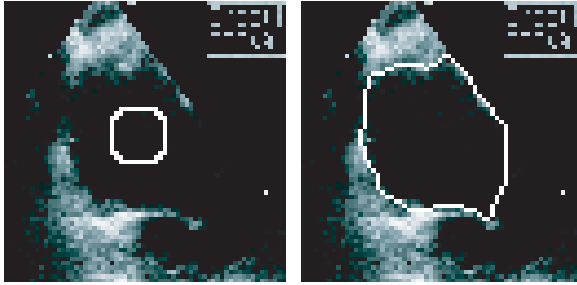
Fig. 36.   Processing on the first frame in a echocardiography image sequence. The initial contour on the left (from the previous frame) and the final contour on the right.
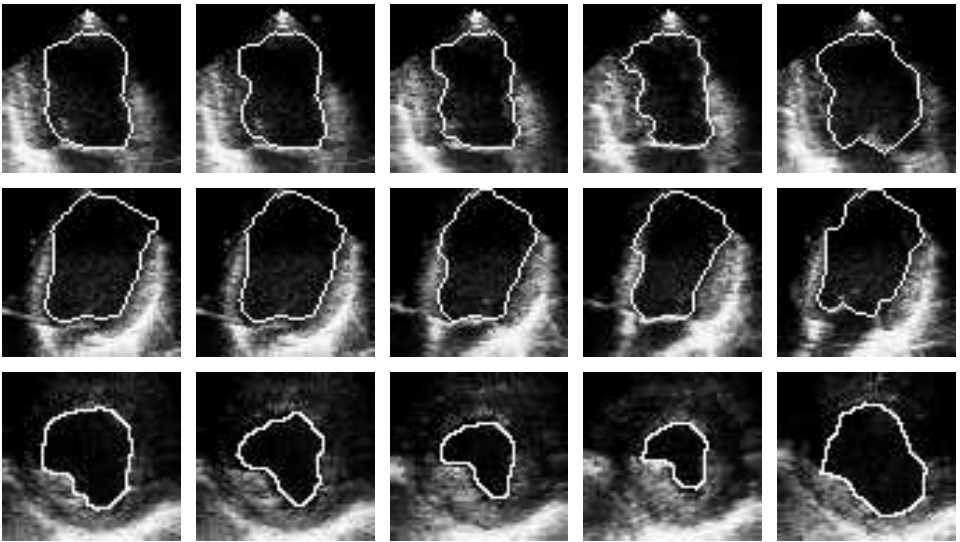
Fig. 37.   Contour tracking of a human left ventricle from four 2D ultrasound sequences corresponding to different orientation planes. (Upside) Apical two chamber view. (Middle) Apical four chamber view. (Bottom) Parasternal short axis view.

5. Evaluation of the Tracking Approach

The good performance of a particular technique of active contour is strongly dependent not only on the particular application but also the parameters which define the method. Thus, for instance, given a parametric model, the successful of the process will strongly depends on:

(1) the external potential terms considered and the weighing of these,
(2) the arrangement of the initial contours more or less close to the objective,
(3) the number of vertices of discretization and the techniques of numerical simulation (B-splines, finite elements, etc.),
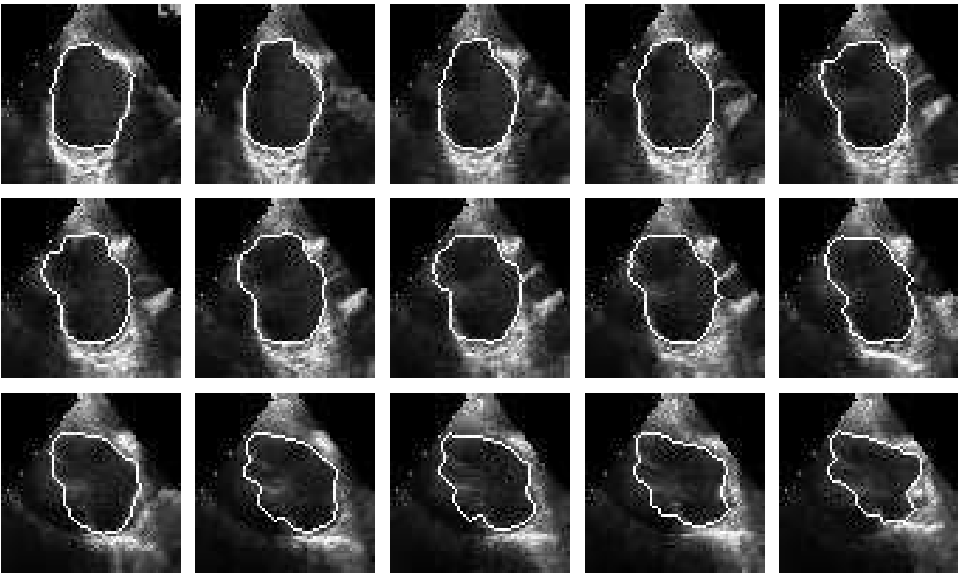
Fig. 38.   Contour tracking in 3D ultrasound echocardiography. The sequence goes from left to right and top to bottom.
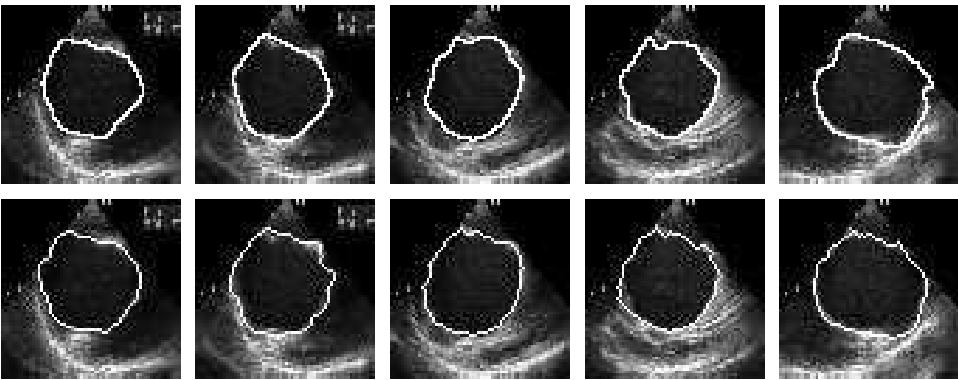
Fig. 39.   Example of contour tracking of a human left ventricle from a sequence of 3D US echocardiography images. (Upside) Contours manually delineated by cardiology specialists. (Bottom) Contours delineated with the on-chip implemented algorithm.

(4) the technique used for the energy minimization (solving the associated Lagrange equations, simulated annealing, neural networks, dynamic programming, genetic algorithms, etc.), etc.

On the other hand, even though the input data have been established, the result obtained does not justify the suitability of a model respect to another since there
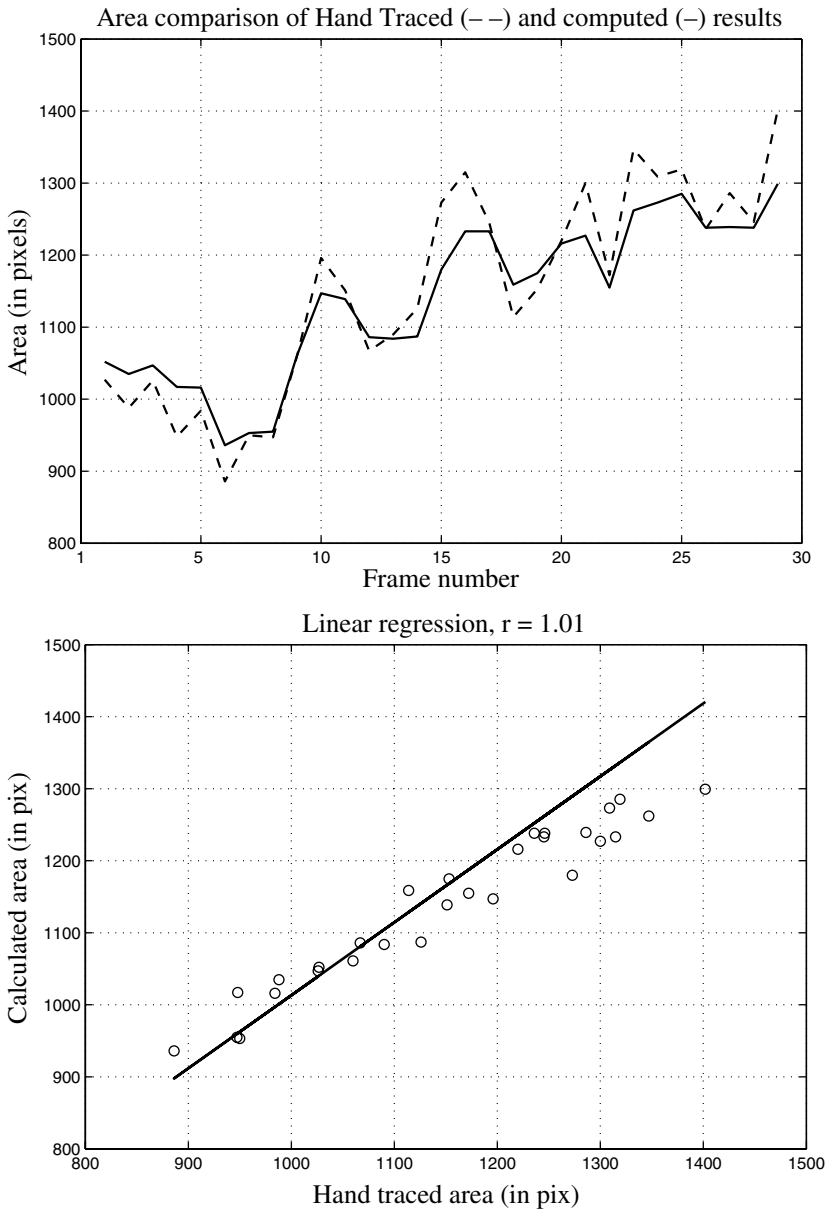
Fig. 40.   (Upside) Area Comparison of the hand-traced and estimated contours for a complete 3D US image sequence. (Bottom) Error function (the area difference relative to the hand-traced area) and the output of a correlation (linear regression) analysis.

is no condition test which can be considered as general. In fact, there will be situations where a technique will be more suitable than other due to its particular characteristics. As a consequence, we consider a quantitative comparison based on simulation with different techniques onto a determined image database has not

much relevance. However, the global characteristics of the parametric and implicit models as well as of the pixel-level snakes are in more or less degree complementary, such as a qualitative comparison of their features could help to determined the application field of each kind of model. Particularly, we will focus on characteristics as control capacity, flexibility, processing speed, precision, convergence and robustness.

5.1. *Control vs flexibility*

The parametric models allow a high control of the mobility by means of the guiding energies as well as by the establishment of a manual control onto local characteristics.[23] In addiction, they provide a better control of the contour deformation in terms of stiffness and elasticity allowing also to penalize deformations going far of *a priori* reference shape. These characteristics make the parametric models particularly suitable for tasks supported on a previous knowledge on the application domain and the accuracy of the initial model. Besides, they present low flexibility such as they cannot approach the segmentation of highly irregular objects with deep and narrow cavities and/or protuberances. Furthermore, they usually cannot handle the topologic transformations required when the number of objects into the scene is different to the number of active contours. This precludes their application where the number of the interesting regions and their approximate locations are not known in advance. There are strategies as the so called topologically-adaptable snakes (T-snakes[30]), derived from the parametric models which permit to perform topologic transformations and to reach irregular contour shapes. This is achieved by iterative processes of evaluation and reparameterization. However, since the energy estimation is strongly dependent on the parameterization, it changes in each of the intermediate reparameterization steps. This certainly makes feasible to reach tubular shapes but at the expense of a penalization of the energy estimation.

The implicit deformable models are characterized by a high flexibility in the contour evolution, allowing to reach highly irregular shapes or making topologic transformations in a natural form. However they present difficulties to impose constrains and specific control terms (propagation speed, suitable stop criteria, etc.) because of the higher dimensionality of the embedding hipersurface and the implicit definition of the contour model.

Finally, the pixel-level snakes are situated midway between parametric and implicit models. Moreover, due to their high level of discretization and the characteristics of evolution (pixel to pixel) they can handle topologic transformations by operations perfectly integrated and accessible in the evolution process. This also allows to control the topologic transformations by preventing the collision between contours in a simple way. Pixel level snakes can also delimit objects with cavities and protuberances keeping a high control of the contour shape due to the definition of deformation potential terms based on the locally measured curvature.

5.2. *Speed vs precision*

Among the techniques based on numerical simulations there are mainly two factors which determine the time required for an active contour to reach stability. One of this is the number of steps needed to reach the final solution, strongly dependent on the initial model. The another factor is the number of vertices on which the model is discretizated (or the embedding surface in the implicit formulations) and which determine to a great extent the precision of the output. A number of approaches reduce the time of evolution at the expense either of limiting the space of variation of the contours or the number of vertices of the discretization, which lead to a lack of precision.

Concerning the parametric models, the computing time must be also multiplied by the number of contour models being considered. Computing time into the implicit formulations is independent on the number of the involved contours, however the higher complexity of these procedures, leads to search alternatives which also constrain their precision.[53]

The pixel level snakes are discretiz to pixel level and can be projected on massively parallel SIMD architectures. This allows to associate each pixel of the contour image to one single processing element which interacts with those associated to the neighbouring pixels. Thus the computational cost is detached from the number of point of discretization of the domain. However, the size of the images to be processed is linked to the capability of integration of processing elements into a single chip.

5.3. *Convergence vs robustness*

The convergence of an active contour towards stable location and shape must be guaranteed. To this end, in the conventional active contour techniques the discretization level in both space (image domain), and time (iteration step), must be suitable to avoid oscillations. On one hand, the finer the discretization, the more guarantees of convergence, but at expenses of higher sensitivity to local minima. On the other hand, the higher the iteration step (or the rougher the discretization of the spatial variable) the lesser possibilities to be anchored in a local minimum, but at expenses of a higher instability of the final solution.

Concerning the pixel-level snakes the spatial and temporal variables are fixed to the minimum level which guarantees a stable final solution. However this makes the pixel-level snakes more sensitive to the local minima and then more dependent on the initiation and the guiding information. This sensitivity is partially attenuated by a higher range of movement and the non-parametric nature of the model in such a way that the number of evolution points is not predetermined.

6. Conclusions

Active contours are widely acknowledged techniques for the segmentation and tracking on a variety of medical images. They can integrate knowledge from different

sources (shape and appearance models, statistics about region and boundary features, motion estimation, etc.), however, the inter- and intra-patient variability, and the different imaging modalities impede the arriving of a unique general method to tackle all these problems. Each approach has its strengths and weaknesses, but the synergetic integration of different abilities is still challenging. Some image analysis tasks need a good delineation of object contours, while others have the main goal of fast following of objects through image frames. Therefore, active contours should be customized for each domain tasks.

The two active contours we have described here, were designed with different goals in mind. The first active contour technique is devoted to tackle segmentation tasks where the desired features of the object of interest could slightly change from slice to slice, while the second active contour technique is dedicated to tracking tasks where the features of the object of interest do not change from frame to frame, but the shape does. The segmentation approach aimed to learn and update the best discriminants from one slice to the next. The tracking approach, however, assumes well established goal features and intends to reach them as fast as possible. Different applications should demand different abilities.

Acknowledgments

References

1. D. Adalsteinsson and J. Sethian, *J. Comput. Phys.* **118** (1995) 269.
2. K. T. Bae, M. L. Giger, C.-T. Chen and C. E. Kahn, Jr, *Med. Phys.* **20** (1993) 71.
3. R. Bajcsy and S. Kovacic, *Comput. Vis. Graph. Image Process.* **46** (1989) 1.
4. A. Blake and M. Isard, *Active Contours* (Springer-Verlag, 1998).
5. V. Caselles, F. Catte and F. Dibos, *Numer. Math.* **66** (1993) 1.
6. A. Chakraborty, L. H. Staib and J. S. Duncan, *IEEE Trans. Med. Imaging.* **15** (1996) 859.
7. V. Chalana, D. T. Linker, D. R. Haynor and K. Yongmin, *IEEE Trans. Med. Imaging.* **15** (1996) 290.
8. K. Cho and P. Meer, *Comput. Vis. Image Und.* **68** (1997) 72.
9. C.-C. Chu and J. K. Aggarwal, *IEEE Trans. Patt. Anal. Machine Intell.* **15** (1993) 1241.
10. L. O. Chua and L. Yang, *IEEE Trans. Circuits Syst.* **35** (1998) 1257.
11. L. D. Cohen and I. Cohen, *IEEE Trans. Patt. Anal. Machine Intell.* **15** (1993) 1131.
12. T. F. Cootes, C. J. Taylor, D. H. Cooper and J. Graham, *Comput. Vis. Image Und.* **61** (1995) 38.
13. A. P. Dhawan and L. Arata, *IEEE Eng. Med. Biol. Mag.* **10** (1991) 30.
14. R. O. Duda and P. E. Hurt, *Pattern Clasification and Scene Analysis* (Wiley-Interscience, New York, 1973).
15. J. H. Elder and S. W. Zucker, *IEEE Trans. Patt. Anal. Machine. Intell.* **20** (1998) 699.

16. L. M. J. Florack, B. M. ter Haar Romeny, J. J. Koenderink and M. A. Viergever, *Image Vis. Comput.* **10** (1992) 376.
17. L. M. J. Florack, B. M. ter Haar Romeny, M. A. Viergever and J. J. Koenderink, *Int. J. Comput. Vis.* **18** (1996) 61.
18. W. T. Freeman and E. H. Adelson, *IEEE Trans. Patt. Anal. Machine. Intell.* **13** (1991) 891.
19. A. Goshtasby, D. A. Turner and L. V. Ackerman, *IEEE Trans. Med. Imaging.* **11** (1992) 507.
20. X. Han, C. Xu and J. L. Prince, *IEEE Trans. Patt. Anal. Machine Intell.* **25** (2003) 755.
21. J. Ivins and J. Porrill, *Image Vis. Comput.* **13** (1995) 431.
22. A. K. Jain, N. K. Ratha and S. Lakshmanan, *Pattern Recognit.* **30** (1997) 295.
23. M. Kass, A. Witkin and D. Terzopoulos, *Int. J. Comput. Vis.* **1** (1988) 321.
24. J. J. Koenderink, *Biol. Cybern.* **50** (1984) 363.
25. A. Lanitis, C. J. Taylor and T. F. Cootes, *Image Vis. Comput.* **13** (1995) 393.
26. M. E. Leventon, O. Faugeras, W. E. L. Grimson and W. M. Wells, in *Proc. IEEE Workshop on Mathematical Methods in Biomedical Image Analysis* (2000), p. 4.
27. G. Liñán, S. Espejo, R. Dominguez-Castro and Rodriguez-Vazquez, *Int. J. Circuit Theory Appl.* **30** (2002) 89.
28. R. Malladi, J. A. Sethian and B. C. Vemuri, *IEEE Trans. Patt. Anal. Mach. Intell.* **17** (1995) 158.
29. J. Malik and P. Perona, *J. Opt. Soc. Am.* **7** (1990) 923.
30. T. McInerney and D. Terzopoulos, in *Proc. Fifth International Conference on Computer Vision* (1995), p. 840.
31. T. McInerney and D. Terzopoulos, *Med. Image Anal.* **1** (1996) 91.
32. S. Menet, P. Saint-Maric and G. Medioni, in *Proc. DARPA Image Understanding Workshop* (1990), p. 720.
33. I. Mikic, S. Krucinski and J. D. Thomas, *IEEE Trans. Med. Imaging* **17** (1998) 274.
34. C. Nastar and N. Ayache, *IEEE Trans. Patt. Anal. Machine. Intell.* **18** (1996) 1067.
35. W. Neuenschwander, P. Fua, G. Szekeli and O. Kubler, *Comput. Vis. Image Und.* **65** (1997) 237.
36. W. J. Niessen, J. S. Duncan, M. Nielsen, L. M. J. Florack, B. M. ter Haar Romeny and M. A. Viergever, *Comput. Vis. Image Und.* **65** (1997) 259.
37. W. J. Niessen, B. M. ter Haar Romeny and M. A. Viergever, *IEEE Trans. Med. Imaging.* **17** (1998) 634.
38. S. Pankanti and A. K. Jain, *IEEE Trans. Patt. Anal. Machine. Intell.* **17** (1995) 831.
39. N. Paragios and R. Deriche, in *Proc. on Computer Vision and Pattern Recognition*, Vol. 2 (Colorado, 1999), p. 300.
40. N. Paragios and R. Deriche, *IEEE Trans. Patt. Anal. Machine Intell.* **22** (2000) 266.
41. X. M. Pardo, D. Cabello and J. Heras, *J. Comput. Inf. Technol.* **7** (1999) 295.
42. X. M. Pardo, M. J. Carreira, A. Mosquera and D. Cabello, *Image Vis. Comput.* **19** (2001) 461.
43. X. M. Pardo, P. Radeva and D. Cabello, *Med. Image Anal.* **7** (2003) 293.
44. A. Paasio, A. Kananen, K. Halonen and V. Porra, *J. VLSI Signal Process. Syst.* **23** (1999) 281.
45. T. Pavlidis and Y. -T. Liow, *IEEE Trans. Patt. Anal. Machine Intell.* **12** (1990) 225.
46. R. P. N. Rao and D. H. Ballard, in *Proc. Int. Joint Conference on Artificial Intelligence* (1995), p. 10.
47. S. P. Raya and J. K. Udupa, *IEEE Trans. Patt. Anal. Machine Intell.* **9** (1990) 32.
48. R. Ronfard, *Int. J. Comput. Vis.* **13** (1994) 229.

49. T. Roska and L. O. Chua, *IEEE Trans. Circuits Syst.* **40** (1993) 163.
50. G. I. Sanchez-Ortiz, G. J. T. Wright, N. Clarke, J. Declerck, A. P. Banning and J. A. Noble, *IEEE Trans. Med. Imaging* **21** (2002) 1069.
51. S. Sandor and R. Leahy, *IEEE Trans. Med. Imaging* **16** (1997) 41.
52. J. Sethian, in *Proc. National Academy of Sciences* (1996), p. 1591.
53. J. A. Sethian, *SIAM Review* **41** (1999) 199.
54. M. Sonka, W. Park and E. A. Hoffman, *IEEE Trans. Med. Imaging.* **15** (1996) 314.
55. L. H. Staib and J. S. Duncan, *IEEE Trans. Patt. Anal. Machine. Intell.* **14** (1992) 1061.
56. L. H. Staib, A. Chakraborty and J. S. Duncan, *Int. J. Patt. Recogni. Artif. Intell.* **11** (1997) 1247.
57. J. B. Subirana-Vilanova and K. K. Sung, *Technical Report AIM-1318* (Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 1992).
58. H. Tek and B. B. Kimia, *Comput. Vis. Image Und.* **65** (1997) 246.
59. M. Turk and A. Pentland, *J. Cogn. Neurosci.* **3** (1991) 71.
60. G. Turk and J. F. O'Brien, in *Proc. ACM SIGGRAPH99* (Los Angeles, California, 1999), p. 335.
61. D. L. Vilariño, D. Cabello, X. M. Pardo and V. M. Brea, *Image Vis. Comput.* **21** (2003) 189.
62. D. L. Vilariño and Cs. Rekeczky, *IEEE Trans. Circuits Syst.* **51** (2004) 885.
63. C. Xu and J. L. Prince, *IEEE Trans. Image Process.* **7** (1998) 359.
64. J. Yang and J. -Y. Yang, *Pattern Recognit.* **36** (2003) 563.
65. A. Yezzi, A. Tsai and A. S. Willsky, in *Proc. Int. Conference on Computer Vision*, Vol. 1 (Greece, 1999), p. 898.
66. R. A. Young, in *Proc. IEEE Conference on Computer Vision and Pattern Recognition* (Miami, FL, 1986), p 564.
67. A. L. Yuille, P. W. Hallinan and D. S. Cohen, *Int. J. Comput. Vis.* **8** (1992).
68. Y. J. Zhang, *Pattern Recognit. Lett.* **18** (1997) 963.
69. L. Zhou, M. Rzeszotarski, L. Singerman and J. Chokreff, *IEEE Trans. Med. Imaging* **13** (1994) 619.
70. D. Zhu, R. Conners and P. Araman, in *SPIE, Visual Communications and Image Processing*, Vol. 1606 (1991).
71. S. Ch. Zhu, T. S. Lee and A. L. Yuille, in *Proc. International Conference on Computer Vision* (1995).
72. Y. Zhu and H. Yan, *IEEE Trans. Med. Imaging.* **16** (1997) 55.

CHAPTER 8

FUNCTIONAL MRI ACTIVITY CHARACTERIZATION: AN ESTIMATION AND DECISION THEORETIC APPROACH

MUKUND DESAI* and RAMI MANGOUBI[†]

C.S. Draper Laboratory
Cambridge Massachusetts, USA
*mdesai@draper.com
[†]rmangoubi@draper.com

HOMER PIEN

Massachusetts General Hospital
Boston, Massachusetts, USA
hpien@nmr.mgh.harvard.edu

Characterizing the response of the brain to a stimulus based on functional MRI data is a major challenge due to the fact that the response time delay or shift of the brain may be different from one stimulus phase to the next and from voxel to voxel. To enhance detectability, this work introduces the use of a curve evolution approach that provides separate estimates of the response time shifts at each phase of the stimulus on a voxel-by-voxel basis. The approach relies on a parsimonious but simple model that is nonlinear in the time shifts of the response relative to the stimulus, and linear in the gains. To effectively use the response time shift estimates in a subspace detection framework, we implement a robust hypothesis test based on a Laplacian noise model. The algorithm provides a voxel-by-voxel functional characterization of the brain's response. The results based on experimental data show that response time shift estimates, when properly implemented, enhance detectability without sacrificing robustness to false alarms or false positives. The chapter also discusses the implications of Laplacian noise model. One such implication is that the simple response model adopted is optimal in the context of a subspace learning problem formulation. Another implication relates to the simple form of the Laplacian detector that enables analysis on the design aspects of stimulus signal.

Keywords: Functional MRI data; robust detection; Laplacian noise; learning stimulus; design of stimulus.

1. Introduction

Since its introduction in the early 1990's,[a] functional magnetic resonance imaging (fMRI) has become an invaluable tool to the neuroscience community by permitting the non-invasive analysis of many perceptual, behavioral, and cognitive activities in the human brain.[5,70,82–85] Using correlates between experimental stimuli and the resulting observed functional responses, these studies have provided considerable insight into purported physical locations of processing centers, propagation of neural

---

[a]Portions of this chapter are based on previously published work,[23,25,27,28] with permission, ©IEEE, 2002–2004.

activity, and the effects of neurodegenerative diseases, to name just a few. Because of the importance of establishing correlations between stimulus and response, the purpose of this paper is to provide a statistical framework within which such correlations may be established in a robust fashion.

Of primary interest to this paper are the detection, estimation, and localization of human cerebral hemodynamic response to repetitive and well-controlled stimuli using blood oxygenation level dependent (BOLD) MRI signals. In this context, the scope of this chapter excludes the related notions of perfusion and diffusion MR imaging.[29,76,105]

The challenges in rigorously characterizing fMRI response signals are significant due to several factors. First, the signal changes are very low in magnitude, typically a few percent of the baseline signal. Second, there is an unknown time delay between the presentation of the stimulus and the observed response. Third, adaptation may occur over time, altering the baseline of the response. Fourth, the waning focus of the subject over time also alters the nature of the observed response. Fifth, neural signals are transient, propagating from loci to loci, thus the temporal nature of the imaging process plays an important role.

In addition to the above challenges, one issue that has not been addressed in past research is that of *unknown* interference, meaning signals that are of unknown origin or whose characteristics have not been determined. It is naturally preferable to learn the interference characteristics so as to determine its presence and account for it at the detection stage. Learning, however, may not always be an option, as for instance sufficient data may not be available. The work in this chapter addresses this issue as well; in fact as we will state later in this section, unknown interference is one issue that has motivated the development by Desai and Mangoubi[25–28] of the new robust or game theoretic detection methodology applied to fMRI in this chapter.

The way an experimenter deals with all these factors is to first make appropriate assumptions about the characteristics and properties of these confounding factors, and utilize a model to capture their essence. Thus the choice of a model within which the BOLD signal is to be processed and interpreted is an important issue, but is one that is often done on an ad hoc basis, or is done implicitly.[96] These models make assumptions about the stimulus, the response, the noise, and spatio-temporal correlation characteristics, the interference, and the hemodynamic delays.

In this chapter, we seek to characterize the deviation from average over a neighborhood response so as to capture the finer variations, both in time and from voxel to voxel. We do so in a *hierarchical* fashion. In the first step we learn the transient characteristics at each voxel that exploits both the spatial and temporal correlations of the transient. Next, in the second step, we determine the steady state response characteristics voxel by voxel, and a decision is made as to whether each voxel responds to the stimulus or not. The transient and steady state characteristics are captured by a parsimonious yet simple response model. This model is compatible with the Laplacian noise model selected for robust detection, which is executed using a newly developed methodology that is insensitive to unknown interference.

We now state our approach towards each of these steps.

1.1. *Objectives*

The chapter presents an estimation and decision theoretic formulation of the fMRI detection problem that has two objectives:

- *Sensitivity to the response signal*, in particular sensitivity at voxels where the response's magnitude is small, and simultaneously,
- *Robustness or insensitivity to learned and unlearned interference*, meaning interference or structured noise that may be learned, such as cardiac effects, as well as interference that has not been learned either for lack of data or for other reasons. Of importance here are errors due to wrong models or unmodeled effects.

As mentioned previously, these objectives have motivated the development of new methodologies for robust or game theoretic signal detection, learning, and estimation. This methodology, developed by Desai and Mangoubi[25–28] for addressing the challenges in fMRI, is more general than the statistical detection and subspace learning methodologies that are based on the Gaussian noise assumption. In particular, the subspace detectors[25,28] applied to fMRI[23] and described in this chapter, generalize the conventional Gaussian based ones in, for instance Refs. 98, 99.

The fMRI application in this chapter is based on the work in Ref. 23. The results described in Sec. 8 compare the new methodology to the Gaussian based ones. A peek at Fig. 4 shows that *the new methods detect more responding voxels, while simultaneously producing less false alarms.*

Although it is always dangerous to segregate the computational from the cognitive or physiological problem, comprehensive treatments of the physics of measurements, the various explanations of physiological phenomena, and the clinical application of fMRI, are beyond the scope of this chapter. The physics and physiology of fMRI measurements are explained in Refs. 64, 80, 79.

It is our view, however, that an understanding of fMRI applications is a prerequisite for any useful work in the signal processing aspects of that field. In fact, the subtleties and challenges found in the various applications motivated the novel aspects of the methodologies we present. We describe these applications in Sec. 2, where we attempt to provide a context for understanding important processing issues. Specifically, we sketch a breadth of issues that are addressed and the BOLD properties that are sought, rather than pursuing an exhaustive review of applications. We also provide a large number of references for the application of fMRI, many of which have guided our thinking. But first we outline the key aspects of the chapter's technical approach.

1.2. *Key analytical aspects of hierarchical approach*

The novel features of this methodology have theoretical implications in statistical detection theory, and, beyond fMRI, in many cases hold promises for other applications such as radar, sonar, biochemical sensing, etc. Some of these features are given below.

1.2.1. *Models and subspace learning*

We provide a simple yet parsimonious model that describes the fMRI stimulus response. This model, which assumes that the response resides in a *binary subspace*, meaning a subspace spanned by a vector consisting of binary elements, or +1 and −1, is described in Sec. 4.2, and has several advantages:

- It permits the estimation of the response magnitude, and more importantly, the *response time shift at each individual voxel and at each phase of the stimulus*. The gains capture the steady state behavior of the response, while the shifts capture the transient behavior. Voxel wise estimation of gain and time shift parameters enhances the sensitivity of the detection, particularly at voxels where the response may not be as strong.
- Though simple, the model is under some noise assumptions, specifically under the Laplacian noise model assumption, the *optimal* choice in the sense that it is the optimal solution to the subspace learning problem where no parametric dependency constraints are explicitly imposed. This is described in Sec. 9.1. The Laplacian noise model has a fatter tail, and provides more robustness.
- The model has implications for fMRI *signal design* as well, since the distributions of the binary elements can be easily selected so as not only to *enhance detectability*, but also help minimize the impact of interference with partially known spectral characteristics; an important concern when dealing with small signal to noise and interference ratio (SNIR) applications such as fMRI. The signal design implications are described in Sec. 9.2.

1.2.2. *Parametric estimation*

Our curve evolution variational approach to the response time-shift parameter estimation (Sec. 6), which generalizes the work in Ref. 100, accounts for the temporal and spatial continuity, meaning the correlation both in time and across voxels, of the shift parameter. The result is an enhancement in sensitivity, or in our case the ability to more accurately estimate the shift parameters. It also allows for the adoption of cost terms that lend themselves to a Laplacian stochastic noise model interpretation.

1.2.3. *Robust non-Gaussian signal detection*

The newly developed robust non-Gaussian detection methodology, which is a minmax or game theoretic generalization of the Gaussian noise based χ^2, t, and F subspace detectors of Ref. 98, 99, is described in Sec. 7 for the fMRI problem.[25,28] Its advantages are:

- It is valid for a large class of non-Gaussian detection problems. Thus, by changing one parameter, the designer has more choice of noise modeling flexibility.

- For the Generalized Gaussian model, it provides easily implementable expressions for the detectors. The Generalized Gaussian model embraces a large family of density functions, including the Gaussian, Laplacian, and Uniform.
- It is robust not only to interference signals that are known, but also those that are totally unknown and have not been learned or are difficult to learn prior to conducting the experiment. Again, this reduces the false alarm rate in fMRI.

1.2.4. *The role of the Laplacian noise density function*

From the above, we see that the *Laplacian noise model* plays an important role in our approach. Specifically, (1) it is compatible with the binary element subspace model selected in the sense that it simplifies the learning process[26,27] (See Sec. 9.1), (2) the curve evolution approach for shift parameter estimation involves Laplacian associated cost terms (1-norm), and (3) forms the basis of the robust detectors of Sec. 7 which, as for instance Fig. 4 shows, provide enhanced detection and reduced false alarm performance when compared to conventional Gaussian based detectors.

1.3. *Organization of the Chapter*

Section 2 discusses the applications of fMRI. In Sec. 3, we provide an overview of our approach, while in Sec. 4 we describe more formally our approach to the fMRI modeling, estimation, and detection problems. The same section presents in details the stimulus response model. Section 5 explains the logic behind some aspects of our approach. The model time shift parameter estimation problem is formulated and solved in Sec. 6. A robust detection paradigm is formulated and explored in Sec. 7, and results are provided in Sec. 8. Further stimulus signal design, theoretical generalizations and implications are given in Sec. 9. Finally, Sec. 10 offers some conclusions.

2. Applications of Functional MRI

Before delving into these applications, it is worthwhile to first consider what exactly is being measured by the fMRI signal. One significant assumption underlying the use of BOLD is that, upon presentation of a stimulus, the subsequent changes observed in the MRI signal of the brain can be uniquely attributed to that stimulus. This signal is measured by observing the differences between a brain at rest and a stimulated brain undergoing some activity which is attributed to that stimulus. In this context, activity is characterized by increased metabolism and the increased demand for oxygen. The control and elimination of extraneous stimuli so as to yield a statistically significant and reproducible BOLD signal is a critical component of any fMRI experimental design.

To observe BOLD fMRI, T2-weighted images are generally used. This is because while oxygenated blood (i.e. oxyhemoglobin) is diamagnetic (essentially

non-magnetic), deoxyhemoglobin is paramagnetic, thus giving rise to shorter T2-signals, hence deoxyhemoglobin appears darker in T2* MR images.[b] In response to a stimulus, there is a corresponding increase in the amount of oxygenated hemoglobin within the microvasculature, which displaces the amount of deoxyhemoglobin and alters the ratio of oxy- to deoxy-hemoglobin. As such, the observed fMRI signal is not the direct result of increased inflow, but instead is an indirect result of the decrease in the amount of deoxyhemoglobin that gives rise to dark MRI signals. Yet the displacement of deoxyhemoglobin is only the result of chain of events: a cascade of biochemicals are released in response to stimulation, leading to increases in blood flow, blood volume, and oxygen consumption.[14,32,34,39] As such, the observed fMRI signal represents a complex choreography of a multitude of molecular and physiological events.

Functional MRI has been widely applied to a number of clinical areas. The purpose of this section is to convey the breadth of these applications without being exhaustive. In broad terms, the application areas can be divided into: (i) mapping of regions of the brain responsible for perception in Sec. 2.1, (ii) understanding the complex neuroanatomy responsible for cognition and emotion in Sec. 2.2, and (iii) elucidating mechanisms of neurological degeneration and disorders in Sec. 2.3. Since our intent is to highlight the applications of fMRI without devoting too much attention to neuroanatomy, those desiring a more extensive treatment of neuroanatomy can refer to, for instance, Refs. 64, 80, 113.

2.1. *Perception*

Many of the earliest human fMRI studies focused on mapping the regions of the brain responsible for perception, most commonly of the visual system.[83,70] These early studies confirmed the role that the primary visual cortex (V1 region) plays during photic stimulation experiments. The V1 region, located at the posterior pole of the occipital lobe, and the nearby V2 and V3 regions, were shown to behave retinotopically via fMRI experiments on patients with defects or lesions in these regions.[60] Other studies have shown specialized regions for color processing,[49] facial image processing,[52,89] and directional sensitivity.[17] Abnormalities in visual information processing have also been implicated in dyslexia patients, where fMRI played a crucial role in identifying differences in visual-stimulation brain activation patterns when compared to normal subjects.[20,36,48,55] Others have studied migraines and the visual illusion known as aura that precede migraines in approximately 20% of migraine sufferers, and confirmed the characteristic retinotopic progression of migraine aura using fMRI.[48] Collectively, these studies point to the retinotopic, distributed, and specialized nature of human visual information processing. Comprehensive and critical reviews of human visual information processing can be found in Refs. 55, 65, 112.

---

[b]Readers desiring a more fundamental introduction to MRIs can refer to, for example Refs. 64, 79.

Functional MRI has also been used to examine other forms of sensory perception. A series of studies has been undertaken to compare differences in regions of brain activation between normal and hyposmia (those with a reduced sense of smell) subjects, as well as hyposmia subjects treated with theophyline for six months.[74,75] These olfactory perception studies confirmed the presence of two types of hyposmia, and demonstrated the role of the frontal cortex in directing olfaction processing to other areas of the brain, such as the temporal and cingulated regions. These studies additionally examined the neuroanatomical relationship between odor recognition and odor memory. The same group of researchers also looked at the sense of taste, and concluded that taste and smell are localized in different parts of the brain.[56] A comprehensive overview of the neuroanatomical structure and function of the olfactory system, as derived from different imaging modalities and analysis techniques, can be found in Ref. 67.

The identification of the regions of the brain responsible for pain is a complex task. One study utilized noxious hot (46-deg C) and cold (5-deg C) stimulation to the dorsum (backside) of the left hand, and identified numerous structures of the brain that were activated, including mid-brain regions thalamus, basal ganglia, and insula, as well as cortical regions cingulate, somatosensory, premotor/motor cortices, prefrontal and inferior parietal cortex.[110] The same study also found that most regions were activated bilaterally, but generally with stronger activation contralateral to the stimulus.[110] The involvement of the somatosensory cortices SI and SII has been controversial, due to inconsistent results. In an attempt to reconcile some of these inconsistencies, one study identified differences in activation patterns between painful heat and non-painful tactile (brush) stimuli.[16] This study found that both groups of subjects activated the SI and SII regions, but the activation showed different peak patterns and time constants. In particular, brush-evoked stimuli showed a consistent, single-peak response at 10-sec, while the heat-evoked response showed two peaks, with the dominant second peak occurring at 17-sec after stimulation,[16] and appears to confirm the persistent nature of the heat-induced pain reported by many subjects. These studies point to the complex nature of our response to pain, and raise many questions about how such experiments control for confounding variables.[19]

2.2. *Cognition*

Language also involves complex interplay among various regions of the brain. The localization of language centers is well established in the history of science — in the mid 1800's, Paul Broca encountered a patient who was able to understand but unable to form speech, subsequent to the patient's death Broca identified a lesion in the inferior frontal gyrus and established the central role of that region to speech. Other scientists of the time, such as Carl Wernicke, similarly found another localized area of the brain responsible for language by noting the relationship between damaged posterior superior temporal region and speech comprehension.[47] The advent

of fMRI has enabled much finer delineations of functional subregions within Broca's region, Wernicke's region, as well as other areas.

In one study, differences between overt and covert speech (i.e. between mentally thinking of a word versus verbalizing a word) were examined. This study found that although similar regions were activated (with the exception that overt speech involved various motor regions), regions activated by overt speech yielded higher magnitude responses.[86] In another study, the encoding of semantic information was investigated.[22] This study found that specific regions within the left inferior prefrontal cortex showed increased activity during semantic encoding as compared to non-semantic encoding. Furthermore, the same study established the "priming" effect — repeated semantic encodings showed decreased activation. Yet another aspect of language is the characterization of semantic versus syntactic information processing. By using syntactically correct but semantically meaningless sentences (so called "Jabberwocky" sentences), one study found that some areas of the brain (such as area 44) are active regardless of the semantic or syntactic nature of the stimuli, whereas other areas (such as temporal lobes) showed differential activation depending on whether the task at hand was a semantic or a syntactic one.[41]

While fMRI has been useful in studying different aspects of language, the complex nature of language and language processing has necessitated increasingly sophisticated experimental paradigms.[8] The same can also be said of a related aspect of cognition — memory. In one study examining activation differences between deep- and shallow-encodings (where shallow-encodings result in low levels of recognition success but high retrieval effort and deep-encodings result in high recognition success and low retrieval effort), many common areas of activation between deep- and shallow-encodings were found. The study also found that bilateral anterior insular regions and a left dorsal prefrontal region were more active after shallow-encoding, while right anterior prefrontal cortex was most active during successful retrieval after deep-encoding.[11] Another study found the that posterior medial temporal lobe is associated with episodic encoding, although inconsistencies between fMRI and positron emission tomography (PET) were found.[97] In yet another study designed to examine familiarity during retrieval, subjects are asked whether they (a) recollect seeing a word, (b) whether they experience familiarity in the absence of recollection, or (c) whether they did not remember seeing it previously. Cases (a) and (b) showed enhanced responses in left prefrontal and left parietal cortices, while the opposite pattern was observed in bilateral temporoccipital regions and amygdala. Furthermore, case (a) showed enhanced response in anterior left prefrontal, left parietal, and posterior cingulated regions.[57] (The interested reader can refer to Ref. 12 for a comprehensive review of the neuroscience surrounding memory.) Collectively, these results point to the contribution of fMRI to understanding the encoding and retrieval of memory, and the localized and distributed nature of cognitive processing within the human brain.

Just as language and memory have benefited from the objectivity provided by fMRI studies in the past decade, so has the study of emotions evolved from

subjective attributes to objective neuroanatomy. For example, fear is associated with the amygdala,[50,73] although others have questioned the specificity of such responses, as well as the role other regions of the brain, such as the insular region, play in fear.[116] Using patients with focal brain damage, one study mapped the regions of the brain responsible for recognition of faces containing emotions.[1] The role of the retrosplenial cortex in responding to emotionally salient stimuli, as well as its involvement in episodic memory, has also been examined.[77] The brain activation patterns of eight obsessive-compulsive disorder subjects with contamination preoccupation were examined, and one study found disgust-inducing tasks activated the right insula and parahippocampal regions.[101] The emotional "high" resulting from a reward has also been examined, both using drugs such as cocaine[9] and using monetary rewards.[10] Recently, a meta-analysis of 55 PET and fMRI studies was performed; this study categorized emotions into five states (happy, fear, anger, sadness, disgust) and examined the level of agreement among them as to the regions of the brain responsible for these emotions.[87] This meta-analysis provides consensus on the regions that have been strongly correlated with each of these five emotional states.

2.3. *Neurodegeneration and disorders*

In addition to advancing our knowledge of neuroanatomical structures responsible for various perceptive and cognitive functions, a significant application of fMRI is to aid in the characterization of neurodegneration, as well as diagnosis of various neurological disorders and mental illnesses.

When examining neurodegeneration in memory or cognitive tasks, it is important to distinguish such changes from normal age-related alterations. It has been shown that neither the BOLD signal shape nor the hemodynamic response functions are affected by age. However, older subjects show significantly reduced signal-to-noise ratio in the BOLD signal when compared to younger subjects, using fMRI images of the primary sensorimotor cortex during a simple reaction time task.[33] Another study found age-related correlations in both signal intensity and the number of activated voxels within the sensorimotor cortex during a finger-tapping experiment.[58] In yet another study, the rise time of the BOLD signal was found to correlate with age.[108] (See Ref. 32 for a review of age-related fMRI studies.) These studies established the baseline from which comparisons can be made between neurodegeneration and aging. For example, in a face-name association experiment, fMRI activities in the hippocampal formation, medial parietal, and posterior cingulated regions were found to be significantly different between Alzheimer's and normal elderly subjects.[106] Interestingly, in another experiment in which drugs are used to induce memory impairment, drugs with purported different mechanisms of action yielded very similar fMRI responses. Specifically, the drug lorazepam is believed to bind to γ-aminobutyric acid (GABA) receptors and impair memory by enhancing the inhibitory influence of GABAergic neurons, while the drug

scopolamine is a potent antagonist of the muscarinic acetylcholine receptor and believed to impair memory by blocking cholinergic transmission.[107] Thus while fMRI is sensitive to memory impairments and neurodegeneration effects, it does not appear to be sensitive to the mechanisms underlying the impairment.

The importance of characterizing normality, and subsequently deviations from normality, is highlighted in the study of attention-deficit hyperactivity disorder (ADHD), since ADHD is considered a learning disability for school-age children, and imaging studies are performed on their still-developing brains. One MRI study found numerous imaging anomalies associated with pediatric ADHD subjects, including reduced volumes of the total brain, superior prefrontal, right superior prefrontal, and other regions.[59] Functionally, ADHD has been associated with dysfunction in arousal, behavioral inhibition, and attention systems, perhaps as a result of structural abnormalities in frontostriatal regions of the brain.[35,51] Schizophrenia has similarly been linked to both structural and functional deviations — hippocampal volume reductions has consistently been associated with schizophrenia subjects, and to a lesser degree in first-degree relatives of schizophrenic probands.[54] Functionally, it has been demonstrated that abnormal levels of hippocampal activity occur at rest, during auditory hallucinations, and during memory retrieval tasks.[54] While functional neuroimaging has contributed to the understanding of schizophrenia, this mental disorder is highly complex in nature, with potential genetic and environmental components, as well as many non-specific biomolecular markers.[88] It will likely be some time before the mechanism of this disorder is fully understood.

It has been hypothesized that interpersonal and social interactions require the recognition of faces.[111] As such, it is not surprising that people with autism may show differences in face-recognition processing when compared to normals.[37] This has been shown experimentally with fMRI. Working with high-functioning autism and Asperger syndrome subjects, one study found that face recognition experiments significantly greater activation of the right inferior temporal gyri and less activation of the right fusiform gyrus as compared to the control group.[102]

2.4. *Summary of applications*

In this section we have taken a cursory view of the numerous and disparate ways in which fMRI is utilized to address compelling scientific and clinical questions. It should be apparent that while significant differences may exist among experimental paradigms or image acquisition methods, the fundamental fMRI parameters researchers utilize to characterize regional activation and abnormalities remain similar: statistical significance of the activated voxels, the regions of the brain associated with these voxels, the hemodynamic delays associated with activation, and the magnitude of the activation. In the context of the methodology introduced in this paper, we have attempted to capture these parameters in a rigorous way, so as to enable the detection of even smaller changes with higher statistical significance.

3. Approaches

A major analytical challenge with the analysis of fMRI data stems from the fact that the observable changes in signal intensity due to changes in the concentration of deoxyhemoglobin are noisy and small; only a few percentage points. As such, considerable research efforts deal with the statistical analysis of fMRI signals collected during experiments when the brain is subjected to stimuli. The objective is to determine whether temporal signal changes at a voxel are due to noise or an indication that the voxel is responding to the stimulus. One of the earlier efforts is that of Bandetini *et al.*,[4] who used cross-correlation between the measured time series, after estimation and removal of effects other than the stimulus, and estimates of the expected response signal. Cox *et al.*[18] provide statistics such as those in Ref. 4, but recursively computed for real time implementation, and examine the threshold selection issue. In Ref. 69, Kim and Ugurbil formulate parametric Student t and non-parametric Kolmogorov–Smirnov hypothesis tests to classify voxels into two classes, responding and not responding. In Refs. 42, 43, Friston *et al.* use time-series methods based on generalized linear models (GLM) for the same purpose, including estimation of hemodynamic response. In a similar vein in other early and continuing work, linear estimation and identification methods have been applied by Worsley *et al.* in for instance Refs. 114, 115.

Other statistical work includes that of Solo *et al.*,[103,104] Nan and Nowak,[81] Ardekani *et al.*,[3] Kershaw *et al.*,[68] Purdon and Weisskoff,[91] and Genovese *et al.*[44] In Refs. 103, 104, prior information is used to model the impulse response and the hemodynamic response delay, with regularization added as a constraint to take advantage of the spatial correlation of the brain response. The approaches in Refs. 3, 68, 81 account for the fact that some regions of the brain may be responding to activities not related to the experiment being conducted. Responses to these activities are assumed to lie in a separate subspace, often called the interference subspace. The hypothesis tests compare the measurement components in each of the response subspace and this separate subspace, which if not known *a priori*, is learned from the experimental data.

3.1. *Our Approach*

Our objective in this chapter is to develop methods that improve the detectability of responding voxels, thus providing enhanced fMRI activation maps. Specifically, we focus on learning the characteristics of the response to the stimulus from the experiment's data. In particular *we focus here on estimating the time varying response time shifts separately at each phase of the stimulus on a voxel-by-voxel basis.* Estimating the time shifts can be shown to have a significant impact on enhanced detectability in low signal to noise regions; errors in the estimate of phase time shifts can lead to errors in the relative orientation of the signal subspace inside the measurement space.

Work on estimating response time shifts include that of Lange and Zeger[71] and Glover,[46] where the hemodynamic response function (HRF) is modeled as a gamma function, whose parameters are estimated by nonlinear regression. Another approach based on nonlinear regression, by Purdon *et al.*[92] Methods that assume periodic stimuli are found in the work of Thierry *et al.*[109] An underlying assumption in the above work is that the delay is constant within the same region.

One important difference between the approach suggested here and those mentioned is that in our underlying HRF model, which is motivated by data on responses to visual stimuli, the time shift is allowed to vary not only at each voxel, but at each phase shift as well. The suggested approach consists of the following steps: (1) modeling the stimulus response with the time shifts as a parameter based on Laplacian subspace learning theory, (2) estimating the response time shifts at each voxel, and (3) implementing an appropriate, meaning a robust Laplacian, matched filter detection test.

We use a binary stimulus time response model for our matched filter. We account separately at each voxel for the variation in time of the *response time shifts* from one stimulus phase to the next. To do so, we adopt a simple time varying stimulus response model that is parametrized by the gain and phase varying time shifts. The time shift parameters enter the response model in a nonlinear fashion. The parsimonious yet simple response model captures the essence of the temporally and spatially non-constant characteristics of the actual response signal, and can therefore handle periodic as well as aperiodic signals.

To *estimate the time shifts*, we use the *curve evolution*[100] approach to image smoothing and segmentation prior to applying a constant false alarm rate (CFAR) detector.[25,99] This nonlinear variational estimation approach has been applied in Ref. 90 to structural MRI, makes use of the response's spatio-temporal correlation and yields the response time shifts separately for each on-off phase of stimulus and at each voxel. The power of the curve evolution approach makes it possible to overcome the complication involved with their estimation for each phase change of the stimulus and at each voxel.

While detectability of responding voxels improves with the introduction of time shifts, we *unavoidably* incur a simultaneous penalty of enhanced false alarms; a well known attribute of all tests where model parameters are estimated. We therefore need to adopt additional measures where we can reduce the false alarms associated with a Gaussian noise model based CFAR detector, while maintaining the improved detectability that the estimation of response time shifts provides. We explore here an alternative CFAR detector, specifically one that replaces the Gaussian noise model with a Laplacian one. Like the Gaussian, the Laplacian density function is unimodal and symmetric about the mean. Its tail event probabilities are, however, larger than those of the Gaussian density function, meaning that it assigns a higher probability to intervals of large values, and thus one of the potential causes of increased false alarms. We then formulate a matched subspace detection test that is appropriate for non-Gaussian noise, specifically the robust detection test derived

by Desai and Mangoubi in Ref. 25, on a voxel by voxel basis. In the presence of non-Gaussian noise, this test protects against modeling errors and unlearned physiological effects. Finally, using image comparison in Fig. 4, we demonstrate that with the Laplacian robust CFAR detector, the enhanced detectability obtained by virtue of estimating the response shift or delay is achievable while maintaining a reduced level of false alarms; no additional postprocessing such as spatial smoothing or background masking, etc., is needed.

3.1.1. *Summary of the approach*

Figure 1 captures the approach. The diagram in (a) describes the measurement model (Sec. 4) and its components — stimulus response (Sec. 4.2), learned and unlearned non-stimulus response and noise (Sec. 4.3). The diagram in (b) illustrates the stimulus response model and its parameters (defined in Sec. 4.2). The diagram in (c) describes how the algorithm processes the measurement (Sec. 5): the temporal characterization (Sec. 6.1) provides estimates for the response time shifts at each stimulus phase transition before the robust detection test (Sec. 7) is implemented to obtain the spatial or functional characterization (Sec. 8) of the brain response.

4. The Measurement Model

4.1. *Measurement components*

When stimulating signals are sent to the brain, some areas respond to the stimulus, while others do not. Our objective is to determine, using hypothesis testing and a sequence of MR images of the brain taken while the brain is receiving the stimulating signals, the probability that the observed change at each voxel is due to the stimulus.

Let Ω represent a discrete two-dimensional region for which a time sequence of image measurements are available. We assume that Ω contains but is not limited to an entire two-dimensional slice of the brain. We note that though our development and implementation is two-dimensional, our approach is applicable to the three-dimensional case as well. Let $i \in \Omega$ be a discrete voxel location. The input signal at time step k over a discrete-time interval $[1, K]$ at each voxel i is assumed to be a known sequence of on/off step functions of equal magnitude. We base our hypothesis test on a measurement vector x_i at each voxel i that consists of different components: the response to the *stimulus*, r_i, the response to known or *learned* non-stimulus effects, $r_{\ell i}$, the response to *unlearned* effects, r_{ui}, and the noise n_i. Thus, we have at each time step k
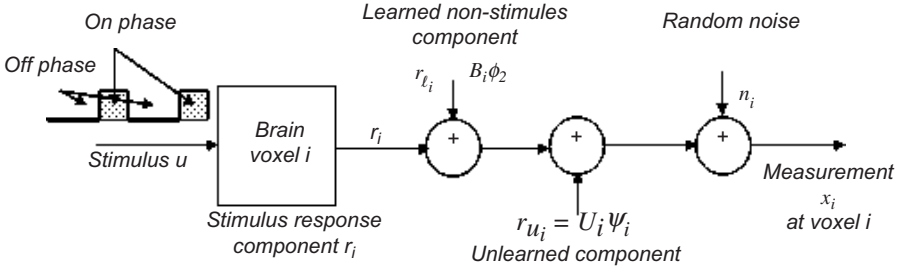
$$x_i[k] = r_i[k] + r_{\ell i}[k] + r_{ui}[k] + n_i[k] \quad \forall k \in [1, K], \ i \in \Omega \tag{1}$$
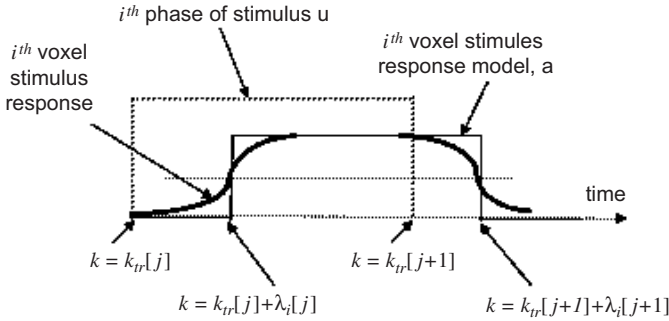
or, in vector notation,

$$x_i = r_i + r_{\ell i} + r_{ui} + n_i \quad \forall i \in \Omega. \tag{2}$$

The above model is illustrated in Fig. 1. We will first discuss the various components of the model, then describe our approach, also illustrated in Fig. 1.
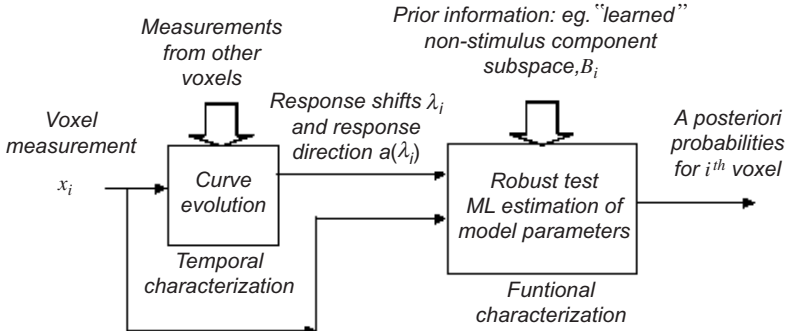
(a) Measurement model components for the i^{th} voxel

(b) Stimulus respone model for i^{th} voxel

(c) The robust detection architecture

Fig. 1.   *The detection model and the architecture.* The diagram in (**a**) describes measurement model (Sec. 4) and its components — stimulus response (Sec. 4.2), learned and unlearned non-stimulus response and noise (Sec. 4.3). The diagram in (**b**) illustrates the stimulus response model and its parameters (defined in Sec. 4.2). The diagram in (**c**) describes how measurements are processed: The temporal characterization (Sec. 6.1) provides estimates for the response time shifts at each stimulus phase transition before the robust detection test (Sec. 7) is implemented to obtain the spatial or functional characterization (Sec. 8) of the brain response. ©2002, IEEE.[23]

4.2. *The stimulus response model, r_i*

An underlying assumption in Sec. 4.2 is that the stimulus response signal time pattern is similar to the stimulus time pattern, except that it is *shifted in time and modulated by a gain*. The actual response is generally more complex, but the time shifts are intended to capture the transient behavior, which may be anticipatory, and to capture transport delays, while the gains capture the steady state behavior. As explained in Ref. 26 and Sec. 9.1, our choice of stimulus response model complements our choice of Laplacian noise model, described shortly in Sec. 4.3. Specifically, if we were to formulate a subspace learning problem in the presence of Laplacian noise without imposing parametric dependency on the form or orientation of the subspace, then the optimal subspace estimate is represented by a vector whose elements are binary valued, such as $+1$ and -1. This is exactly the type of model we describe in this section.

In our parametric subspace estimation problem here, where again we assume the binary form, the time shifts and the gains are the two parameters to be estimated. In this study, the gains are assumed constant in time at each voxel, but vary from voxel to voxel. The time shifts vary from voxel to voxel and, in particular, *they also vary from response phase to response phase*. In previous work, the gains and shifts are also unknown, and spatially varying time delays are studied and explained by, for instance, Lange and Zeger,[71] and Solo *et al.*[103] These authors, as mentioned in the introduction, assume that the time shifts at a given voxel are constant, while we do not. Time delays are allowed to vary because they have a dominant role in identifying the signal response subspace. For example, for a 66 sample signal of 6 on and 7 off phases in a one-dimensional signal space, a misclassification consisting of a delay of one sample at each phase would yield a response space that is separated by 51 degrees from the actual response space.

As mentioned earlier, the input signal at time step k over a discrete-time interval $[1, K]$ at each voxel i is assumed to be a known sequence of on/off step functions of equal magnitude. It models the on and off phases of stimulus with onset of the jth phase, $j = 1, \dots, J$, given by $k_{\mathrm{tr}}[j]$ as shown in Fig. 1(b). The subscript "tr" is for transition.

We assume a simple two-parameter linear, time varying response model; a shifted version of the stimulus signal that is modulated by a gain. The change in response at the ith voxel due to a unit step applied at time k_0 is then modeled as

$$r_i[k] = \theta_i h \left[k - k_0 - \lambda_i[k_0] \right], \quad r_i[1] = 0, \tag{3}$$

where θ_i and $\lambda_i[k_0]$ represent the gain and time delay, respectively, and $h[k]$ is a unit step function defined as $h[k] = 1$ for $k \geq 0$ and zero otherwise. We assume that the time of response onset $k_0 + \lambda_i[k_0]$ occurs after step $k = 1$.

We can obtain in vector notation the following response model for the alternating on and off stimulus signal of Fig. 1(b),

$$r_i = a\left(\lambda_i\right)\theta_i \tag{4}$$

where λ_i is the vector of individual phase delays of length J, the number of phases, while the vectors r_i and $a(\lambda_i)$ are of length K. The vector $a(\lambda_i)$ is given by

$$a(\lambda_i) = [\delta[1]e_{b_i}[1], \delta[2]e_{b_i}[2], \ldots, \delta[J]e_{b_i}[J]]' \tag{5}$$

where the superscript $'$ denotes the transpose. Here each $e_{b_i}[j]$ is a row vector of ones whose length equals the number of time steps of the response in the jth stimulus phase. Specifically, $e_{b_i}[j]$ is a vector of 1's of length $T_i[j]$, where

$$T_i[j] = k_{\mathrm{tr}}[j+1] - k_{\mathrm{tr}}[j] + \lambda_i[j+1] - \lambda_i[j], \quad \lambda_i[1] = 0, k_{\mathrm{tr}}[J+1] = \lambda_i[J+1] = K. \tag{6}$$

We also have

$$\delta[j] = \begin{cases} 1 & \text{if on phase} \\ 0 & \text{if off phase} \end{cases}. \tag{7}$$

For example, if $k_{\mathrm{tr}}[2] = 6$ and $k_{\mathrm{tr}}[3] = 10$, and $\lambda_i[2] = 2$, $\lambda_i[3] = 1$ for a voxel i, then the stimulus $u = a(0)$ and $a(\lambda_i)$ would respectively start as

$$u = [\, 0\ 0\ 0\ 0\ 0\ 0\ 1\ 1\ 1\ 1\ 0\ 0\ \ldots \,]'$$
$$a(\lambda_i) = [\, 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 1\ 1\ 1\ 0\ \ldots \,]'.$$

The simple two-parameter step response model of Eq. (4) therefore implies that the response at each voxel $i \in \Omega$ is a $K \times 1$ vector residing in a *one-dimensional subspace* parametrized by the response time shifts λ_i, and whose magnitude is scaled by the gain θ_i. Estimating the gain completely fixes the stimulus response vector.

We note that the above model is a generalization of the box-car shaped HRF used in early fMRI work by Bandetini.[4] More generally, the one-dimensional subspace model $a(\lambda_i)$ is one that embraces many shapes associated with hemodynamic responses and thus introduces additional degrees of freedom. Further, since the model is time domain and not frequency domain based, it can accommodate aperiodic signals as well.

In Sec. 9, we provide a more general response model, where we allow the gain magnitudes to vary from phase to phase. As mentioned earlier, we also briefly describe an attractive feature of the on-off response model. Specifically, as shown in Ref. 26 for subspace learning in the presence of Laplacian noise, where no parametric dependency constraints are imposed on either the form or the orientation of the subspace, this model form naturally obtains as it is the optimal subspace estimate.

4.3. *The measurement model,* x_i

In addition to the stimulus response, the measurement vector may also include responses to other known or learned *interference*, unknown or unlearned effects, and random noise, as Fig. 1(a) shows.

4.3.1. *Learned non-stimulus effects, r_{li}*

By learned interference, we mean that we have determined the subspace where they reside. In addition to bias, this can include low frequency drift phenomena such as a ramp, or cardiac and breathing effects. These effects, if not specified *a priori*, can be learned, for instance, over a larger group of voxels assuming local spatial stationarity; an aspect of the fMRI problem we do not address in this study. We denote by B_i the learned non-stimulus effect matrix of dimension $K \times M$ whose columns span an M-dimensional subspace of the K dimensional measurement space at voxel i. Some columns of B_i, such as those representing drift phenomena, may be the same for many or all voxels. Thus, our learned response model is given by

$$r_{\ell i} = B_i \phi_i \quad \forall i \in \Omega \tag{8}$$

where ϕ_i is an unknown parameter vector representing the strength of the learned interference. Again we do not deal in this paper with the learning aspect of the fMRI problem; we simply assume that the matrix B is known and has been learned *a priori*.

4.3.2. *Unlearned effects, r_{ui}*

We designate as unlearned effects those effects that we did not learn, or ignore in the model. Unlearned effects can be due to undersampling, modeling errors, errors in learning of physiological effects, unforeseen or ignored phenomena, etc. One can only unambiguously see this unlearned component in the space orthogonal to the stimulus and learned effect subspaces. Therefore, the presence of these effects can only be inferred from the component of the measurement that lies in the *null space* of the space jointly spanned by the columns of $a(\lambda_i)$ and B_i. This null space contains those unlearned effects that are decoupled from the subspaces spanned by the columns of a and B. We denote the subspace of unlearned effects by $\mathcal{U}_i(\lambda_i)$ and we assume it is spanned by the columns of the $K \times (K - M - 1)$ matrix $U_i(\lambda_i)$. Thus we have

$$r_{ui} = U_i(\lambda_i)\psi_i \quad \forall i \in \Omega \tag{9}$$

where ψ_i represents the unlearned interference strength, and needs to be estimated as well.

4.3.3. *The noise, n_i*

The noise is a $K \times 1$ random vector modeled as

$$n_i = \sigma_i v_i$$

where σ_i is an unknown scalar that represents the time-invariant voxel dependent standard deviation, and v_i is a $K \times 1$ random vector of zero mean and unit covariance whose elements are independent and identically distributed. We explicitly show

the time invariant voxel dependent standard deviation σ_i in the equations of the hypothesis test to emphasize the fact that the covariance is unknown as well. We will consider two density functions for v_i, the Gaussian and the Laplacian. As mentioned in the introduction, we consider a Laplacian distribution model for the noise because the heavy tailed nature of the Laplacian has been shown in other applications to provide additional robustness to modeling errors.[66] Aside from robustness considerations, the noise itself may be Laplacian, or generally non-Gaussian.[53] We note that in other work,[7,45] non-Gaussianity is dealt with at the analysis or thresholding level, as opposed to the decision statistic synthesis level, as done here. The theory of non-Gaussian subspace learning is developed in Ref. 25. Moreover, Descombes *et al.*,[31] and Everitt and Bullmore,[38] use non-Gaussian density functions for spatial processing, whereas we focus on temporal processing. We also note that our independent noise assumption may be an idealization, as the noise may well be spatially or temporally correlated. Temporal correlations are considered by Purdon and Weisskoff[91] and Worsley *et al.* in Ref. 114.

4.3.4. *The overall measurement model, x_i*

Based on the above, our measurement model (Fig. 1) for a responding voxel i has the following form

$$
\begin{aligned}
x_i &= r_i + r_{\ell i} + r_{ui} + n_i \\
&= a(\lambda_i)\theta_i + B_i\phi_i + U_i(\lambda_i)\psi_i + \sigma_i v_i \quad \forall i \in \Omega.
\end{aligned} \tag{10}
$$

The unknown parameters in the above model are again the response shift parameter λ_i, the voxel's response magnitude θ_i, the learned interference strength vector ϕ_i, the unlearned interference strength vector ψ_i, and the noise standard deviation σ_i. The measurement is linear in all these parameters, except for the λ_i's, which enter the model in a nonlinear fashion. We note that, with the parametrization of the signal subspace $a(\lambda_i)$, we are lifting the assumption that the stimulus response subspace is known *a priori*. We also note that our stimulus response subspace is only one dimensional, though we provide a more general extension to the subspace of piecewise continuous functions in Sec. 9, albeit preferably under some constraints on parameters to reduce the attendant increase in false alarms. Finally, not all measurement components in Eq. (10) need be simultaneously present; the actual measurement model depends on the hypotheses of the detection test that is posed, as we explain in Sec. 7.

5. Model Parameter Estimation: Time Shift (λ's) versus Gain (θ's)

The linear parameters θ will be estimated in a robust generalized likelihood ratio test framework in Sec. 7. In that section, the estimation of the gain θ is imbedded in the derivation of the detector's statistics, and there is no need for an explicit expression.[25–28]

As noted in Fig. 1, we tackle the detection and gain parameter estimation issue only after curve evolution is used for estimation of the λ_i's. We now explain the motivation behind this two-step hierarchical approach.

The spatial correlation in the stimulus response time shifts can be significant among responsive voxels. Therefore, when estimating those shifts, we need an approach that takes into account their spatial continuity. The curve evolution approach is particularly attractive for this purpose. We mention three advantages: (1) it can segment the image into regions in space and time, and simultaneously smooth only within regions, (2) the resulting time shift estimates $\hat{\lambda}_i$ this method leads to are relatively less sensitive to low frequency phenomena such as bias and ramp, and (3) we do not need to be concerned with implementing constraints on other parameters such as the gain.

We do not, however, need the curve evolution method for obtaining the stimulus response gains because, at least when compared to the time shifts, their spatial continuity is not as consistent, for reasons such as partial volume effects. Moreover, the curve evolution approach does not easily accommodate parameter constraints should these prove necessary, such as when they are assumed to be constant. The relative inconsistency of the spatial continuity of the gains and the possible need to impose constraints on them leads us to formulate voxel specific hypothesis tests. Moreover, characterization of the stimulus response lends itself readily to the imposition of constraints on the gains, either explicitly, or implicitly through our choice of response model.

The next section describes the use of curve evolution to estimate the response time shifts, while Sec. 7 describes the voxel-by-voxel hypothesis test used in functional characterization.

6. Response Time Shift (λ's) Estimation

Our objective here is to temporally characterize the stimulus response subspace parameters $\lambda_i, \forall i \in \Omega$ to enhance detectability of response to stimulus. To do so, we use the curve evolution[100] approach to the simultaneous smoothing and segmentation of a time sequence of images. As mentioned earlier, this approach is particularly adept at exploiting the possible spatial correlation among neighboring responding voxels.

6.1. *Curve evolution*

The stimulus response model and the hypothesis test is presented in the previous section in discrete time and discrete space. The reason for this choice is that our measurements are in discrete time and space to begin with. To present our variational problem, however, we will switch to the continuous time and space domains, both because the continuous formulation provides more conceptual insight and, in any case, Shah's original formulation[100] is continuous. For implementation, the resulting necessary conditions are of course discretized.

For the continuous domain formulation, we consider the image region Ω introduced in Sec. 4 to be a discretization of a continuous region, which we denote $\Omega_c \subset R^2$, with elements $s \in \Omega_c$. Likewise, the time variable k in the time interval $[1, K]$ is a discretization of a continuous time variable $t \in T \equiv [t_0, t_f]$. The analog of our measurement vectors $x_i, i \in \Omega$ will be the functions $x_c(s, t)$, respectively, for $s \in \Omega_c$, $t \in T$. In the approach of Ref. 100, we have two unknown space *and time varying* functions, $\zeta \equiv \zeta(s, t)$ and $\nu \equiv \nu(s, t)$, where $\zeta(s, t)$ represents the piecewise smooth model of x_c, meaning the *image field*, and $\nu(s, t)$ represents points in $\Omega_c \times T$ with significant change in the measurement. Both ζ and ν are to be estimated so as to minimize the energy cost functional $\mathcal{E}(\zeta, \nu)$, given by

$$
\mathcal{E}(\zeta, \nu) = \iint_{\Omega_c \times T} \Big((1 - \nu)^2 (\alpha_1 \|\nabla_s \zeta\| + \alpha_2 |\zeta_t|) + \beta |\zeta - x_c|
$$
$$
+ \frac{\rho}{2} \big(\|\nabla_s \nu\|^2 + \nu_t^2 \big) + \frac{\nu^2}{2\rho} \Big) ds\, dt \tag{11}
$$

where α_1, α_2, β, and ρ are constants representing the chosen weights on the accompanying cost components, ∇ denotes the gradient operator, and ζ_t, ν_t denote partial derivatives with respect to time for the functions ζ and ν respectively. The term $\|\nabla_s \zeta\|$ and $\|\nabla_s \nu\|$ represent, respectively, the magnitude or Euclidean norm of the gradient of ζ and ν with respect to the spatial variable s at the point s, t.

The first term in the above functional provides smoothness in the space and time domains at all interior points of a region where $\nu \ll 1$, the second term is for data fidelity, the third term provides smoothness in both domains for the edge field, while the fourth term is intended to penalize the excessive presence of edges.[90,100] Note that the measurement fidelity cost term $|\zeta - x_c|$ is based on the 1-norm; a stochastic interpretation of which can be associated with Laplacian measurement noise.

We note again that the variational formulation just given is independent of both the input stimulus and our choice of response and measurement models in the hypothesis test. It also can provide a segmented image in space and time, with the regions in space and time between edges and the edges themselves smoothed, but, as we explained in the previous section, we choose to use it to extract estimates of the response time shifts.

6.2. *The gradient descent equations*

The Euler equations that are the necessary conditions associated with the minimization of the energy functional of Eq. (11) can be solved by the gradient descent method. If ς is the rate of descent variable, then the gradient descent equations are given by

$$
\frac{\partial \zeta}{\partial \varsigma} = -2(\alpha_1 |\zeta_t| \langle \nabla_s \nu, \nabla_s \zeta \rangle + \alpha_2 \|\nabla_s \zeta\| \nu_t \zeta_t)
$$
$$
+ (1 - \nu)\|\nabla_s \zeta\| |\zeta_t| \Big(\alpha_1 \Big\langle \nabla_s, \frac{\nabla_s \zeta}{\|\nabla_s \zeta\|} \Big\rangle + \alpha_2 \frac{\partial}{\partial t} \Big(\frac{\zeta_t}{|\zeta_t|} \Big) \Big)
$$
$$
- \frac{\beta}{1 - \nu} \|\nabla_s \zeta\| |\zeta_t| \frac{\zeta - x_c}{|\zeta - x_c|} \tag{12}
$$

$$\frac{\partial \nu}{\partial \varsigma} = \nabla_s^2 \nu - \frac{\nu}{\rho^2} + \frac{2}{\rho}(1-\nu)(\alpha_1 \|\nabla_s \zeta\| + \alpha_2 |\zeta_t|) \tag{13}$$

with conditions on the boundary $\partial \Omega_c$ and ∂T given by

$$\left.\frac{\partial \zeta}{\partial \eta}\right|_{\partial \Omega_c} = 0; \quad \left.\frac{\partial \nu}{\partial \eta}\right|_{\partial \Omega_c} = 0; \quad \left.\frac{\partial \zeta}{\partial t}\right|_{\partial T} = 0; \quad \left.\frac{\partial \nu}{\partial t}\right|_{\partial T} = 0; \tag{14}$$

where η is the direction normal to $\partial \Omega_c$. In the above equations, $\langle \cdot, \cdot \rangle$ denotes the dot product and ∇_s^2 is the Laplacian.

The solutions to the above equations, $\hat{\zeta}$ and $\hat{\nu}$, provide estimates for the measurement and edge strength signals at each voxel location. Since $\nu \in [0,1]$,[90,100] it can be interpreted as the probability that a particular voxel is an edge voxel. If we look at the history of the measurement signal at each voxel, the discretized solution to the above gradient descent equations thus provides us with sufficient information to estimate the time of significant shifts in the measurement signals, meaning the λ_i's. We chose to determine the time shifts at each phase transition and for each voxel based on applying likelihood considerations to the data over an interval surrounding each transition. Specifically, the transition point is allowed to shift from its nominal position in the stimulus within a certain range, and the likelihood function is calculated for each shifted position. Finally, we note again that the solution to the above problem makes no use of the response or measurement model introduced in the previous section.

6.3. *Results of shift estimation*

Figure 2(a) shows, for a particular responsive voxel, bands representing on and off phases of the stimulus signal, and the measurement signal, while Fig. 2(b) shows the same measurement signal as in (a) (dashed), and the *measurement* signal estimate (solid) from the curve evolution approach. We see that the estimate of the measurement signal is *deblurred* between phase transitions, as the model of Eq. (4) dictates. The plots in (c) show the measurement signal and time shifted on and off bands of the measurement signal. Note the ability of the curve evolution approach to preserve sharp discontinuities in time, making it possible to obtain accurate estimates of the response time shifts, as can be seen from the placement of the bands with respect to the actual measurement in (c). We note that Eq. (12) is hyperbolic[100] in the direction of the gradient of ζ, and is therefore capable of creating "shocks" or sharp discontinuities.

Figure 3 shows the average delay times, or the difference between the transition times of the measurement signal and the stimulus signal, for all active voxels for 13 phase transitions. In this experiment, all the time shifts are delays. The variation in average delays shown in Fig. 3 underlines the importance of estimating individual delay times at each transition.
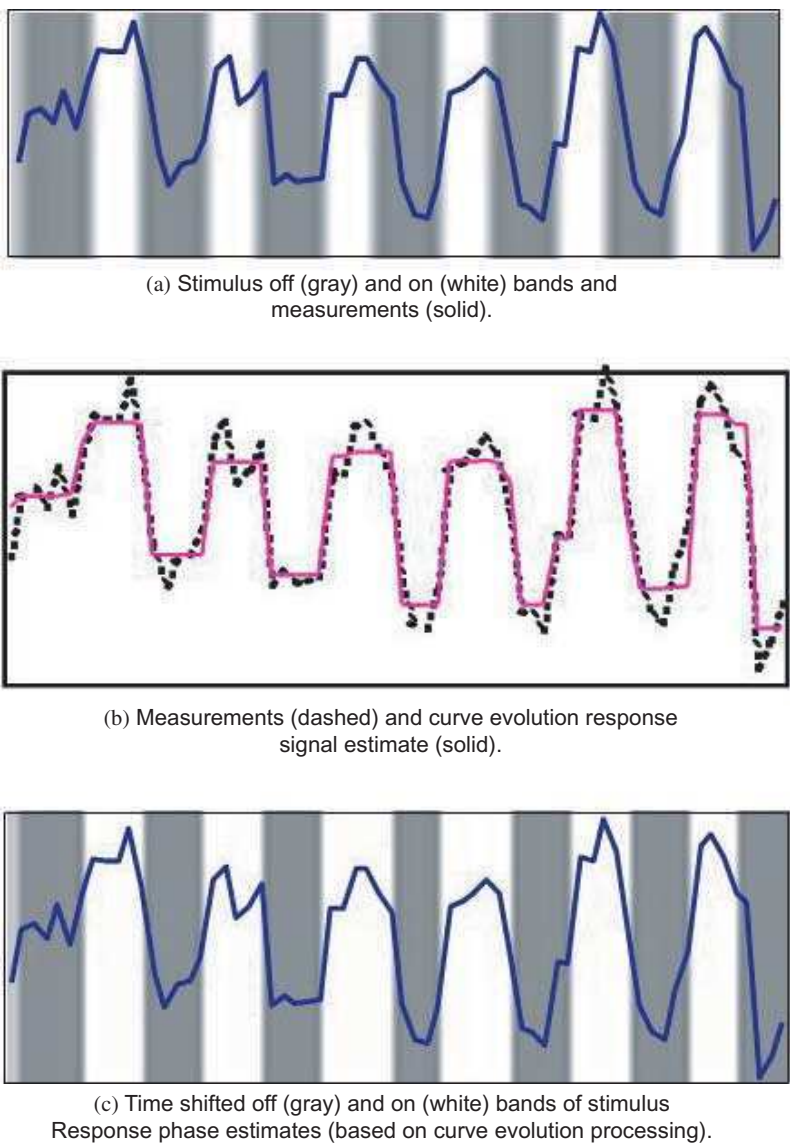
(a) Stimulus off (gray) and on (white) bands and
measurements (solid).

(b) Measurements (dashed) and curve evolution response
signal estimate (solid).

(c) Time shifted off (gray) and on (white) bands of stimulus
Response phase estimates (based on curve evolution processing).

Fig. 2.   *Temporal characterization of the brain response.* Plots in **(a)** show actual measurements
for a voxel with a significant response to an external stimulus with bands showing the onset and
time duration of on (white) and off (gray) phases of the stimulus. Response to stimulus is delayed
with variable delays at each phase transition. In **(b)**, the dashed curve shows the measurement of
(a) along with its estimate (solid) from the curve evolution approach based processing. Estimate
is seen to be piece-wise constant with sharp transitions, a characteristic that readily leads to the
estimation of the time shifts at each phase transition of the input stimulus. Information of **(c)** is
analogous to that of (a) with a difference. The bands represent now the estimated on/off phases
of the voxel response and not of the input stimulus and band boundary placement is seen to be
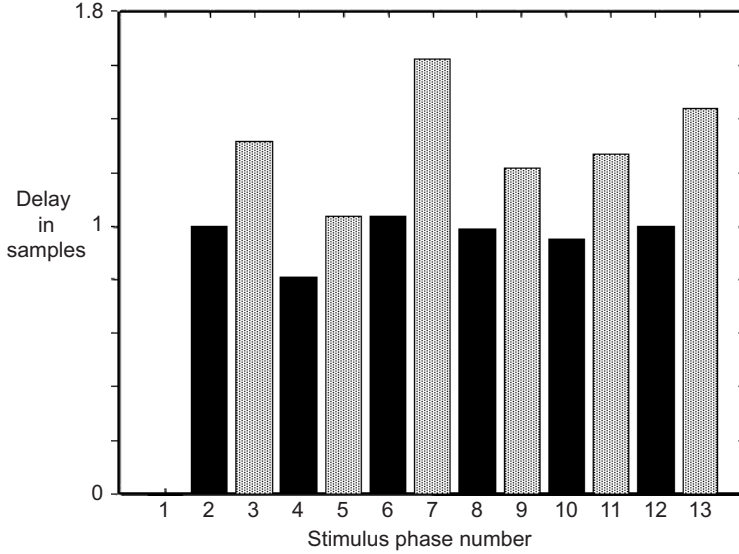at the middle of the measurement transitions as desired from modeling considerations. ©2002,
IEEE[23].

Fig. 3. *Shifts in response in sample times for thirteen phases.* Off to on (shaded) and on to off (dark) phase transitions, are averaged over most active voxels. The figure shows that the delay time can be different from phase to phase. ©2002, IEEE[23].

6.4. *Summary of response time shifts*

The time shift estimates enable the determination of the subspaces for stimulus response. The *Temporal Characterization* referred to in the first block of Fig. 1c is thus accomplished. We are ready to proceed to the final step of our hierarchical algorithm, labeled *Functional Characterization* in the second block of Fig. 1c.

7. The Robust Detection Test

We first formulate the problem for general signal subspaces in Sec. 7.1.[25] When the subspaces are spanned by vectors that are binary valued, such as $+1$ or -1, the detectors possess a special structure, which we describe in Sec. 7.4.[28]

7.1. *Robust hypothesis test formulation*

We pose a hypothesis test at each voxel $i \in \Omega$, where the null hypothesis H_{i0} states that the voxel is not responding, while the alternate hypothesis H_{i1} states that it is. Specifically,

$$H_{i0}: x_i = U_i(\lambda_i)\psi_i + B_i\phi_{i0} + \sigma_{i0}v_{i0}, \quad \mathcal{U}_i(\lambda_i) = \mathcal{N}([a(\lambda_i) \ B_i]) \tag{15}$$

$$H_{i1}: x_i = a(\lambda_i)\theta_i + B_i\phi_{i1} + \sigma_{i1}v_{i1}. \tag{16}$$

In effect, we are testing whether the response, separately at each voxel i, is due to the unlearned effects, meaning $\theta_i = 0$, or to the stimulus, meaning $\psi_i = 0$. This test is *robust* to unlearned interference because it explicitly accounts for the presence of unlearned interference in the hypothesis H_{i0}. The robust test uses a game theoretic approach.[25] We also note that in the statistics literature, the term robust hypothesis testing often refers to work by, for instance, Huber[62] that specifically addresses the issue of statistical outliers. Also, in detection work related to the broader domain of dynamic plants, robustness is usually with respect to plant model and noise model uncertainties for example by Mangoubi.[78] Both these methodologies are motivated by different applications than the approach presented here.

Our choice of this test is data driven and is intended to reduce false alarms that arise in the presence of non-Gaussian noise with a nominal test where $U_i(\lambda_i) = 0$ in hypothesis H_{i0}. We note that in the Gaussian case, our robust test yields the F-statistic.

Moreover, when the noise is Gaussian, we do not need to explicitly account for the presence of unlearned interference, meaning that we can set $\psi_i \equiv 0$ and obtain a statistic whose ROC performance is equivalent to that of the F-statistic. But this equivalence obtains only in the Gaussian case, and only for CFAR tests, as we explain in detail in Ref. 25 When the noise is Laplacian, a test where we set $\psi_i = 0$ *a priori* will yield a different statistic with different performance characteristics than a test where ψ_i is estimated from the data. The difference between the effects on the output images of a Laplacian and a Gaussian noise model is illustrated in Fig. 4 of Sec. 8.

7.2. *The robust detectors for general subspaces*

We now provide expressions for the Generalized likelihood ratio (GLR) robust detectors that obtain from the test of Eqs. (15–16) for different density functions. We note that in these detectors, the estimate of the gain and noise variance parameter, are implicitly imbedded in the final expression for the robust detection's statistic, and no explicit expression is needed. Derivations as well as other and more general detectors are found in Ref. 25.

7.2.1. *Gaussian detector*

For the Gaussian case, the Generalized likelihood ratio based F-distributed statistic is used. At each voxel $i \in \Omega$, we have

$$F_i(x_i; \hat{\lambda}_i) = \frac{x_i' P_{((I - P_{B_i})a(\hat{\lambda}_i))} x_i / L}{x_i' P_{U_i(\hat{\lambda}_i)} x_i / (K - L - M)} \tag{17}$$

where, for an arbitrary matrix W, the projection matrix is given by $P_W \equiv W(W'W)^{-1}W'$, with the superscript $'$ indicating a matrix transpose operation, and $L = 1$, the dimension of the response space $a(\hat{\lambda}_i)$. Since the signal response space is one-dimensional, this statistic is closely related to a two-sample t test. We note that the early works of Bandetini[4] and Cox,[18] which makes use of box-car HRF's, relies on statistics that are related to the t statistic as well.
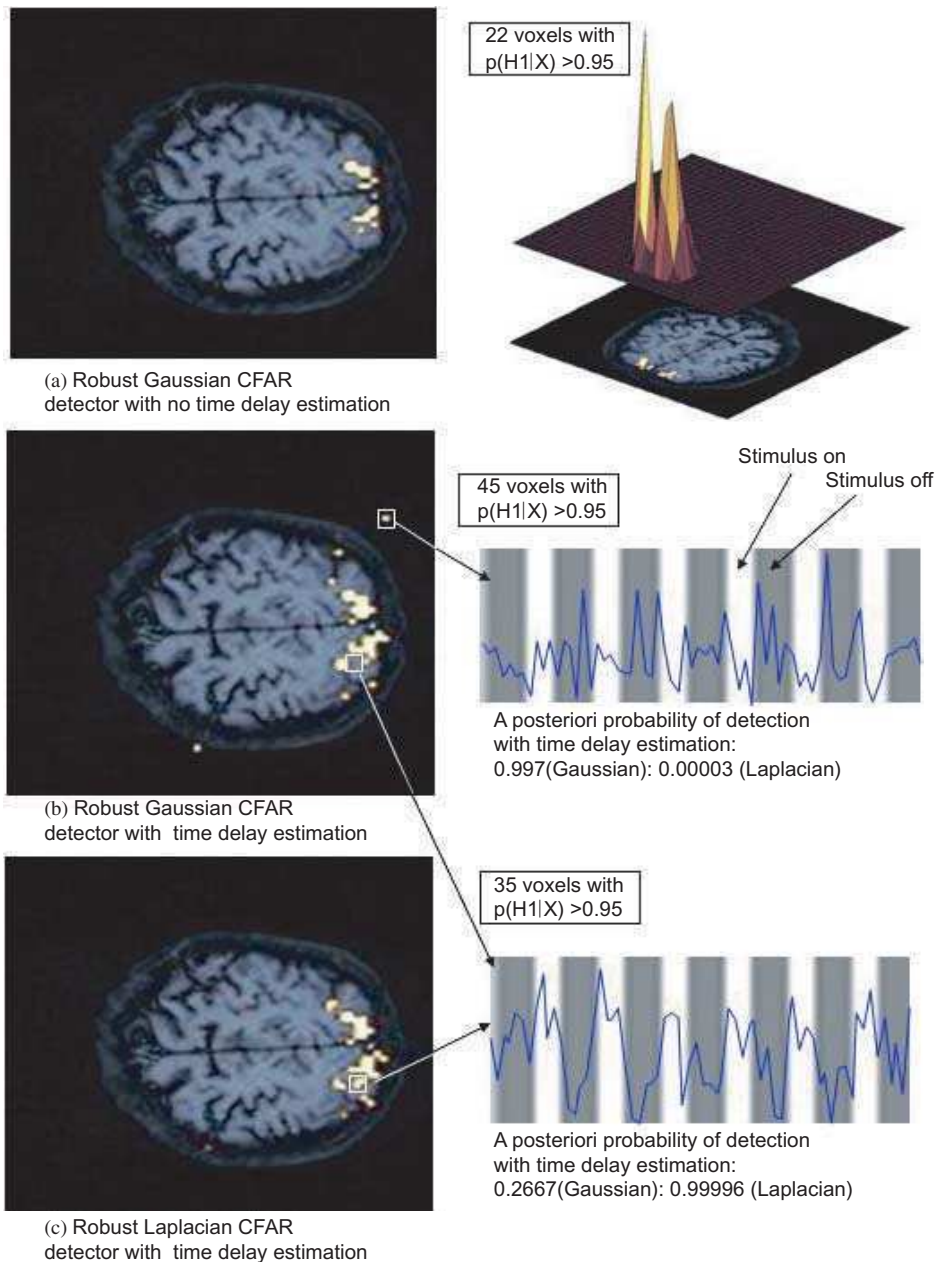
(a) Robust Gaussian CFAR detector with no time delay estimation

22 voxels with p(H1|X) >0.95

45 voxels with p(H1|X) >0.95

Stimulus on
Stimulus off

A posteriori probability of detection with time delay estimation: 0.997(Gaussian): 0.00003 (Laplacian)

(b) Robust Gaussian CFAR detector with time delay estimation

35 voxels with p(H1|X) >0.95

A posteriori probability of detection with time delay estimation: 0.2667(Gaussian): 0.99996 (Laplacian)

(c) Robust Laplacian CFAR detector with time delay estimation

Fig. 4. *Detectors' outputs.* The image in the LHS of **(a)** is the output of a detector based on a Gaussian noise model, with response time shift or delay set to zero. On the RHS is the *a posteriori* probability map. In **(b)**, response time shifts estimates are introduced: additional voxels are classified as responding. In the image in **(c)**, the Laplacian noise model replaces the Gaussian one: some of the voxels declared in (b) as responding are withdrawn, while others are added. On the RHS of the images in (b) and (c) are the time responses of two different voxels. The plots indicate that the Laplacian detector correctly classified the two voxels, while the Gaussian detector did not. A close examination of all voxels' time response indicated that the Laplacian detector incorporating response time shift estimates is more accurate than a Gaussian one. ©2002, IEEE[23].

7.2.2. *Laplacian detector*

When the noise is Laplacian, we consider the generalized likelihood ratio statistic for the test of Eqs. (15) and (16). When the signal space is one-dimensional such as $a(\hat{\lambda}_i)$, this statistic is[25]

$$\Lambda_i(x_i, \hat{\lambda}_i) = \frac{|x_i'(I - P_{B_i})a(\hat{\lambda}_i)|}{\|(I - P_{B_i})a(\hat{\lambda}_i)\|_\infty \, \|x_i - a(\hat{\lambda}_i)\hat{\theta}_i - B_i\hat{\phi}_{i1}\|_1} \tag{18}$$

where for an arbitrary scalar α, $|\alpha|$ denotes its absolute value, and for an arbitrary vector β with elements $\beta_j, j = 1, \ldots, J$, $\|\beta\|_\infty \equiv \max_j |\beta_j|$ is the infinity norm of β, and $\|\beta\|_1 \equiv \sum_{j=1}^{J} |\beta_j|$ is the 1-norm β. Finally, $\hat{\theta}_i$ and $\hat{\phi}_{i1}$ are maximum likelihood estimates of the gains θ_i and ϕ_{i1}, respectively, when the noise is Laplacian. We note that the computation of Λ_i in Eq. (18) involves searches for maximum likelihood estimates, but with the increasing power of computers, the required computation is a doable task.

7.2.3. *The a posteriori probability of detection*

Tabulated values for the F-statistic distribution are readily available for the null hypothesis when no unlearned effects are present. Tables for the statistic of Λ_i of Eq. (18), however, are not available. For this reason, and for ease of comparison of performance using the Gaussian and Laplacian noise models, we instead use the *a posteriori* probabilities for characterizing the response. Specifically, we have for the Laplacian and Gaussian case, respectively,

$$p_{\ell i}(H_1|x_i) = \frac{\Lambda_i}{1 + \Lambda_i} \tag{19}$$

$$p_{gi}(H_1|x_i) = \left(1 + \left(\frac{K - L - M}{L}\frac{1}{F_i}\right)^{K/2}\right)^{-1}. \tag{20}$$

We emphasize that the *a posteriori* probability of detection is introduced here so as to provide a common basis for comparing response maps provided by different detectors. They are not used as a basis for threshold selection. A data adaptive approach for threshold selection based on false discovery rate is developed by Benjamini and Hochberg.[7]

7.3. *Generalization of the t and F statistics for Generalized Gaussian (GG) noise*

The family of Generalized Gaussian densities is a large family of densities, of which the Gaussian and the Laplacian are but two members. For an arbitrary random variable x, this density function is given by

$$p(x) = \left(\frac{p}{2\omega\Gamma(1/p)}\right)\exp-\frac{\|x - m\|_p^p}{\omega^p} \tag{21}$$

where m, ω, and p are, respectively, the location, width, and shape parameters of the density functions. The Gaussian density function obtains with $p = 2$, while the Laplacian one obtains with $p = 1$.[25,28]

The detectors given above for the Gaussian and the Laplacian noise density functions are but two members of the large Generalized Gaussian (GG) family of detectors.[25] The GG family of detectors are generalizations of the well known t and F statistics. Specifically, in the absence of known interference and when $p = 2$, all the detectors above have the same performance as the CFAR detectors based on the t and F Gaussian based statistics.[25,99] One implication is that the t and F statistics are robust to interference whose subspace is completely unknown.

In the next section, we provide expressions for these detectors for binary subspaces,[28] the model we assume for our fMRI response and known interference. These expressions are simpler and more computationally efficient than those for general subspaces.

7.4. *Specific subspaces: binary valued subspaces and Generalized Gaussian detectors*

We now illustrate the simplifications that obtain when we consider the case where a has binary valued elements $+1$ or -1, and b represents a bias vector whose elements are 1. We rearrange the elements of a as follows

$$a = \left[\mathbf{1}'_{K_+}, -\mathbf{1}'_{K_-} \right]' \tag{22}$$

where $\mathbf{1}_{K_+}$ and $\mathbf{1}_{K_-}$ are column vectors of 1's with length K_+ and K_- respectively. Naturally $K = K_+ + K_-$, and the rearranged x can be compatibly partitioned as $x = [x_+, x_-]$.

With U also compatibly partitioned as $U = [U_+, U_-]$, and defining $\zeta_+ = \theta + \phi_1$ and $\zeta_- = -\theta + \phi_1$, the robust test is then rewritten as follows

$$H_0 : x = \begin{bmatrix} \mathbf{1}_{K_+} \\ -\mathbf{1}_{K_-} \end{bmatrix} \phi_0 + \begin{bmatrix} U_+ \\ U_- \end{bmatrix} \psi_0 + \begin{bmatrix} \omega_0 v_{0+} \\ \omega_0 v_{0-} \end{bmatrix} \tag{23}$$

$$H_1 : x = \begin{bmatrix} \mathbf{1}_{K_+} \zeta_+ \\ -\mathbf{1}_{K_-} \zeta_- \end{bmatrix} + \begin{bmatrix} \omega_1 v_{1+} \\ \omega_1 v_{1-} \end{bmatrix} \tag{24}$$

where ω_0 and ω_1 are the Generalized Gaussian width parameters (Eq. (21)) under hypothesis H_0 and H_1 respectively, and are proportional to the noise standard deviations σ_0 and σ_1, respectively, introduced in Eqs. (15–16). The special structure of the likelihood ratio detectors have been derived in terms of ρ_0 and ρ_1, the residuals for hypotheses H_0 and H_1, respectively, that are implicitly defined in Eqs. (27–26). Specifically, for an arbitrary Generalized Gaussian noise with shape parameter p, we have for the case where both gains and noise variance are unknown, the GLR decision statistic[28]

$$\lambda_p = \frac{\|\rho_{0p}\|_p}{\|\rho_{1p}\|_p} \tag{25}$$

where

$$\|\rho_1\|_p^p = \|x - a\hat{\theta}_p - b\hat{\phi}_{1p}\|_p^p$$
$$= \|x_+ - \mathbf{1}_{K_+}\hat{\zeta}_{+,p}\|_p^p + \|x_- - \mathbf{1}_{K_-}\hat{\zeta}_{-,p}\|_p^p \qquad (26)$$

and

$$\|\rho_0\|_p = \|x - b\hat{\phi}_{0p}\|_p^p$$
$$= \frac{|\bar{x}_+ - \bar{x}_-|}{\|\delta\|_{q-1}^{1/p}} \qquad (27)$$

where $\delta = [1/K_+, 1/K_-]$.

It may be observed that the two-dimensional subspace spanned by a and b possesses the special *separability* structure shown in Eq. (24). Such a structure leads to detectors that are computationally efficient,[28] as also can be seen from Eq. (26). Moreover, the simplicity of Eqs. (26) and (27) has significant implications for signal design, and we discuss these in Sec. 9.2.

8. Results

Our objective here is to examine the simultaneous effect of the various enhanced detectability and robustness features we discuss above using experimental fMRI data. We compare images and *a posteriori* probability maps. We caution that the empirical observations we offer are relative and subjective, as no ground truth for the functional characterization of the brain is available.

8.1. *Data*

Functional MRI measurements were taken on a normal, 27-year-old male subject. Imaging was performed using a 1.5 Tesla General Electric Signal scanner with an ANMR, Inc. retrofit for echo-planar imaging. Asymmetric spin echo imaging was performed with a TR = 2500 ms, TE = 50 ms, Offset = 30 ms. Images are acquired with 3.125 mm in-plane resolution and 7 mm slice thickness. The subject alternately views a flashing checkerboard (7 Hz. flicker rate) and a stationary fixation point. A total of 66 images was taken. In our implementation, the learned effect matrix B in Eq. (10) has one column to account for the presence of measurement bias.

8.2. *Image and a posteriori probability comparison*

For the images on the LHS of Fig. 4, brain voxels in gray have an *a posteriori* probability of response $P(H_1|x)$ lower than 0.95, equivalent under the Gaussian noise assumption to an F statistic value of 69.97 with degrees of freedom 1 and 64 (Eq. (17)). We emphasize again that the *a posteriori* probability $P(H_1|x)$ is used here to facilitate comparison between images output by various detectors, and is not suggested for threshold selection. A surface probability plot is shown in the

RHS of the image of Fig. 4(a). The intensity of the gray voxels in the brain images corresponds to the structural MR image for the same subject. Bright voxels are those that have an *a posteriori* probability of response larger than 0.95 in the surface probability plots shown in (b). The brightness of these voxels are proportional to their actual probability.

8.2.1. *Gaussian detector, no shift estimation*

For the image in Fig. 4(a), a Gaussian noise model is assumed and a nominal time shift of zero is used for all voxels and all phases. The conditions for the image in Fig. 4(b) are the same as those for the image in (a) of the same figure, except that the response time shift or delay estimates as given by the curve evolution approach are included. The algorithm is seen to be more sensitive since 45 voxels in (b) have $p(H_1|x) > 0.95$, characterized in this discussion as "responding" for the sake of comparison, versus only 22 in (a). At many of the added voxels, the response gain is weaker, and our ability to detect them is due to the sensitivity of the curve evolution approach to discontinuities, and its ability to localize these in time.

8.2.2. *Gaussian detector, with shift estimation*

It is clear, however, that at least one of the voxels that are additionally characterized in Fig. 4(b) as "responding" is a false alarm, specifically, the one falling outside the brain area. This is an indication of the vulnerability of the Gaussian based detector to false alarms. This Gaussian detector of Fig. 4(b) has assigned an *a posteriori* probability of response of $p(H_1|x) = 0.997$ to this voxel. The response measurement time history of this voxel is provided on the RHS of the brain image in (b). The white bands indicate the on phases of the stimulus signal, and the gray band indicates the off phases. A close examination of the time response shows clearly that the measurement signal of the voxel outside the brain does not indicate a response. Though it is obvious that we have a false alarm from the fact that the voxel is located outside the brain, this example serves to emphasize that, were a non-responding voxel located inside the brain to provide a similar time response, within bias and scale factor, a false alarm would have been declared by the Gaussian detector corresponding to Fig. 4(b).

8.2.3. *Laplacian detector, with shift estimation*

The image in Fig. 4(c) replaces the Gaussian noise model with a Laplacian one. As a result, there are only 35 instead of the previous 45 "responding" voxels as 13 of the "responding" voxels in (b) are now eliminated, and 3 new ones are declared "responding". For example, the voxel outside the brain discussed earlier has been assigned an *a posteriori* probability of response of $p(H_1|x) = 0$ versus a value of 0.997 for the same probability assigned by the Gaussian detector in (b). This is a possible indication that the Laplacian detector is less vulnerable to false alarms. At

the same time, there are indications that the Laplacian detector can also be more sensitive to responding voxels. Consider for instance the voxel whose time response is shown next to Fig. 4(c). The Laplacian detector assigns a $p(H_1|x)$ greater than 0.999, declaring the voxel responding, while the Gaussian detector corresponding to Fig. 4(b) assigns a value of 0.267 only, declaring it non-responding, as does the Gaussian detector corresponding to the image in Fig. 4(a) which does not use the time response estimates. A look at the response time history, however, indicates that we have a stimulus response at the voxel.

8.2.4. *Detector comparison*

A meticulous visual examination of the response time history of each voxel in the collected data indicates that of the three brain images in Fig. 4, the one in (c), output by a Laplacian detector that incorporates the time shift or delay estimates provided by curve evolution, is the most reliable one. Moreover, the differences in the number of responding voxels is not trivial. Of the two detectors incorporating the response time shifts or delays, the Laplacian detector (Fig. 4c) identifies 20% less responding voxels than the Gaussian one (Fig. 4b). Also, the same Laplacian detector identifies more than 50% additional responding voxels than the Gaussian detector (Fig. 4a) that does not incorporate response time shifts. A possible explanation for the reduction in false alarms when the Laplacian noise model is assumed may lie in the large probability that this density function assigns to tail events, when compared to the Gaussian density function. There is therefore less need to rely on post-processing, such as background masking, etc. to reduce the number of possible false alarms.

Finally, we note that though our model in the B matrix of Eq. (10) allows for incorporating physiological effects, etc., we did not include them. If correctly learned, including these effects would enhance performance, but we do not expect inclusion of such effects would change the conclusion as far as the relative performance of the detectors discussed above is concerned.

9. Learning, Signal Design, and Generalizations

In this section we discuss several generalizations and implications of our choice of subspace and noise models. In Sec. 7.4 we presented, for the case of binary valued subspaces, detectors with a special structure that is computationally efficient and easier to employ for analysis such as signal design. In Secs. 9.1 and 9.2, we summarize the relevance of binary valued subspaces to two additional problems relevant to fMRI: subspace learning,[26] and signal design,[28] respectively. We also discuss other generalizations in Sec. 9.3. Specifically, in Sec. 9.3.1, we generalize the response model of Sec. 4.2, and in Sec. 9.3.2, we introduce the problem of gain constraints. Robust detection in the presence of additional types of constraints on signals and interference are summarized in Sec. 9.4. Finally, additional observations are given in Sec. 9.5.

9.1. *Learning: Laplacian noise and binary subspaces*

In Secs. 6, we estimated or identified the shift parameter of the stimulus response model. This was a parametric approach to determining the subspace, meaning that the form of the subspace, in our case the binary valued form presented in Sec. 4.2, was assumed *a priori*, and the only unknowns were the shift parameters.

The binary form of the subspace relies on an essential result regarding learning in the presence of Laplacian noise,[26,27] which we restate without derivation. Specifically, in an optimal subspace learning problem formulation no parametric dependency constraints are explicitly imposed on either the form or the orientation of the subspace, meaning on the form of the vector a in Eq. (10). With such a formulation, when the noise is Laplacian, the form of the subspace that obtains is exactly the form that we have assumed. In other words, the optimal subspace is one spanned by a vector whose elements assume binary values of equal magnitude only, such as $+1$ and -1.

The implication is that there is no loss of generality when assuming a binary form, provided the noise is Laplacian. This result justifies the binary model of Sec. 4.2, where only the shift parameters need to be estimated.

If the noise is not Laplacian, however, then the binary form of the subspace is no longer necessarily optimal. One can still use the data to determine the entire subspace's form, and not just the shift parameter vectors λ associated with the assumed structure of the subspace. Such subspace learning is applicable when sufficient data is available. In the absence of sufficient data, subspace learning requires assumptions of spatial stationarity of the response over a suitable neighborhood, thus enabling estimation of an average response subspace a of Eq. (10). As mentioned in the introduction, variability of the response magnitude, varying time delays, adaptation over time of the subject to the stimulus, and the transient nature of the response involving propagation from loci to loci, are all aspects of non stationarity that add to the difficulties of learning.

We provide a formulation of the subspace learning problem below.

9.1.1. *Maximum likelihood subspace learning*

Consider the *multidimensional* subspace model

$$x = A\theta + B\phi + n \tag{28}$$

where both A and B may have more than one column. We assume that B is known, but A, θ, and ϕ are unknowns to be determined. Using the maximum likelihood approach for a sequence of N observations $x_i, i = 1, \ldots, N$, we have under the assumption that the noise vectors n_i are independent,

$$\max_{A} \max_{\theta_1,\ldots,\theta_N} \max_{\phi_1,\ldots,\phi_N} \prod_{i=1}^{N} p(x_i - A\theta_i - B\phi_i). \tag{29}$$

The approach of Eq. (29) is the conventional one, where a function of the residuals $x - A\theta_i - B\phi_i, i = 1, \ldots, N$ is minimized. An alternate approach is to estimate $A_\perp$, the space orthogonal to the signal space A by minimizing, instead of maximizing, the likelihood function with respect to $A_\perp$,

$$\min_{A_\perp} \max_{\theta_{\perp 1}, \ldots, \theta_{\perp N}} \max_{\phi_1, \ldots, \phi_N} \prod_{i=1}^{N} p(x_i - A_\perp \theta_{\perp i} - B\phi_i). \tag{30}$$

We now restrict our attention to the Generalized Gaussian class of noise density functions of Eq. (21). Again, since no component of A residing in the interference space spanned by the columns of B can be estimated, we will restrict our search for $A_\perp$ in the space orthogonal to the combined column space of $[A, B]$. Thus, after taking the log, we have the log-likelihood maximization with respect to $A_\perp$ and minimization with respect to the $\theta_{\perp i}$'s, and the ϕ_i's, $i = 1, \ldots, N$,

$$\max_{A_\perp} \min_{\theta_{\perp 1}, \ldots, \theta_{\perp N}} \min_{\phi_1, \ldots, \phi_N} \sum_{i=1}^{N} \|(x_i - A_\perp \theta_{\perp i} - B\phi_i\|_p^p. \tag{31}$$

We will refer to the two minimization problems in Eq. (31) as the *projection* problem, while the maximization with respect to $A_\perp$ will be denoted the *learning* problem.

Defining the residual

$$\eta_i = x_i - A_\perp \theta_{\perp i} - B\phi_i$$

and choosing matrix $\bar{A}$ such that

$$\text{span}(\bar{A}) = \text{Null Space of } [A_\perp, B]$$

we are led to an equivalent problem

$$\max_{\bar{A}} \min_{\eta_1, \ldots, \eta_N} \sum_{i=1}^{N} \|\eta_i\|_p^p \tag{32}$$

subject to the constraints

$$\bar{A}'\eta_i = \bar{A}'x_i, \quad i = 1, \ldots, N \tag{33}$$
$$\bar{A}'B = 0. \tag{34}$$

Note that the above equation implies that

$$\bar{A} = (I - P_B)A \tag{35}$$

where $P_B = B(B'B)^{-1}B'$ is the projection matrix onto the range space of B. We consider this problem further in Ref. 27.

Finally, as mentioned earlier, when the noise is Laplacian with $p = 1$, and the subspace $\bar{A}$ is one dimensional, the maximizing subspace vector consists of binary elements $+1$ or -1.[26]

9.2. *Signal design: implications of binary subspaces*

The special forms and the parametric dependencies that the detectors exhibit when the signal of interest is binary has implications for signal design. Specifically, we see from the previous section that the likelihood ratio expressions for all four detectors is a function of K_+ and K_-, or the size of the signal's positive and negative "parts" as explained in Sec. 7.4, but not of their distribution in time. A question then arises as to what is the effect of the parameters K_+ and K_- on the detector's performance. We first consider the case when the noise is Gaussian. As mentioned in Sec. 5, the detector $\lambda_{2,ru}$ is equivalent in performance to a t statistic with $K - 2$ degrees of freedom and noncentrality parameter μ given by

$$\mu = \frac{2\theta}{\sigma} K \sqrt{\epsilon(1 - \epsilon)}$$

where $\epsilon = K_+/K$ is a measure of the asymmetry of the binary signal a. The noncentrality parameter μ, and hence the probability of detection for any given false alarm probability, is maximized with $\epsilon = 1/2$ when the signal is symmetric, meaning $K_+ = K_-$. For this reason, a symmetric signal offers superior ROC performance. This can be explained by the fact that in the Gaussian case, it is only possible to estimate the component of the signal residing in the space orthogonal to that of the bias. When the signal is symmetric, it resides entirely in that space, since $a'b = 0$, and the bias does not degrade the estimate. In the case of an arbitrary Generalized Gaussian noise, meaning any shape parameter p, the objective would be to maximize the residual ρ_0's p norm in terms of ϵ. This maximum obtains for $\epsilon = 1/2$, meaning when the signal is symmetric as well. Figures 5 and 6 illustrates, for Laplacian noise, the detector's performance as a function of the symmetry index ϵ.

Finally, we note that there are sufficient degrees of freedom for a given ϵ to choose the signal waveform, and these degrees of freedom can be exploited for performance robustness with respect to expected interference sources that are only vaguely characterized. For instance, aperiodicity may be employed when interference is known to be periodic, as is the case for instance for the cardiac related interference effects when performing fMRI.

9.3. *Response model generalizations*

9.3.1. *A more general response model*

We provide a more general response model, and describe how we arrive at the response model of Sec. 4. The general model accommodates the temporal variability of not only the duration periods of on and off phases, but also of the gains and response time shifts that occur at a time scale larger than the time scale of stimulus phase variation. In short, the model can be used for a finer characterization of the voxel's responsiveness. In particular, we can use it for detecting anomalous responses, the presence of low frequency drift effects, etc. On the other hand, more detailed models lead to over-parametrization and associated problems, as well as
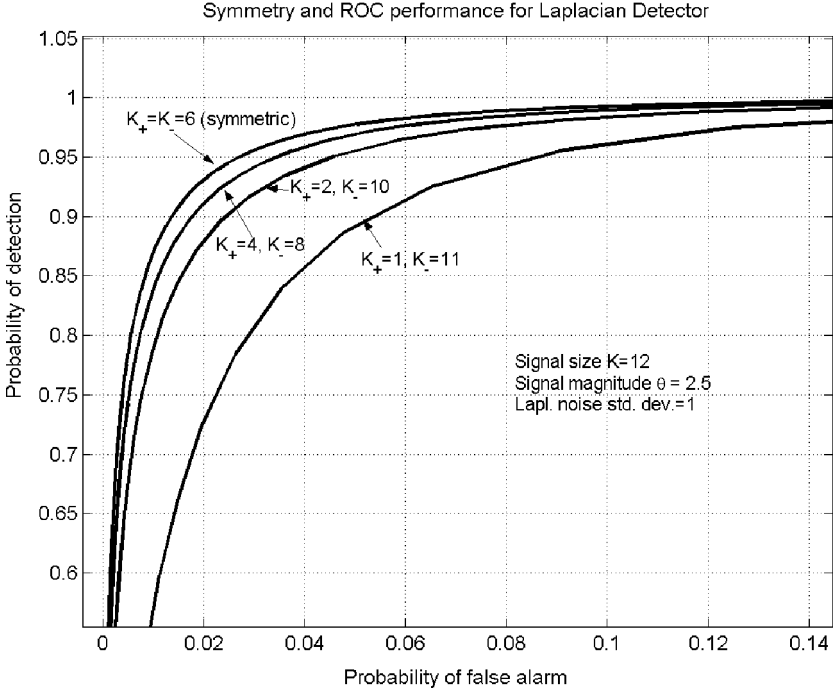
Fig. 5. *ROC curves for various levels of asymmetry.* A symmetric signal $(K_+ = K_- = 6)$ provides superior ROC performance, as shown here for the CFAR Laplacian detector $\lambda_{1,ru}$, but note that symmetric signals are not necessarily periodic. ©2004 IEEE.[28]

increased false alarms, alleviation of which by ad hoc and non-ad hoc measures becomes a difficult task. We first generalize the unit step response model of Eq. (3) so as to accommodate phase varying gains. Specifically,

$$r[k] = r[1] + \theta\,[k_0]\,h[k - k_0 - \lambda[k_0]]. \tag{36}$$

Then, for input stimuli of the type shown in Fig. 1, usage of Eq. (36) leads to

$$r[k] = r[1] + \sum_{j=2}^{J} \theta[j]\delta[j]h[k - k_{\mathrm{tr}}[j] - \lambda[j]], \quad k = 1, \ldots, K. \tag{37}$$

where $\delta[j]$ is as defined in Eq. (7). We assume the initial phase $j = 1$ is an off phase. We note that the subscript i representing the voxel is dropped. The above stimulus response model allows for piecewise constant response to an on-off stimulus pattern with time varying on and off phase durations. More precisely, we can simplify the time dependency of Eq. (37) in terms of a smaller number of phase dependent variables as follows

$$r[k] = \vartheta[j], \quad k \in \{j\text{th phase of response}\} \tag{38}$$
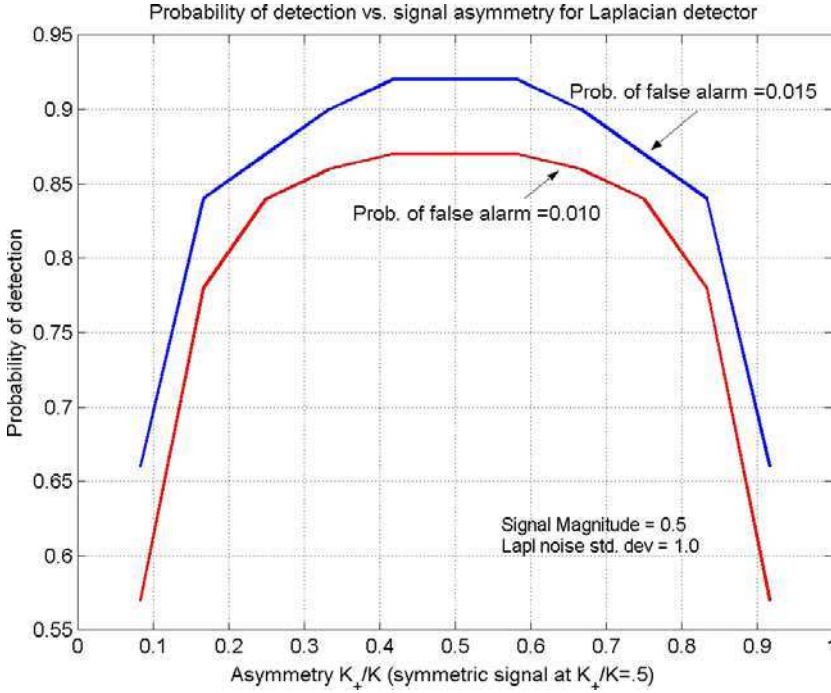
Fig. 6.   *Probability of detection vs. level of symmetry in signal design.* A symmetric signal ($K_+ = K_-$) provides superior ROC performance. The results shown are for a Laplacian CFAR detector, or $\lambda_{1,ru}$. ©2004 IEEE.[28].

where

$$\vartheta[j] = \vartheta[1] + \sum_{\ell=2}^{j} \theta[\ell]\delta[\ell], \quad j = 2, \ldots, J, \quad \vartheta[1] = r[1]. \tag{39}$$

The jth phase of response is identified once time shifts are determined. In vector notation, we can rewrite Eq. (38) as

$$r = A(\lambda)\theta \tag{40}$$

where r and θ are now the response and gain *vectors* of length J, respectively. Thus,

$$\theta = [r[1], \theta[2], \ldots, \theta[J]]'. \tag{41}$$

The $K \times J$ response matrix A is given by the row partitioning

$$A = [A'[1], \ldots, A'[J]]' \tag{42}$$

where $A[j], j = 1, \ldots, J$, is of dimension $T[j] \times J$ and is given by

$$A[j] = e_b[j]\chi[j]$$

with $e_b[j]$ a column vector of ones of length $T[j]$ as given by Eq. (6), and $\chi[j]$ a row vector of length J whose ℓth element is given by

$$\chi[\ell] = \begin{cases} 1 & \text{if } \ell = 1 \\ (-1)^\ell & \text{if } 2 \le \ell \le j \\ 0 & \text{otherwise.} \end{cases}$$

The above model implicitly assumes that the shifts $\lambda_i(j)$, $j = 2, \ldots, J$ are small relative to the duration of neighboring stimulus phases $j - 1$ and j, such that the following constraints are not violated:

(1) *Time shift localization.* We constrain each stimulus response time shift to occur within the time interval of adjacent stimulus phases. Specifically, we require for $j = 2, \ldots, J$,

$$k_{\text{tr}}[j - 1] < k_{\text{tr}}[j] + \lambda[j] < k_{\text{tr}}[j],$$

where, by definition $k_{\text{tr}}[1] = 1$ and $k_{\text{tr}}[J + 1] = K$.

(2) *Time shift precedence.* This constraint is expressed as

$$k_{\text{tr}}[j - 1] + \lambda[j - 1] < k_{\text{tr}}[j] + \lambda_i[j]$$

where, by definition $\lambda[1] = 0$.

We note that the relaxation of constraints on response model parameters *generally* leads to increases in both the probabilities of detection and false alarm. False alarms can arise due to the presence of non-stimulus drift effects with frequencies lower than those of the stimulus response. This can lead to temporal variability in the response, particularly in the gain parameters. Relaxation of constraints can lead to including these effects in the estimation of the response model parameters. If unmodeled low frequency effects are bounded, however, then the probability of detection would increase without an equally large increase in the probability of false alarm if we properly specify the bounded gain magnitude variations between phases. Limiting the relaxation of constraints on the parameters to the extent where it is physically justifiable is crucial to restricting false alarms, particularly if the dimension of the response subspace is greater than one.

9.3.2. *Response models with gain constraints*

The stimulus response models of Eqs. (37) and (40) can be specialized to give simpler models, including the one we describe in Sec. 4 and implement in Sec. 8. Specifically, we can impose the following constraints on the *response gains*:

(1) *Polarity.* The polarity constraints require a positive gain, or $\theta(j) \ge 0$, which implies from Eq. (37) that the on and off stimuli respectively lead to an increase and decrease in the response.

(2) *Time invariant response gains.* There are two cases, one where symmetry obtains between on and off phase gains, meaning the on response gains are

equal to the off response gains, and the other case, where the on and off phase response gains are asymmetric, meaning they are unequal.

(a) *Symmetric on and off gain responses*:
The gain's magnitude, which can be different for different voxels, are nevertheless constrained to be the same from one phase to the next, or $\theta[j] = \bar{\theta}$. This implies from Eq. (37) that the overall amplitude of response is equal to $r[1]$ for off phases and $\bar{\theta} + r[1]$ for on phases.

(b) *Asymmetric on and off gain responses*:
If we allow for unequal amplitudes for the on and off response phases with $\theta[j] = \bar{\theta}$ for on phases, and $\theta[j] = \bar{\theta} + m_r$ for off phases, then an additional term appears in Eq. (4) to give

$$r = a(\lambda)\theta + e_r m_r \tag{43}$$

where e_r is given by

$$e_r = [e_r[1], \ldots, e_r[J]]$$

and represents a piecewise constant staircase function, with

$$e_r[j] = -\text{int}\left[\frac{j-1}{2}\right] e_b[j]$$

where $\text{int}[w]$ is the largest integer smaller than or equal to w.

Under the assumptions of polarity and symmetric response time invariant gains, we have the response model of Eq. (4). Under the assumptions of polarity and asymmetric response time invariant gains, we are led to an additional term to the same response model involving a piecewise constant staircase term. It should be noted that the staircase term is different from a secular ramp term that may arise from non-stimulus effects.

9.4. *Robust detection with additional model constraints*

We have considered in this chapter robust detection that takes into account possible presence of interference where we make an assumption about the relative orientation of the subspace in which interferencce can lie with respect to the response subspace. Alternate type of partial characterization of interference may be available and incorporated in the design of the robust detector. One such example of partial characterization is bounds on the magnitude of the inteference that may reside in the signal subspace as well. Algorithms for robust detection in the presence of constrained magnitude interference have been earlier developed and successfully demonstrated on the failure detection and identification of flight sensors for NASA F8 fly-by-aircraft,[24] and to the M.I.T. nuclear reactor.[95] Further developments are found in Ref. 94. Additional more general design considerations in the context of min-max approach to robust detection to uncertainties in the different components of the model have been developed,[26] and in particular for signal subspaces.[93]

9.5. *Additional observations*

Potentially beneficial modifications of our approach include:

- *Noise density function.* Regarding the choice of density functions, we have seen in Sec. 7 that the Gaussian and the Laplacian are two members of a larger family of parametric distributions. The effect of this parameter value on detection is discussed in Ref. 25.
- *A posteriori model learning.* Once the detection test results are available, they can be used for improving the models used. For instance the matrices $a(\lambda)$, B, and other parameters in Sec. 4 can be updated. Moreover, this learning could be used to construct priors for follow on testing and processing.

The above generalizations would not alter the proposed approach in a significant way. Regarding *a posteriori* model learning (the second bullet above), we note that in a general context functional MRI is concerned with how various parts of the brain function. One challenge with which we deal here is the problem of estimating the probability that each voxel in the brain responds to a particular class of stimuli. Another challenge, which we do not directly address, is the estimation of the temporal and spatial nature of the response variation. This is the focus of other work such as that of Solo.[103,104] The two problems just mentioned are related in that the solution of one is of help with the other, and one can iterate between the two. For instance, the relative amplitude characteristics of on and off responses can be further examined for the presence of anomalous response, low drift, etc, using a more general response model than used here. Since we focus in this work on the functional characterization aspect of the problem, we only examine the difference between the average of responses during on and off phases as a first step of what would be an iterative procedure.

10. Conclusion

The fMRI detection problem is particularly challenging due to the low signal to noise and interference ratio at hand. To address this challenge, it is important to capture the variations in the transient and steady state characteristics of the response both in time and from voxel to voxel. The chapter addresses the accompanying modeling, estimation, and detection challenges in a hierarchical fashion that addresses separately the transient and steady state characteristics. This is facilitated by the adoption of a simple response model parametrized by the response gains and the time shifts, usually delays. The gains are intended to capture the response's steady state behavior, while the time shifts capture the transport delays and the transient behavior. We estimate both the gains, which vary from voxel to voxel but are constant for each voxel, and the shifts, which vary spatially from voxel to voxel *and temporally* from phase to phase of stimulus.

The transient characteristics are learned at each voxel using the curve evolution approach to the image segmentation problem, which provides temporal characterization and detects the presence of low contrast edges that may be blurred

in the images. The approach permits the exploitation of the spatial correlation of transient characteristics. The transient characteristics are estimated first prior to learning the steady state characteristics, and implementing a robust matched subspace detection test based on a Laplacian noise model. The test functionally characterizes the response to produce the a posteriori probability that each voxel is responding.

We note that the robust detection test we formulate enables the derivation, based on likelihood ratio considerations, of a decision statistic for the large class of Generalized Gaussian density functions. The new statistic is analogous to the t, F-statistic used for subspace based detection when the noise is Gaussian.

The results indicate that accounting for response time shifts in a simple response model produces enhanced images of the response's likelihood; more responding voxels are attributed a higher probability of stimulus response. Furthermore, the use of the time shift estimates in a robust detection test based on a Laplacian noise model reduces the number of potential false alarms without having to resort to spatial smoothing or background masking. In our view, the proposed features are an attempt to simultaneously address the two critical challenges of enhanced detectability and reduced false alarms.

The chapter also provides various implications of Laplacian noise models. One such implication is that in the presence of such noise, the simple response model adopted implies no loss of generality as it is the optimal choice. Another implication relates to the simple form of the Laplacian detector that enables analysis on the design aspect of stimulus signal.

Finally, the fMRI problem motivated the newly developed robust non-Gaussian detection approach that offered superior detection and false alarm performance when compared to traditional Gaussian approaches. This new approach is a theoretical and practical contribution to signal detection in other applications as well, such as radar and sonar.

Acknowledgments

References

1. R. Adolphs, H. Damasio, D. Tranel and A. R. Damasio, Cortical systems for the recognition of emotion in facial expressions, *J. Neurosci.* **16**(23) (1996) 7678–87.
2. H. Akaike, A new look at statistical model identification, *IEEE Trans. on Autom. Contr.* **19** (1974) 716–723.

3. B. A. Ardekani, J. Kershaw, K. Kashikura and I. Kanno, Activation detection in functional MRI using subspace modeling and maximum likelihood estimation, *IEEE Trans. Med. Imag.* **18**(2) (1999) 101–114.

4. P. A. Bandetini, A. Jemanowicz, E. C. Wong and J. S. Hyde, Processing strategies for time-coarse data sets in functional MRI of the human brain, *Magn. Reson. Med.* **30** (1993) 161–173.

5. P. A. Bandettini, E. C. Wong, R. S. Hinks, R. S. Tikofsky and J. S. Hyde, Time course EPI of human brain function during task activation, *Magn. Reson. Med.* **25**(2) (1992) 390–7.

6. J. W. Belliveau, D. N. Kennedy, R. C. McKinstry, B. R. Buchbinder, R. M. Weisskoff, M. S. Cohen, J. M. Vevea, T. J. Brady and B. R. Rosen, Functional mapping of the human visual cortex by magnetic resonance imaging, *Science* **254** (1991) 716–719.

7. Y. Benjamini and Y. Hochberg, Controlling the false discovery rate: a practical and powerful approach to multiple testing, *Journal of the Royal Statistical Society*, Series B **57** (1995) 289–300.

8. S. Bookheimer, Functional MRI of language: new approaches to understanding the cortical organization of semantic processing, *Annu. Rev. Neurosci.* **25** (2002) 151–88.

9. H. C. Breiter, R. L. Gollub, R. M. Weisskoff, D. N. Kennedy, N. Makris, J. D. Berke, J. M. Goodman, H. L. Kantor, D. R. Gastfriend, J. P. Riorden, R. T. Mathew, B. R. Rosen and S. E. Hyman, Acute effects of cocaine on human brain activity and emotion, *Neuron.* **19**(3) (1997) 591–611.

10. H. C. Breiter, I. Aharon, D. Kahneman, A. Dale and P. Shizgal, Functional imaging of neural responses to expectancy and experience of monetary gains and losses, *Neuron.* **30**(2) (2001) 619–39.

11. R. L. Buckner, W. Koutstaal, D. L. Schacter, A. M. Dale, M. Rotte and B. R. Rosen, Functional-anatomic study of episodic retrieval using fMRI. I. Retrieval effort versus retrieval success, *Neuroimage* **7**(3) (1998) 151–62.

12. R. L. Buckner and M. E. Wheeler, The cognitive neuroscience of remembering. *Nat. Rev. Neurosci.* **2**(9) (2001) 624–34.

13. E. Bullmore, M. Brammer, S. C. R. Williams, S. Rabe-Hesketh, N. Janot, A. David, J. Mellers, R. Howard and P. Sham, Statistical methods of estimation and inference for functional MR image analysis, *Magn. Reson. Med.* **35** 261–277.

14. R. B. Buxton and L. R. Frank, A model for the coupling between cerebral blood flow and oxygen metabolism during neural stimulation, *J. Cereb. Blood Flow Metab.* **17**(1) (1997) 64–72.

15. R. B. Buxton, *An Introduction to Functional Magnetic Resonance Imaging: Principles and Techniques* (Cambridge University Press, Cambridge, 2001).

16. J. I. Chen, B. Ha, M. C. Bus hnell, B. Pike and G. H. Duncan, Differentiating noxious- and innocuous-related activation of human somatosensory cortices using temporal analysis of fMRI, *J. Neurophysiol.* **88**(1) (2002) 464–74.

17. L. Cornette, P. Dupont, A. Rosier, S. Sunaert, P. Van Hecke, J. Michiels, L. Mortelmans and G. A. Orban, Human brain regions involved in direction discrimination, *J. Neurophysiol.* **79**(5) (1998) 2749–65.

18. R. W. Cox, A. Jemanowicz and J. S. Hyde, Real-time functional magnetic resonance imaging, *Magn. Reson. Med.* **33** (1995) 230–236.

19. K. D. Davis, Neurophysiological and anatomical considerations in functional imaging of pain, *Pain* **105**(1–2) (2003) 1–3.

20. J. B. Demb, G. M. Boynton and D. J. Heeger, Brain activity in visual cortex predicts individual differences in reading performance *Proc. Natl. Acad. Sci. USA* **94**(24) (1997) 13363–6. 21.

21. J. B. Demb, G. M. Boynton and D. J. Heeger, Functional magnetic resonance imaging of early visual pathways in dyslexia, *J. Neurosci.* **18**(17) (1998) 6939–51.

22. J. B. Demb, J. E. Desmond, A. D. Wagner, C. J. Vaidya, G. H. Glover and J. D. Gabrieli, Semantic encoding and retrieval in the left inferior prefrontal cortex: a functional MRI study of task difficulty and process specificity, *J. Neurosci.* **15**(9) (1995) 5870–8.

23. M. Desai, R. Mangoubi, J. Shah, W. Karl, H. Pien, A. Worth and D. Kennedy, Functional MRI activity characterization using curve evolution, *IEEE Trans. on Med. Imag.* Nov., 2002, pp. 1402–1413.

24. M. Desai and A. Ray, Failure detection system design-theory and application, *Proc. of 20th Conf. on Decision and Control (CDC)*, San Diego, California (December 1981).

25. M. Desai and R. Mangoubi, Robust Gaussian and non-Gaussian matched subspace detection, *IEEE Trans. on Signal Processing* **51**(12) (2003) 3115–3127.

26. M. Desai and R. Mangoubi, Subspace learning and innovation characterization in generalized Gaussian noise, in *Proc. 36th Asilomar Conf. Signals, Systems, and Computers* (November, 2003).

27. M. Desai and R. Mangoubi, Non-Gaussian subspace learning in the presence of interference, *Proc. IEEE Sensor, Array, and Multichannel Sig. Proc. Workshop*, Barcelona, Spain (18–21 July 2004).

28. M. Desai and R. Mangoubi, Non-Gaussian robust matched subspace detectors and signal design analysis for a class of separable subspaces, *Proc. IEEE Sensor, Array, and Multichannel Sig. Proc. Workshop*, Barcelona, Spain (18–21 July 2004).

29. M. Desai, R. Mangoubi, D. Kennedy, J. Shah, W. Karl, N. Markis and A. Worth, Diffusion tensor model based smoothing and segmentation, *IEEE Conf. Medical imaging*, Washington DC (Aug 2002).

30. M. Desai, J. Deckert and J. Deyst, Dual-sensor failure identification using analytic redundancy, *Journal of Guidance and Control,* May–June 1979.

31. X. Descombes, F. Kruggel and D. Y. von Cramon, fMRI signal restoration using a spatio-temporal Markov random field preserving transitions, *NeuroImage* **8**(4) (1998) 340–349.

32. M. D'Esposito, L. Y. Deouell and A. Gazzaley, Alterations in the BOLD fMRI signal with ageing and disease: A challenge for neuroimaging, *Nat. Rev. Neurosci.* **4**(11) (2003) 863–72.

33. M. D'Esposito, E. Zarahn, G. K. Aguirre and B. Rypma, The effect of normal aging on the coupling of neural activity to the bold hemodynamic response, *Neuroimage* **10**(1) (1999) 6–14.

34. A. Devor, A. K. Dunn, M. L. Andermann, I. Ulbert, D. A. Boas and A. M. Dale, Coupling of total hemoglobin concentration, oxygenation, and neural activity in rat somatosensory cortex, *Neuron.* **39**(2) (2003) 353–9.

35. S. Durston, A review of the biological bases of ADHD: what have we learned from imaging studies? *Ment. Retard. Dev. Disabil. Res. Rev.* **9**(3) (2003) 184–95.

36. G. F. Eden, J. W. VanMeter, J. M. Rumsey, J. M. Maisog, R. P. Woods and T. A. Zeffiro, Abnormal processing of visual motion in dyslexia revealed by functional brain imaging, *Nature* **382**(6586) (1996) 66–9.

37. K. Elgar and R. Campbell, Annotation: the cognitive neuroscience of face recognition: implications for developmental disorders, *J. Child Psychol Psychiatry* **42**(6) (2001) 705–17.

38. B. Everitt and E. T. Bullmore, Mixture model maping of brain activation in functional magnetic resonance images, *Human Brain Mapping* **7** (1999) 1–14.

39. P. T. Fox and M. E. Raichle, Focal physiological uncoupling of cerebral blood flow and oxidative metabolism during somatosensory stimulation in human subjects, *Proc. Natl. Acad. Sci. USA* **83**(4) (1986) 1140–4.

40. L. R. Frank, R. B. Buxton and E. C. Wong, Probabilistic analysis of functional magnetic resonance imaging data, *Mag. Reson. Med.* **39** (1998) 132–148.

41. A. D. Friederici, Syntactic, prosodic, and semantic processes in the brain: evidence from event-related neuroimaging, *J. Psycholinguist Res.* **30**(3) (2001) 237–50.

42. K. J. Friston, P. Jezzard and R. Turner, Analysis of functional MRI time-series, *Human Brain Mapping* **1** (1994) 153–171.

43. K. J. Friston, A. P. Holmes, K. J. Worsley, J.-B. Poline, C. D. Frith and R. S. J. Frackowiak, Statistical parametric maps in functional imaging: a general linear approach, *Human Brain Mapping* **2** (1995) 189–210.

44. C. R. Genovese, D. C. Noll and W. F. Eddy, Estimating test-retest reliability in fmri I: statistical methodology, *Mag. Resonance in medicine* **38** (1997) 497–507.

45. C. R. Genovese, N. A. Lazar and T. Nichols, Thresholding of statistical maps in functional neuroimaging using the false discovery rate, submitted.

46. G. H. Glover, Deconvolution of impulse response in event-related BOLD fMRI, *NeuroImage* **9** (1999) 416–429.

47. P. Gray, *Psychology* (Worth Publishing, New York, 1994).

48. N. Hadjikhani, D. R. M. Sanchez, O. Wu, D. Schwartz, D. Bakker, B. Fischl, K. K. Kwong, F. M. Cutrer, B. R. Rosen, R. B. Tootell, A. G. Sorensen and M. A. Moskowitz, Mechanisms of migraine aura revealed by functional MRI in human visual cortex, *Proc. Natl. Acad. Sci. USA* **98**(8) (2001) 4687–92.

49. N. Hadjikhani, A. K. Liu, A. M. Dale, P. Cavanagh and R. B. Tootell, Retinotopy and color sensitivity in human visual cortical area V8, *Nat. Neurosci.* **1**(3) (1998) 235–41.

50. N. Hadjikhani and B. de Gelder, Seeing fearful body expressions activates the fusiform cortex and amygdala, *Curr. Biol.* **13**(24) (2003) 2201–5.

51. T. S. Hale, A. R. Hariri and J. T. McCracken, Attention-deficit/hyperactivity disorder: perspectives from neuroimaging, *Ment. Retard. Dev. Disabil. Res. Rev.* **6**(3) (2000) 214–9.

52. E. Halgren, A. M. Dale, M. I. Sereno, R. B. Tootell, K. Marinkovic and B. R. Rosen, Location of human face-selective cortex with respect to retinotopic areas, *Hum. Brain Mapp.* **7**(1) (1999) 29–37.

53. S. J. Hanson and B. M. Bly, The distribution of bold susceptibility effects in the brain is non-Gaussian, *NeuroReport* **12** (2001) 1971–76.

54. S. Heckers, Neuroimaging studies of the hippocampus in schizophrenia, *Hippocampus* **11**(5) (2001) 520–8.

55. D. J. Heeger, Linking visual perception with human brain activity, *Curr. Opin. Neurobiol* **9**(4) (1999) 474–9.

56. R. I. Henkin, L. M. Levy and C. S. Lin, Taste and smell phantoms revealed by brain functional MRI (fMRI), *J. Comput. Assist. Tomogr.* **24**(1) (2000) 106–23.

57. R. N. Henson, M. D. Rugg, T. Shallice, O. Josephs and R. J. Dolan, Recollection and familiarity in recognition memory: an event-related functional magnetic resonance imaging study, *J. Neurosci.* **19**(10) (1999) 3962–72.

58. V. Hesselmann, W. O. Zaro, C. Wedekind, T. Krings, O. Schulte, H. Kugel, B. Krug, N. Klug and K. J. Lackner, Age related signal decrease in functional magnetic resonance imaging during motor stimulation in humans, *Neurosci. Lett.* **308**(3) (2001) 141–4.

59. D. E. Hill, R. A. Yeo, R. A. Campbell, B. Hart, J. Vigil and W. Brooks, Magnetic resonance imaging correlates of attention-deficit/hyperactivity disorder in children, *Neuropsychology* **17**(3) (2003) 496–506.

60. J. C. Horton and W. F. Hoyt, The representation of the visual field in human striate cortex. A revision of the classic Holmes map, *Arch. Ophthalmol.* **109**(6) (1991) 816–24.

61. X. Hu, T. H. Le, T. Parrish and P. Erhard, Retrospective estimation and correction of physiological fluctuation in functional MRI, *Magn. Res. Med.* **34** (1995) 201–212.

62. P. Huber, *Robust Statistics* (Wiley, New York, 1981).

63. A. K. Jain, *Fundamentals of Digital Image Processing* (Prentice Hall, Englewood Cliffs, N.J. 1989).

64. P. Jezzazd, P. Matthews and S. Smith, *Functional MRI: An Introduction to Methods*, International Institute for Strategic Studies, New Ed edition (2003).

65. N. Kanwisher and E. Wojciulik, Visual attention: insights from brain imaging, *Nat. Rev. Neurosci.* **1**(2) (2000) 91–100.

66. S. A. Kassam, *Signal Detection in Non-Gaussian Noise* (Springer Verlag, 1990).

67. J. S. Kauer and J. White, Imaging and coding in the olfactory system, *Annu. Rev. Neurosci.* **24** (2001) 963–79.

68. J. Kershaw, B. A. Ardekani and I. Kanno, Application of Bayesian inference to fMRI data analysis, *IEEE Trans. Med. Imag.* **18**(12) (1999) 1138–1153.

69. S. G. Kim and K. Ugurbil, Functional magnetic resonance imaging of the human brain, *J. Neurosc. Meth.* **74** (1997) 229–243.

70. K. K. Kwong, J. W. Belliveau, D. A. Chesler, I. E. Goldberg, R. M. Weisskoff, B. P. Poncelet, D. N. Kennedy, B. E. Hoppel, M. S. Cohen and R. Turner, Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation, *Proc. Natl. Acad. Sci. USA* **89**(12) (1992) 5675–79.

71. N. Lange and S. L. Zeger, Non-linear Fourier time series analysis for human brain mapping by functional magnetic resonance imaging, *J. Appl. Statistics* **16**(1) (1997) 1–29.

72. T. H. Le and X. Hu, Retrospective estimation and correction of physiological artifacts in fMRI by direct extraction of physiological activity from MR data, *Magn. Reson. Med.* **35** (1996) 290–298.

73. J. E. LeDoux, Emotion circuits in the brain, *Annu. Rev. Neurosci.* **23** (2000) 155–84.

74. L. M. Levy, R. I. Henkin, C. S. Lin and A. Finley, Rapid imaging of olfaction by functional MRI (fMRI): identification of presence and type of hyposmia, *J. Comput. Assist. Tomogr.* **23**(5) (1999) 767–75.

75. L. M. Levy, R. I. Henkin, C. S. Lin, A. Hutter and D. Schellinger, Odor memory induces brain activation as measured by functional MRI, *J. Comput. Assist. Tomogr.* **23**(4) (1999) 487–98.

76. R. Luypaert, S.Boujraf, S. Sourbron and M. Osteaux, Diffusion and perfusion MRI: basic physics, *Eur. J. Radiol.* **38**(1) (2001) 19–27.

77. R. J. Maddock, The retrosplenial cortex and emotion: new insights from functional neuroimaging of the human brain, *Trends Neurosci.* **22**(7) (1999) 310–6.

78. R. Mangoubi, *Robust Estimation and Failure Detection: A Concise Treatment* (Springer Verlag, 1998).

79. D. W. McRobbie, E. A. Moore Graves, M. J. and M. R. Prince, *MRI: From Picture to Proton* (Cambridge University Press, New York, 2003).

80. C. T. W. Moonen and P. A. Bandetini, *Functional MRI* (Springer Verlag, 2000).

81. F. Y. Nan and R. D. Nowak, Generalized likelihood ratio detection for fMRI using complex data, *IEEE Trans. Med. Imag.* **18**(4) (1999) 320–329.

82. S. Ogawa and T. M. Lee, Magnetic resonance imaging of blood vessels at high fields: in vivo and in vitro measurements and image simulation, *Magn. Reson. Med.* **16**(1) (1990) 9–18.

83. S. Ogawa, T. M. Lee, A. R. Kay and D. W. Tank, Brain magnetic resonance imaging with contrast dependent on blood oxygenation, *Proc. Nat'l Acad. Sci.* **87** (1990) 9868–9872.

84. S. Ogawa, D. W. Tank, R. Menon, J. M. Ellermann, S. G. Kim, H. Merkl and K. Ugurbil, Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging, *Proc. Natl. Acad. Sci. USA* **89** (1992) 5951–5.

85. S. Ogawa, D. W. Tank, R. Menon, J. M. Ellermann, S. G. Kim, H. Merkle and K. Ugurbil, Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields, *Magn. Reson. Med.* **14**(1) (1990) 68–78.

86. E. D. Palmer, H. J. Rosen, J. G. Ojemann, R. L. Buckner, W. M. Kelley and S. E. Petersen, An event-related fMRI study of overt and covert word stem completion, *Neuroimage* **14**(1 Pt 1) (2001) 182–93.

87. K. L. Phan, T. Wager, S. F. Taylor and I. Liberzon, Functional neuroanatomy of emotion: a meta-analysis of emotion activation studies in PET and fMRI, *Neuroimage* **16**(2) (2002) 331–48.

88. G. D. Pearlson, Neurobiology of schizophrenia, *Ann. Neurol.* **48**(4) (2000) 556–66.

89. M. T. Posamentier and H. Abdi, Processing faces and facial expressions, *Neuropsychol Rev.* **13**(3) (2003) 113–43.

90. H. Pien, M. Desai and J. Shah, Segmentation of MR images using curve evolution and prior information, *Int'l J. Pattern Recognition and Artificial Intelligence* **11**(8) (1997) 1233–1245.

91. P. L. Purdon and R. M. Weisskoff, Effect of temporal autocorrelation due to physiological noise and stimulus paradigm on pixel-level false-positive rates in fMRI, *Human Brain Mapping* **6**(4) (1998) 239–249.

92. P. L. Purdon, V. Solo, E. Brown, R. Buckner, M. Rotte and R. M. Weisskoff, fMRI noise variability across subjects and trials: insights for noise estimation and methods, *NeuroImage* **7** (1999) S617.

93. S. Ramprashad, T. W. Parks and R. Shenoy, Signal modeling and detection using cone classes, *IEEE Trans. on Signal Processing* **44**(2) (1996).

94. A. Ray and M. Desai, Fault detection using a linear algebraic approach, in *Linear Algebra in Signals, Systems, and Control*, eds. B. N. Datta, C. R. Johnson, M. A. Raashoek, R. Plemmons and E. Sontag (SIAM, 1988).

95. A. Ray, R. Geiger, M. Desai and J. Deyst, Analytic redundancy for on-line fault diagnosis in a nuclear reactor, *AIAA Journal of Energy* (1983) 367–373.

96. M. Razavi, T. J. Grabowski, W. P.Vispoel, P. Monahan, S. Mehta, B. Eaton and L. Bolinger, Model assessment and model building in fMRI, *Hum. Brain. Mapp.* **20**(4) (2003) 227–38.

97. D. L. Schacter and A. D. Wagner, Medial temporal lobe activations in fMRI and PET studies of episodic encoding and retrieval, *Hippocampus* **9**(1) (1999) 7–24.

98. L. L. Scharf and B. Friedlander, Matched subspace detectors, *IEEE Trans. Signal Processing* **42**(2) (1994) 2146–2157.

99. L. L. Scharf, *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis* (Addison-Wesley, 1990).

100. J. Shah, A Common Framework for Curve Evolution, Segmentation and Anisotropic Diffusion, *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, San Francisco (June, 1996), pp. 136–142.

101. N. A. Shapira, Y. Liu, A. G. He, M. M. Bradley, M. C. Lessig, G. A. James, D. J. Stein, P. J. Lang and W. K. Goodman, Brain activation by disgust-inducing pictures in obsessive-compulsive disorder, *Biol. Psychiatry* **54**(7) (2003) 751–6.

102. R. T. Schultz, I. Gauthier, A. Klin, R. K. Fulbright, A. W. Anderson, F. Volkmar, P. Skudlarski, C. Lacadie, D. J. Cohen and J. C. Gore, Abnormal ventral temporal cortical activity during face discrimination among individuals with autism and Asperger syndrome, *Arch. Gen. Psychiatry* **57**(4) (2000) 331–40.

103. V. Solo, E. Brown and R. Weisskoff, A signal processing approach to functional MRI for brain mapping, *Proc. IEEE. Int. Conf. Image Proc.* **2** (1997) 121–123.

104. V. Solo, P. Purdon, R. Weisskoff and E. Brown, A signal estimation approach to functional MRI, *IEEE Trans. Med. Imag.* **20**(1) (2001) 26–35.

105. A. G. Sorensen and P. Reimer, *Cerebral MR Perfusion Imaging: Principles and Current Applications* (Thieme, New York, 2000).

106. R. A. Sperling, J. F. Bates, E. F. Chua, A. J. Cocchiarella, D. M. Rentz, B. R. Rosen, D. L. Schacter and M. S. Albert, fMRI studies of associative encoding in young and elderly controls and mild Alzheimer's disease, *J. Neurol. Neurosurg. Psychiatry* **74**(1) (2003) 44–50.

107. R. Sperling, D. Greve, A. Dale, R. Killiany, J. Holmes, H. D. Rosas, A. Cocchiarella, P. Firth, B. Rosen, S. Lake, N. Lange, C. Routledge and M. Albert, Functional MRI detection of pharmacologically induced memory impairment, *Proc. Natl. Acad. Sci. USA* **99**(1) (2002) 455–60.

108. T. Taoka, S. Wasaki, H. Uchida, A. Fukusumi, H. Nakagawa, K. Kichikawa, K. Takayama, T. Yoshioka, M. Takewa M and H. Ohishi, Age correlation of the time lag in signal change on EPI-fMRI, *J. Comput. Assist. Tomogr.* **22**(4) (1998) 514–7.

109. G. Thierry, K. Boulanouar, F. Kherif, J.-P. Ranfeva and J.-F. Démonet, Temporal sorting of nerural components underlying phonological processing, *Neuroreport* **10** (1999) 2599–2603.

110. I. Tracey, L. Becerra, I. Chang, H. Breiter, L. Jenkins, D. Borsook and R. G. Gonzalez, Noxious hot and cold stimulation produce common patterns of brain activation in humans: a functional magnetic resonance imaging study, *Neurosci. Lett.* **288**(2) (2000) 159–62.

111. F. R. Volkmar and D. Pauls, *Autism*, Lancet **362**(9390) (2003) 1133–41.

112. B. A. Wandell, Computational neuroimaging of human visual cortex, *Annu. Rev. Neurosci.* **22** (1999) 145–73.

113. J. J. Warner, *Atlas of Neuroanatomy* (Butterworth Heinemann, Boston, MA, 2001).

114. K. J. Worsley, C. H. Liao, J. Aston, V. Petre, G. H. Duncan, F. Morales and A. C. Evans, A general statistical analysis for fMRI data, *NeuroImage* **15**(1) (2002) 1–15.

115. K. J. Worsley, A. C. Evans, P. C. Marrett and P. Neelin, Detecting and estimating the regions of activation in CBF activation studies in human brain, in *Qualification of Brain Function: Tracer Kinetics and Image Analysis in Brain PET*, eds. K. Umeura, N. Lassen, T. Jones and I Kanno (Excerpta Medica, London, 1993), pp. 535–48.

116. C. I. Wright, B. Martis, K. McMullin, L. M. Shin and S. L. Rauch, Amygdala and insular responses to emotionally valenced human faces in small animal specific phobia, *Biol. Psychiatry* **54**(10) (2003) 1067–76.

This page intentionally left blank

CHAPTER 9

TECHNIQUES FOR DETECTION OF SPECTRAL SIGNATURES IN MR IMAGES AND THEIR APPLICATIONS

CLAYTON CHI-CHANG CHEN[1,2,3], CHUIN-MU WANG[4], CHEIN-I CHANG[5],
CHING-WEN YANG[6], JYH WEN CHAI[2], YI-NUNG CHUNG[7] and PAU-CHOO CHUNG[8]

[1]*Department of Radiological Technology,*
Chungtai Institue of Health Science and Technology,
Taichung, Taiwan, 406, ROC

[2]*Department of Radiology, Taichung Veterans General Hospital,*
Taichung, Taiwan, 407, ROC

[3]*Institute of Physical Therapy, National Yangming University,*
Taipei, Taiwan, 112, ROC

[4]*Department of Electronic Engineering, National Chinyi Institute of Technology,*
Taichung, Taiwan, 411, ROC

[5]*Remote Sensing Signal and Image Processing Laboratory*
Department of Computer Science and Electrical Engineering
University of Maryland, Baltimore County, Baltimore, MD 21250

[6]*Computer Center Hospital, VACRS, Taichung,*
Taiwan, 407, ROC

[7]*Department of Electrical Engineering, Da-Yeh University,*
Changhua, Taiwan, ROC

[8]*Department of Electrical Engineering, National Cheng Kung University,*
Tainan, Taiwan, ROC

This chapter presents two major hyperspectral imaging techniques, orthogonal subspace projection (OSP) and constrained energy minimization (CEM), which have recently found success in spectral signature detection in magnetic resonance (MR) images. Unlike classical image classification techniques which are primarily designed on a pure pixel basis, the proposed OSP and CEM are mixed pixel classification techniques which estimate abundance fractions of different material substances assumed to be present in the image data so as to achieve classification. Three scenarios are considered for MR image classification based on different levels of object information provided *a priori*. The OSP is applied to the case that the complete prior object knowledge is known compared to the CEM which is used for the case that only desired object knowledge is given while other knowledge can be discarded. When no prior object knowledge is available, the OSP is extended to an unsupervised OSP (UOSP), which obtains the necessary object knowledge directly from the image data in an unsupervised manner for OSP classification. In order to demonstrate the utility of these three scenarios in MR image classification, a series of experiments are conducted and compared to the commonly used c-means method for performance evaluation. The results show that both OSP and CEM are promising and effective spectral techniques for MR image classification.

Keywords: Magnetic resonance image; orthogonal subspace projection; constrained energy minimization.

1. Introduction

Nuclear magnetic resonance (NMR) can be used to measure the nuclear spin density, the interactions of the nuclei with their surrounding molecular environment and those between close nuclei, respectively. It produces a sequence of multiple spectral images of tissues with a variety of contrasts using three magnetic resonance parameters, spin-lattice (T1), spin-spin (T2) and dual echo-echo proton density (PD). By appropriately choosing pulse sequence parameters such as echo time (TE) and repetition time (TR), a sequence of images of specific anatomic areas can be generated by pixel intensities that represent characteristics of different types of tissues throughout the sequence. Additionally, the spectral correlation among the image sequence also produces information that spatial correlation cannot provide. As a result, magnetic resonance imaging (MRI) becomes a more useful image modality than X-ray computerized tomography (X-CT) when it comes to analysis of soft tissues and organs since the information about T1 and T2 offers a more precise picture of tissue functionality than that produced by X-CT.[1,2]

Over the past years many computer-assisted diagnostic (CAD) methods have been reported in the literature[3–20] such as PCA,[6] eigenimage analysis,[7–12] neural networks,[13–16] fuzzy c-means methods,[17–18] hybrid methods,[19] knowledge-based techniques,[20] etc. For example, eigenimage filtering-based approach has shown a promise in segmentation and feature extraction. Hybrid methods combine imaging processing and model-based techniques to improve segmentation. Knowledge-based techniques further allow one to make more intelligent classification and segmentation decisions. As an alternative, neural networks are also proposed to demonstrate their superior performance in segmentation of brain tissues to classical maximum likelihood methods.

Most recently, remote sensing image processing techniques have been also explored where two hyperspectral image processing techniques, orthogonal subspace projection (OSP) and constrained energy minimization (CEM) have been shown promise in applications of MR image classification.[21–24] Unlike traditional image classification methods which are pure pixel-based techniques, the OSP[25–27] and CEM[25,27–30] are mixed pixel classification techniques which consider an image pixel as a pixel linearly mixed by several substances resident in the pixel. The classification is then carried out by the estimated abundance fractions of these substances present in the pixel. These remote sensing image processing-based techniques have changed the way that classical image processing techniques classify MR images where an image pixel is assumed to be a pure pixel and the classification is primarily performed on the basis of spatial correlation. The OSP was developed based on the fact that the prior knowledge about the spectral signatures of different tissues such as white matter (WM), gray matter (GM), cerebral spinal fluid (CSF) is known so that an image pixel in an MR image can be modeled as a linear mixture of these signatures from which each of the signatures can be classified in accordance with their estimated abundance fractions. One of major drawbacks from

which the OSP suffers is the requirement of the complete prior knowledge about the spectral signatures of interest in MR images. In order to resolve this dilemma, two approaches have been proposed to mitigate this issue. One is to make the OSP unsupervised so that the required signature knowledge can be obtained directly from the data. The resulting OSP is called an unsupervised OSP (UOSP).[22] In this case, there is no need of knowing signatures *a priori* as required by the OSP. The other is to only focus on the knowledge of the spectral signature of interest, while discarding all other information such as background information. One such an approach, referred to as the constrained energy minimization (CEM) was investigated.[23] In this chapter, we revisit these techniques for MR image classification. In particular, we will investigate their strengths and weaknesses in the sense of different levels of information provided *a priori*.

The remainder of this chapter is organized as follows. Section 2 describes the OSP approach. Section 3 presents the CEM approach. Section 4 extends the OSP to an (UOSP). Section 5 briefly describes the c-means method and a modified version of the c-means method for comparison. Sections 6 and 7 conduct a series of computer simulations and real MR image experiments respectively to evaluate performance of OSP, CEM and UOSP in classification performance and also compare the results to those produced by the CM method. Section 8 concludes some comments.

2. Orthogonal Subspace Projection (OSP)

The OSP technique was originally developed for hyperspectral image classification.[26] It models an image pixel as a linear mixture of finite number of known signatures assumed to be present in the image. More specifically, suppose that L is the number of spectral bands (channels). In our case, an MR image sequence is actually a collection of co-registered L spectral bands. So, an i-th image pixel in an MR image sequence can be considered as an L-dimensional pixel vector. Let $\mathbf{r}$ be an $L \times 1$ column pixel vector in a hyperspectral image where the bold face is used for vectors. Assume that there is a set of p targets of interest present in the image and $\mathbf{m}_1, \mathbf{m}_2, \ldots, \mathbf{m}_p$ are their corresponding signatures. Let $\mathbf{M}$ be an $L \times p$ target signature matrix denoted by $\mathbf{M} = [\mathbf{m}_1 \mathbf{m}_2 \cdots \mathbf{m}_p]$ where $\mathbf{m}_j$ is an $L \times 1$ column vector represented by the signature of the j-th target resident in the image scene. Let $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \ldots, \alpha_p)^T$ be a $p \times 1$ abundance column vector associated with $\mathbf{r}$, where α_j denotes the fraction of the j-th signature $\mathbf{m}_j$ present in the pixel vector $\mathbf{r}$. Then the spectral signature of $\mathbf{r}$ can be represented by the following linear mixture model

$$\mathbf{r} = \mathbf{M}\boldsymbol{\alpha} + \mathbf{n} \tag{1}$$

where $\mathbf{n}$ is noise or can be interpreted as either a measurement error or a model error.

Equation (1) assumes that the knowledge of $\mathbf{M}$ must be given *a priori*. Without loss of generality we further assume that $\mathbf{d} = \mathbf{m}_p$ is the desired target signature to

be detected or classified and $\mathbf{U} = [\mathbf{m}_1 \mathbf{m}_2 \cdots \mathbf{m}_{p-1}]$ is the undesired target signature matrix made up of the remaining $p-1$ undesired target signatures in $\mathbf{M}$. Then, we rewrite (1) as

$$\mathbf{r} = \mathbf{d}\alpha_p + \mathbf{U}\gamma + \mathbf{n} \tag{2}$$

where γ is the abundance vector associated with $\mathbf{U}$. Model specified by (2) will be referred to as OSP-model since it is the first model proposed to separate the desired target signature $\mathbf{d}$ from the undesired target signatures in $\mathbf{U}$. This allows us to design the following orthogonal subspace projector, denoted by $P_{\mathbf{U}}^{\perp}$ to annihilate $\mathbf{U}$ from $\mathbf{r}$ prior to classification

$$P_{\mathbf{U}}^{\perp} = \mathbf{I} - \mathbf{U}\mathbf{U}^{\#} \tag{3}$$

where $\mathbf{U}^{\#} = \left(\mathbf{U}^T \mathbf{U}\right)^{-1} \mathbf{U}^T$ is the pseudo-inverse of $\mathbf{U}$. Applying $P_{\mathbf{U}}^{\perp}$ in (3) to (2) results in

$$P_{\mathbf{U}}^{\perp}\mathbf{r} = P_{\mathbf{U}}^{\perp}\mathbf{d}\alpha_p + P_{\mathbf{U}}^{\perp}\mathbf{n}. \tag{4}$$

Equation (4) represents a standard signal detection problem. If the signal-to-noise ratio (SNR) is chosen as the criterion for optimality, the optimal solution to (4) is given by a matched filter $\mathbf{M_d}$ defined by

$$\mathbf{M_d}\left(P_U^{\perp}\mathbf{r}\right) = \kappa \mathbf{d}^T P_{\mathbf{U}}^{\perp}\mathbf{r} \tag{5}$$

where the matched signal is specified by $\mathbf{d}$ and κ is a constant. Setting $\kappa = 1$ in (5) yields the following OSP filter $\delta^{\mathrm{OSP}}(\mathbf{r})$[26]

$$\delta^{\mathrm{OSP}}(\mathbf{r}) = \mathbf{d}^T P_{\mathbf{U}}^{\perp}\mathbf{r} = (\mathbf{d}^T P_{\mathbf{U}}^{\perp}\mathbf{d})\alpha_p + \mathbf{d}^T P_{\mathbf{U}}^{\perp}\mathbf{n}. \tag{6}$$

The OSP specified by (6) only detects the abundance fraction which does not reflect true abundance fraction of the desired target signature $\mathbf{d}$. In order to reliably estimate the abundance fraction of $\mathbf{d}$ a least squares OSP, $\delta^{\mathrm{LSOSP}}(\mathbf{r})$, given by[31−33]

$$\delta^{\mathrm{LSOSP}}(\mathbf{r}) = \left(\mathbf{d}^T P_{\mathbf{U}}^{\perp}\mathbf{d}\right)^{-1} \mathbf{d}^T P_{\mathbf{U}}^{\perp}\mathbf{r} = \left(\mathbf{d}^T P_{\mathbf{U}}^{\perp}\mathbf{d}\right)^{-1} \delta^{\mathrm{OSP}}(\mathbf{r}) \tag{7}$$

where the constant $\left(\mathbf{d}^T P_{\mathbf{U}}^{\perp}\mathbf{d}\right)^{-1}$ is included in (7) to account for the least-squares estimation error.[31−33]

3. Constrained Energy Mimization (CEM)

Assume that $\mathbf{r}_i = (r_{i1}, r_{i2}, \ldots, r_{iL})^T$ is the i-th image pixel vector with r_{ij} representing the pixel of the i-th pixel vector in the j-th spectral band. Suppose that $\{\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_N\}$ is a set of all image pixel vectors in an MR image sequence where N is the total number of pixels in the image. Let $\mathbf{d}$ be the spectral signature of an object of interest. The goal is to design a finite impulse response (FIR) linear filter specified by an L-dimensional vector $\mathbf{w} = (w_1, w_2, \ldots, w_L)^T$ that passes the desired signature $\mathbf{d}$ by constraining its direction while minimizing its output energy that are caused by signal source vectors with directions other than the constrained direction.

More specifically, let y_i denote the output of the designed FIR filter resulting from the i-th MR image pixel $\mathbf{r}_i$. Then y_i can be expressed by

$$y_i = \sum_{l=1}^{L} w_l r_{il} = \mathbf{w}^T \mathbf{r}_i = \mathbf{r}_i^T \mathbf{w}. \tag{8}$$

The average filter output energy resulting from $\{\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_N\}$ is given by

$$\frac{1}{N} \left[\sum_{i=1}^{N} y_i^2 \right] = \frac{1}{N} \left[\sum_{i=1}^{N} \left(\mathbf{r}_i^T \mathbf{w} \right)^T \left(\mathbf{r}_i^T \mathbf{w} \right) \right]$$
$$= \mathbf{w}^T \left[\frac{1}{N} \left(\sum_{i=1}^{N} \mathbf{r}_i \mathbf{r}_i^T \right) \right] \mathbf{w} = \mathbf{w}^T \mathbf{R}_{L \times L} \mathbf{w} \tag{9}$$

where $\mathbf{R}_{L \times L} = \frac{1}{N} \left[\sum_{i=1}^{N} \mathbf{r}_i \mathbf{r}_i^T \right]$ is the auto-correlation sample matrix of the MR image sequence. So, the CEM filter is one solving the following linearly constrained optimization problem

$$\min_{\mathbf{w}} \{ \mathbf{w}^T \mathbf{R}_{L \times L} \mathbf{w} \} \quad \text{subject to } \mathbf{d}^T \mathbf{w} = 1. \tag{10}$$

The solution to (10) is given by[25,27–30,32]

$$\mathbf{w}^{\mathrm{CEM}} = \frac{\mathbf{R}_{L \times L}^{-1} \mathbf{d}}{\mathbf{d}^T \mathbf{R}_{L \times L}^{-1} \mathbf{d}}. \tag{11}$$

Substituting the optimal weight given by (11) for $\mathbf{w}$ in (1) yields the CEM filter which implements a detector, $\delta^{\mathrm{CEM}}(\mathbf{r})$ on an image pixel vector $\mathbf{r}$ and is given by

$$\delta^{\mathrm{CEM}}(\mathbf{r}) = \left(\mathbf{w}^{\mathrm{CEM}} \right)^T \mathbf{r} = \left(\frac{\mathbf{R}_{L \times L}^{-1} \mathbf{d}}{\mathbf{d}^T \mathbf{R}_{L \times L}^{-1} \mathbf{d}} \right)^{-1} \mathbf{r} = \frac{\mathbf{d}^T \mathbf{R}_{L \times L}^{-1} \mathbf{r}}{\mathbf{d}^T \mathbf{R}_{L \times L}^{-1} \mathbf{d}}. \tag{12}$$

As we can see from (12), when $\mathbf{r} = \mathbf{d}$, $\delta^{\mathrm{CEM}}(\mathbf{d}) = 1$ which satisfies the constraint in (10). In this case, the $\mathbf{r}$ is considered to be the desired object pixel and will be extracted by the CEM filter. Despite that the primary task of the CEM filter is object detection, as demonstrated in the experiments it can perform as a classifier by detecting different types of objects, one at a time. In this case, separate images are produced for each type of targets.

A comment is noteworthy. The value of $\delta^{\mathrm{CEM}}(\mathbf{r})$ resulting from (12) represents the estimated abundance fraction of the object signature $\mathbf{d}$ contained in the image pixel $\mathbf{r}$. So, unlike most spatial-based classification methods that can be considered as label (class)-assignment techniques, the CEM filter detects a desired object by estimating its abundance fraction using (10). As a result, the image generated by the CEM filter is generally gray scale where the gray level value of each image pixel reflects the detected amount of the abundance fraction of the desired object present in the pixel. The object detection is then performed based on the resulting gray scale image and classification is carried out by detecting the desired objects in separate images.

4. Unsupervised Orthogonal Subspace Projection (UOSP)

The automatic target generation process (ATGP) was previously developed to find potential target pixels that can be used to generate a signature matrix used in an orthogonal subspace projection (OSP) approach. It is one of two processes used in the automatic target detection and classification algorithm (ATDCA).[33−34]

The ATGP repeatedly makes use of (3) to find target pixel vectors of interest from the data without prior knowledge regardless of what types pixels of these targets. It can be briefly described as follows. Assume that $\mathbf{t}_0$ is an initial target pixel vector. The ATGP begins with the initial target pixel vector $\mathbf{t}_0$ by applying an orthogonal subspace projector $P_{\mathbf{t}_0}^{\perp}$ specified by (3) with $\mathbf{U} = \mathbf{t}_0$ to all image pixel vectors. It then finds a target pixel vector, denoted by $\mathbf{t}_1$ with the maximum orthogonal projection in the orthogonal complement space, denoted by $\langle \mathbf{t}_0 \rangle^{\perp}$ that is orthogonal to the space, $\langle \mathbf{t}_0 \rangle$ linearly spanned by $\mathbf{t}_0$. The reason for this selection is that the selected $\mathbf{t}_1$ generally has the most distinct features from $\mathbf{t}_0$ in the sense of orthogonal projection because $\mathbf{t}_1$ has the largest magnitude of the projection in $\langle \mathbf{t}_0 \rangle^{\perp}$ produced by $P_{\mathbf{t}_0}^{\perp}$. A second target pixel vector $\mathbf{t}_2$ can be found by applying an orthogonal subspace projector $P_{[\mathbf{t}_0\,\mathbf{t}_1]}^{\perp}$ with $\mathbf{U} = [\mathbf{t}_0\ \mathbf{t}_1]$ to the original image and a target pixel vector that has the maximum orthogonal projection in $\langle \mathbf{t}_0, \mathbf{t}_1 \rangle^{\perp}$ is selected as $\mathbf{t}_2$. The above procedure is repeated over and over again to find a third target pixel vector $\mathbf{t}_3$, a fourth target pixel vector $\mathbf{t}_4$, etc. until a certain stopping rule is satisfied. In this paper, the stopping rule is determined by the number of target pixel vectors required to generate, p which is predetermined *a priori*. Using the p as a stopping criterion, the ATGP can be implemented in the following steps.

Automatic Target Generation Process (ATGP)

(1) Initial condition:
    Select an initial target pixel vector of interest denoted by $\mathbf{t}_0$. In order to initialize the ATGP without knowing $\mathbf{t}_0$, we select a target pixel vector with the maximum length as the initial target $\mathbf{t}_0$, namely, $\mathbf{t}_0 = \arg\{\max_{\mathbf{r}} \mathbf{r}^T \mathbf{r}\}$, which has the highest intensity, i.e. the brightest pixel vector in the image scene. Set $k = 1$ and $\mathbf{U}_0 = [\mathbf{t}_0]$. (It is worth noting that this selection may not be necessarily the best selection. However, according to our experiments it was found that the brightest pixel vector was always extracted later on, if it was not used as an initial target pixel vector in the initialization.)
(2) At k-th iteration, apply $P_{\mathbf{t}_0}^{\perp}$ via (3) to all image pixels $\mathbf{r}$ in the image and find the k-th target $\mathbf{t}_k$ generated at the k-th stage which has the maximum orthogonal projection as follows.

$$\mathbf{t}_k = \arg\left\{ \max_{\mathbf{r}} \left[\left(P_{[\mathbf{t}_0\,\mathbf{U}_{k-1}]}^{\perp} \mathbf{r} \right)^T \left(P_{[\mathbf{t}_0\,\mathbf{U}_{k-1}]}^{\perp} \mathbf{r} \right) \right] \right\} \tag{13}$$

where $\mathbf{U}_{k-1} = [\mathbf{t}_1 \mathbf{t}_2 \cdots \mathbf{t}_{k-1}]$ is the target matrix generated at the $(k-1)^{\text{st}}$ stage.

(3) Stopping rule:
   If $k < p - 1$, let $\mathbf{U}_k = [\mathbf{U}_{k-1}\mathbf{t}_k] = [\mathbf{t}_1\mathbf{t}_2\cdots\mathbf{t}_k]$ be the k-th target matrix, go to Step 2. Otherwise, continue.
(4) At this stage, the ATGP is terminated. The resulting target matrix is $\mathbf{U}_{p-1}$, contains $p - 1$ target pixel vectors as its column vectors, which do not include the initial target pixel vector $\mathbf{t}_0$.

As a result of the ATGP, the final set of target pixel vectors produced by the ATGP at step 4 is the final target set which comprises p target vectors, $\{\mathbf{t}_0, \mathbf{t}_1, \mathbf{t}_2, \ldots, \mathbf{t}_{p-1}\} = \{\mathbf{t}_0\} \cup \{\mathbf{t}_1, \mathbf{t}_2, \ldots, \mathbf{t}_{p-1}\}$ which were found by repeatedly using (13). Now, by virtue of the ATGP, we can extend the OSP for (6) to an unsupervised OSP (UOSP) which produces its own signature matrix $\mathbf{M} = [\mathbf{t}_0\mathbf{t}_1\cdots\mathbf{t}_{p-1}]$ for (1) that is obtained directly from the image and is made up of the p target pixel vectors $\mathbf{t}_0, \mathbf{t}_1, \mathbf{t}_2, \ldots, \mathbf{t}_{p-1}$ generated by the ATGP. The implementation of the UOSDP can be briefly summarized as follows.

UOSP Algorithm

(1) Preset the p, number of signatures of interest.
(2) Apply the ATGP to generate p target signatures, $\mathbf{t}_0, \mathbf{t}_1, \mathbf{t}_2, \ldots, \mathbf{t}_{p-1}$.
(3) Form a desirable signature matrix $\mathbf{M} = [\mathbf{t}_0\mathbf{t}_1\cdots\mathbf{t}_{p-1}]$ for (1).
(4) Due to the unavailability of prior knowledge of signatures, the LSOSP classifier, $\delta^{\text{LSOSP}}(\mathbf{r})$ described by (7) must be applied to classify each of p signatures $\mathbf{t}_0, \mathbf{t}_1, \mathbf{t}_2, \ldots, \mathbf{t}_{p-1}$. In other words, the LSOSP classifier must be performed p times to classify all the p signatures. In this case, in order to classify the j-th signature $\mathbf{t}_j$, the $\mathbf{d}$ and $\mathbf{U}$ in $\delta^{\text{LSOSP}}(\mathbf{r})$ are specified by $\mathbf{t}_j$ and $\mathbf{U}_j = [\mathbf{t}_0\cdots\mathbf{t}_{j-1}\mathbf{t}_{j+1}\cdots\mathbf{t}_{p-1}]$.

5. C-Means Methods

In order to evaluate performance of the CEM approach, the widely used c-means (CM) method[35] (also known as k-means[36]) is used for comparative analysis. The reason to select the CM method is two-fold. One is that it allows us to generate background signatures in an unsupervised manner for classification. Another is that it is basically a spatial-based pattern classification technique. As opposed to the CEM approach that only classifies objects of interest, the CM method classifies all MR image pixel vectors including background pixel vectors into pattern classes.

The CM method to be implemented here for experiments has two versions. One is completely unsupervised, also referred to as ISODATA,[35−36] which will be used to compare with the unsupervised OSP since both require no prior knowledge about the MR images to be processed. The other is a modified version of CM, referred to as MCM, which will be used to compare the CEM with the desired object knowledge provided *a priori*. In the latter case, the MCM method includes into its clustering procedure the same knowledge of objects of interest that is required by the CEM

approach. The detailed implementation of the CM method can be described as follows.

CM (ISODATA) Method

1. Initialization:
   Determine the number of pattern classes, c and randomly select p class means $\boldsymbol{\mu}_i^{(k)}$ for $1 \leq i \leq c$ and let $k = 0$.
2. At the $k \geq 0$ iteration, compute the distance of each sample pixel vector from all class means, $\boldsymbol{\mu}_i^{(k)}$ for $1 \leq i \leq c$ and assign the sample vector to the class whose mean has the shortest distance to the sample vector.
3. Compute the means of the reclustered sample vectors for each class, $\hat{\boldsymbol{\mu}}_i^{(k)}$ for $1 \leq i \leq c$.
4. If there is any mean changed, that is $\hat{\boldsymbol{\mu}}_i^{(k)} \neq \boldsymbol{\mu}_i^{(k)}$ for some $1 \leq i \leq c$, let $\boldsymbol{\mu}_i^{(k)} \leftarrow \hat{\boldsymbol{\mu}}_i^{(k)}$ and $k \leftarrow k + 1$. Go to step 2. Otherwise, the obtained $\{\boldsymbol{\mu}_i^{(k)}\}_{i=1}^c$ are the desired class means and the algorithm is terminated.

In order to include the desired object knowledge into the clustering process described in the above CM method, let the spectral signatures of p objects of interest be denoted by $\{\mathbf{d}_i\}_{i=1}^p$ where $\mathbf{d}_i$ is the spectral signature of the i-th object. A modified CM (MCM) method can be further developed as follows.

MCM Method

1. Initialization:
   Determine the number of pattern classes, $c \geq p$ and let $\{\boldsymbol{\mu}_i^{(k)}\}_{i=1}^c$ be their corresponding class means. Let the first p class means are fixed at $\boldsymbol{\mu}_i^{(k)} = \mathbf{d}_i$ for $1 \leq i \leq p$ and all other class means $\{\boldsymbol{\mu}_i^{(k)}\}_{i=p+1}^c$ are selected randomly, but must be different from the $\{\mathbf{d}_i\}_{i=1}^p$ which already selected for $\{\boldsymbol{\mu}_i^{(k)}\}_{i=1}^p$. Set $k = 0$.
2. At the $k \geq 0$ iteration, compute the distance of each sample pixel vector from all class means, $\{\boldsymbol{\mu}_i^{(k)}\}_{i=1}^c$ and assign the sample vector to the class whose mean has the shortest distance to the sample vector.
3. For each class i with $p + 1 \leq i \leq c$, recompute its class mean by averaging the sample vectors in the class, denoted by $\hat{\boldsymbol{\mu}}_i^{(k)}$.
4. If there is any mean changed, that is $\hat{\boldsymbol{\mu}}_i^{(k)} \neq \boldsymbol{\mu}_i^{(k)}$ for some $p + 1 \leq i \leq c$, let $\boldsymbol{\mu}_i^{(k)} \leftarrow \mathbf{d}_i$ for $1 \leq i \leq p$ and $\boldsymbol{\mu}_i^{(k)} \leftarrow \hat{\boldsymbol{\mu}}_i^{(k)}$ for $p + 1 \leq i \leq c$. Let $k \leftarrow k + 1$ and go to step 2. Otherwise, the obtained $\{\boldsymbol{\mu}_i^{(k)}\}_{i=1}^c$ are the desired class means and the algorithm is terminated.

It should be noted that the knowledge of $\{\mathbf{d}_i\}_{i=1}^p$ is given *a priori*. Therefore, the first p class means are fixed during iterations. However, the class means, $\{\boldsymbol{\mu}_i^{(k)}\}_{i=p+1}^c$ are regenerated at each iteration by the MCM method in an unsupervised manner using the minimum distance as a criterion. These generated class

means are considered to be signatures of unknown signal sources, which are not provided by prior knowledge and may include background signatures.

Since the CM and MCM methods are unsupervised pattern classification techniques, one of its weaknesses is the determination of c, i.e. the number of pattern classes. If c is too small, the number of c pattern classes may not well represent the data, in which several distinct classes may be merged into one class. If c is too large, the number of c pattern classes may over-represent the data, in which a class may be forced to be broken up into several classes. The CEM resolves this dilemma by performing object classification without using any information other than that provided by $\{\mathbf{d}_i\}_{i=1}^{p}$. As a result, has an advantage over the CM, MCM and UOSP methods in the sense that it does not require the knowledge of c.

6. Computer Simulations for Phantom Experiments

In this section, we present a series of experiments based on computer-generated phantom images. These phantom image experiments enable us to conduct a quantitative study and error analysis for various methods presented in this chapter. A set of computer simulations was performed to conduct a quantitative study and performance analysis of the OSP, CEM and UOSP in comparison with the CM and MCM methods described in Section IV with number of classes $p = c = 4$ representing four classes of WM, GM, CSF and image background of interest. The computer-generated phantom images used for simulations are shown in Fig. 1 which have five bands, each of which was made up of six overlapped ellipses with their spectral signatures shown in Fig. 2.

(a) band 1　　　　　　(b) band 2　　　　　　(c) band 3

(d) band 4　　　　　　(e) band 5
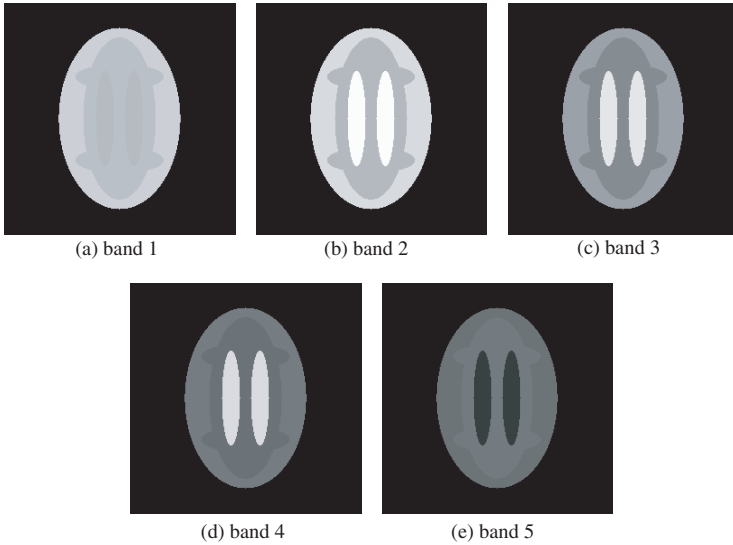
Fig. 1.　Five bands of phantom images used for simulations.
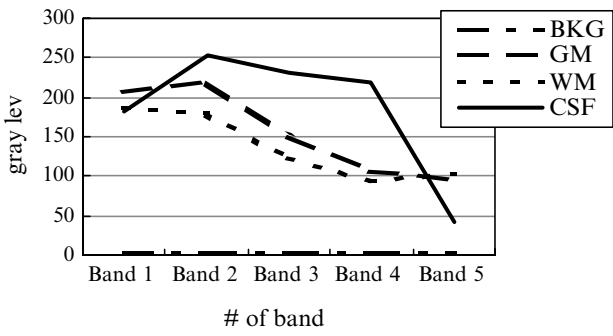
Fig. 2.   GM, WM, CSF and BKG spectral signatures.

These ellipses represent structure areas of three interesting cerebral tissues corresponding to gray matter (GM), white matter (WM) and cerebral spinal fluid (CSF). From the periphery to the center are background (BKG), GM, WM and CSF simulated by the signatures given in Fig. 2. The gray level values of these areas in each band were simulated in such a fashion that these values reflect the average values of their respective tissues in real MR images shown in Fig. 3.

Table 1 tabulates the values of the parameters used by the MRI pulse sequence and the gray level values of the tissues of each band used in the experiments.
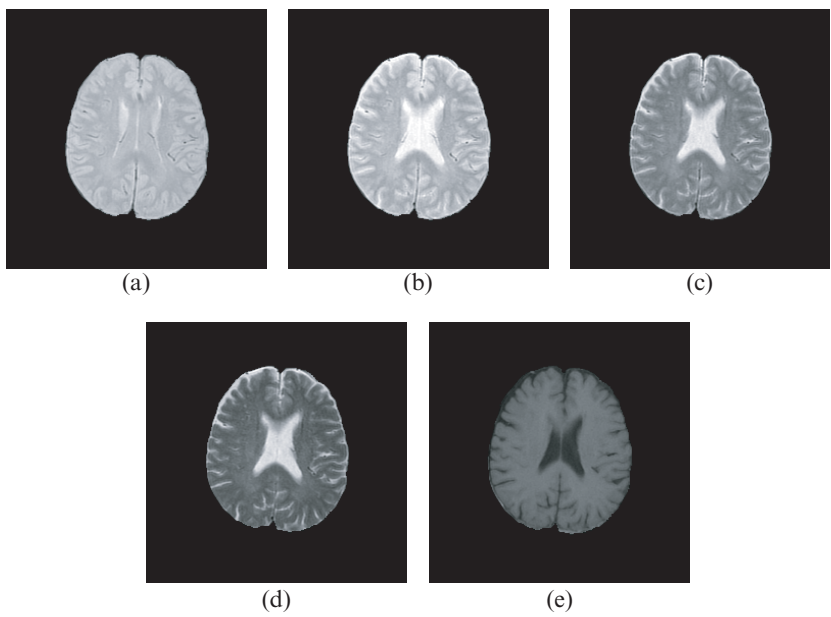
Fig. 3.   Five band real brain MR images used for experiments. (a) TR/TE = 2500 ms/25 ms; (b) TR/TE = 2500 ms/50 ms; (c) TR/TE = 2500 ms/75 ms; (d) TR/TE = 2500 ms/100 ms; (e) TR/TE = 500 ms/11.9 ms.

Table 1. Gray level values used for the five bands of the test phantom in Fig. 1.

Band #	MRI Parameter	GM	WM	CSF
Band 1	TR/TE = 2500 ms/25 ms	207	188	182
Band 2	TR/TE = 2500 ms/50 ms	219	180	253
Band 3	TR/TE = 2500 ms/75 ms	150	124	232
Band 4	TR/TE = 2500 ms/100 ms	105	94	220
Band 5	TR/TE = 500 ms/11.9 ms	95	103	42

A zero-mean Gaussian noise was added to the phantom images in Fig. 1 so as to achieve different levels of signal-to-noise ratios (SNR) ranging from 5 db to 20 db. Despite the fact that such MR phantom images may be unrealistic, they only serve as a purpose for illustration of the proposed CEM, OSP and UOSP techniques and demonstration of its advantages.

One of unique features of the remote sensing image processing techniques presented in this chapter is that the produced images are fractional abundance images, each of which shows the estimated abundance fractions of a particular spectral signature assumed to be present in the MR image. As a result, these fractional abundance images are generally gray scale as opposed to the class-map images produced by classical image classification techniques. In other words, the former is classified by soft decisions compared to the latter classified by hard decisions. In order to convert a soft decision to a hard decision, a thresholding technique is needed. In what follows, we describe one thresholding method developed in Ref. 37.

6.1. *Abundance percentage thresholding method*

Since the fractional abundance images are real-values, we first normalize the abundance fractions of these images to the range of $[0, 1]$. More specifically, let $\mathbf{r}$ be the image pixel vector and $\hat{\alpha}_1(\mathbf{r}), \hat{\alpha}_2(\mathbf{r}), \ldots, \hat{\alpha}_p(\mathbf{r})$ are the estimated abundance fractions of $\alpha_1, \alpha_2, \ldots, \alpha_p$ present in the $\mathbf{r}$. Then for each estimated abundance fraction $\hat{\alpha}_j(\mathbf{r})$ its normalized abundance fraction, $\tilde{\alpha}_j(\mathbf{r})$ can be obtained by

$$\tilde{\alpha}_j(\mathbf{r}) = \frac{\hat{\alpha}_j(\mathbf{r}) - \min_{\mathbf{r}} \hat{\alpha}_j(\mathbf{r})}{\max_{\mathbf{r}} \hat{\alpha}_j(\mathbf{r}) - \min_{\mathbf{r}} \hat{\alpha}_j(\mathbf{r})}. \tag{14}$$

Suppose that $a\%$ is used for the cut-off abundance fraction threshold value. If the normalized abundance fraction of a pixel vector is greater than or equal to $a\%$, then the pixel is detected as the desired object pixel and will be assigned by a "1"; otherwise, the pixel is assigned by a "0", in which case the pixel does not match the desired object signature. Using (14) with $a\%$ as the cut-off threshold value to threshold a fractional abundance image is referred to as $a\%$ thresholding method.

6.2. *3-D ROC analysis*

By virtue of the above $a\%$ threshold method, we can actually tally the number of pixels detected in its generated fractional abundance images. This subsection

further develops a 3-D ROC analysis based on $a\%$. First of all, let $\{\mathbf{d}_i\}_{i=1}^{P}$ be a set of objects of interest, which we would like to classify. We define $N(\mathbf{d}_i)$, $N_D(\mathbf{d}_i)$ and $N_F(\mathbf{d}_i)$ to be the total number of pixels specified by the i-th object signature $\mathbf{d}_i$, the total number of pixels that are specified by the object signature $\mathbf{d}_i$ and actually detected as the $\mathbf{d}_i$, and the total number of false alarm pixels that are not specified by the object signature $\mathbf{d}_i$ but detected as the $\mathbf{d}_i$ respectively. For example, the desired object signature $\mathbf{d}_i$ can be chosen to be one of GM, WM or CSF. Using the definitions of $N(\mathbf{d}_i)$, $N_D(\mathbf{d}_i)$ and $R_F(\mathbf{d}_i)$ we further define the detection rate $R_D(\mathbf{d}_i)$, false alarm rate $R_F(\mathbf{d}_i)$ for $\mathbf{d}_i$ and mean detection rate R_D, mean false alarm rate, R_F by

$$R_D(\mathbf{d_i}) = \frac{N_D(\mathbf{d_i})}{N(\mathbf{d_i})} \tag{15}$$

$$R_F(\mathbf{d_i}) = \frac{N_F(\mathbf{d_i})}{N - N(\mathbf{d_i})} \tag{16}$$

and

$$R_D = \sum_{i=1}^{p} R_D(\mathbf{d}_i)p(\mathbf{d}_i) \tag{17}$$

$$R_F = \sum_{i=1}^{p} R_F(\mathbf{d}_i)p(\mathbf{d}_i) \tag{18}$$

where N is the total number of pixels in the image and $p(\mathbf{d}_i) = N(\mathbf{d}_i)/\sum_{i=1}^{p} N(\mathbf{d}_i)$. It is worth noting that the mean detection rate R_D defined by (17) is the average of detection rates over the detected objects. In order to classify p objects $\{\mathbf{d}_i\}_{i=1}^{p}$, we must calculate its mean detection rate. Similarly, the mean false alarm R_F defined by (18) is the average of false alarm rates over the detected objects. Using (15–18) each fixed $a\%$ produces a pair of R_D and R_F. As a consequence, varying $a\%$ from 0% up to 100% generates a set of pairs (R_D, R_F) where each pair results from a particular $a\%$ being used as a cutoff threshold value. In this case, we use an approach proposed in Ref. 33 and 38 to plot a 3-D ROC curve based on three parameters, $(R_D, R_F, a\%)$, where the (x, y)-coordinate corresponds $(R_F, a\%)$ and z-axis is specified by R_D. By means of such a 3-D ROC curve we can further plot three 2-D curves of (R_D, R_F), $(R_D, a\%)$ and $(R_F, a\%)$ where the 2-D curve of (R_D, R_F) can be viewed as the traditional ROC curve.[39]

6.3. *Phantom image experiments*

Now we can use this 3-D ROC curve along with three 2-D curves to analyze the performance of five different methods, OSP, CEM, UOSP, CM and MCM methods with different SNRs in classification of GM, WM and CSF. Since these five methods require different levels of object information, their performance will be different. In this case, a fair comparison among these five methods may be difficult.

The experiments presented in this section only demonstrate advantages and disadvantages of each of these five methods.

In our experiments, the number of classes was assumed to be 4, i.e. $c = 4$. The spectral signatures of GM, WM, CSF and background (BKG) used as prior knowledge for the OSP and CEM provided by Fig. 2. When the OSP was implemented, the signature matrix $\mathbf{M}$ consisted of four signatures, GM, WM, CSF and BKG. When the CEM was implemented, only the desired signature $\mathbf{d}$ was selected to be one of the three GM, WM and CSF signatures to be detected. Similarly, when the MCM method was implemented, the $p = 1$ and $\boldsymbol{\mu}_1^{(k)}$ was always fixed at the desired signature $\mathbf{d}$ during the course of iterations, i.e. $\boldsymbol{\mu}_1^{(k)} = \mathbf{d}$. Since the UOSP and CM method generate the signatures of GM, WM, CSF and BKG directly from the data, no prior knowledge about the GM, WM, CSF and BKG is required for both of them. Figures 4–8 show the results of the OSP, CEM, MCM method, UOSP and CM method in classification of GM, WM and CSF respectively.

As we can see from Figs. 4–8, it is difficult to conduct a comparative analysis since the fractional abundance images in Figs. 4–5 and 7 produced by the OSP, CEM and UOSP are gray scale compared to the images produced by the CM and MCM methods in Figs. 6 and 8 which are class maps, but not gray scale images. Therefore, we first analyze performance among the three hyperspectral image processing techniques, OSP, CEM and UOSP in terms of 3-D ROC curves in

(a) GM   (b) WM   (c) CSF
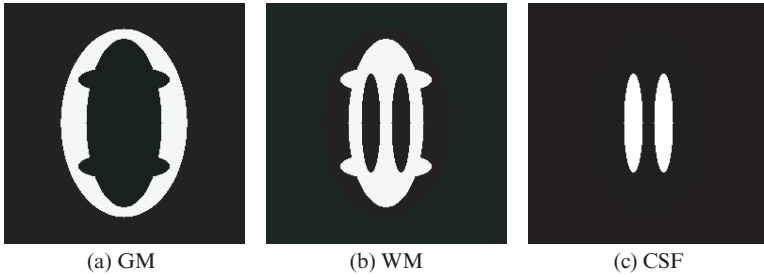
Fig. 4.   Classification results of GM, WM and CSF produced by the OSP.

(a) GM   (b) WM   (c) CSF
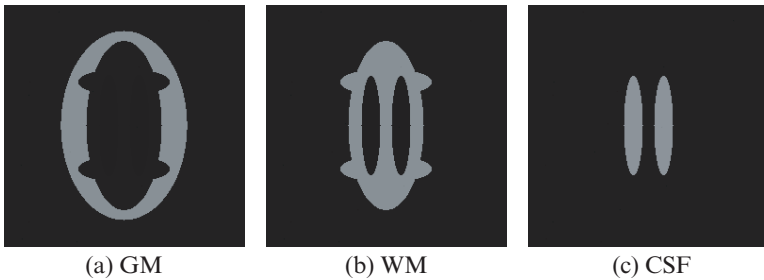
Fig. 5.   Classification results of GM, WM and CSF produced by the CEM.

(a) GM   (b) WM   (c) CSF
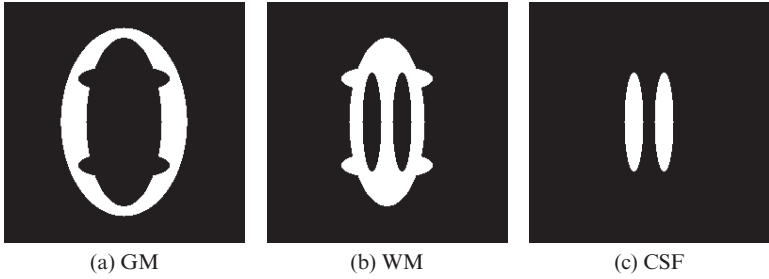
Fig. 6.   Classification results of GM, WM and CSF produced by the MCM method.

(a) GM   (b) WM   (c) CSF
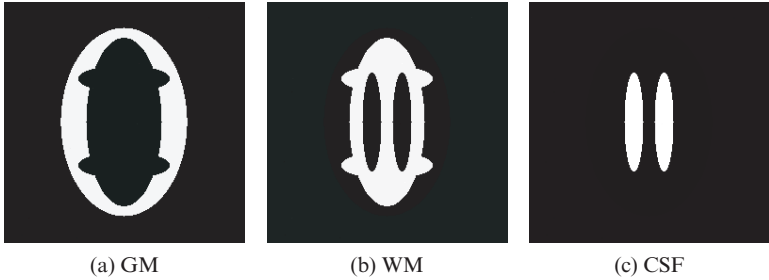
Fig. 7.   Classification results of GM, WM and CSF produced by the UOSP.

(a) GM   (b) WM   (c) CSF
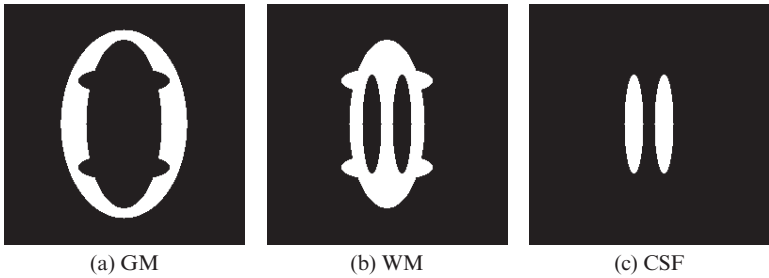
Fig. 8.   Classification results of GM, WM and CSF produced by the CM method.

Figs. 9–11, then calculate their mean classification rates in Tables 2–3 which will be further compared to the classification rates produced by the CM and MCM methods in Tables 4–5. Finally, a summary of performance evaluation in classification among these five methods is provided in Table 6.

In order to produce 3-D ROC curves for the gray scale fractional abundance images in Figs. 4, 5 and 7, we applied the $a\%$ threshold method to these images. Figures 9–11 show their 3-D ROC curves and their associated 2-D curves for SNR = 5 db, 10 db, 15 db and 20 db respectively where figures labeled by (a)–(d) are 3-D ROC curves of $(R_D, R_F, a\%)$, 2-D curves of (R_D, R_F), $(R_D, a\%)$ and $(R_F, a\%)$ respectively.
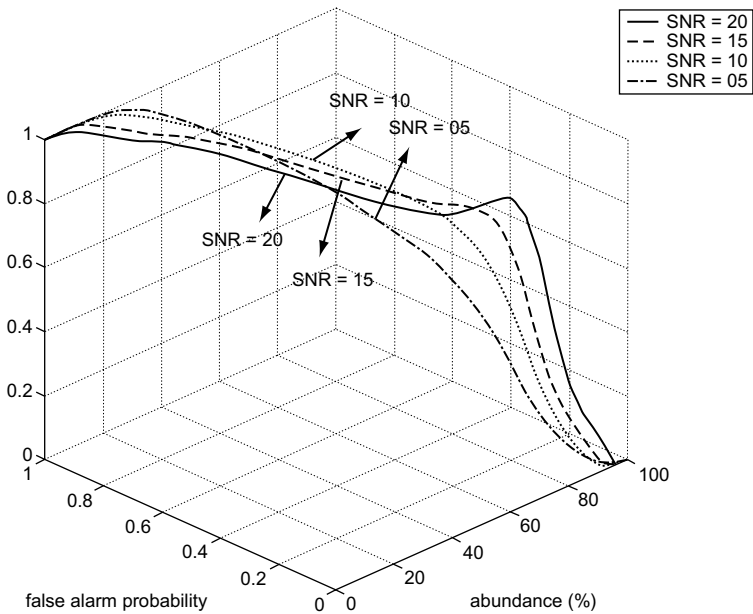
(a) 3-D ROC curve
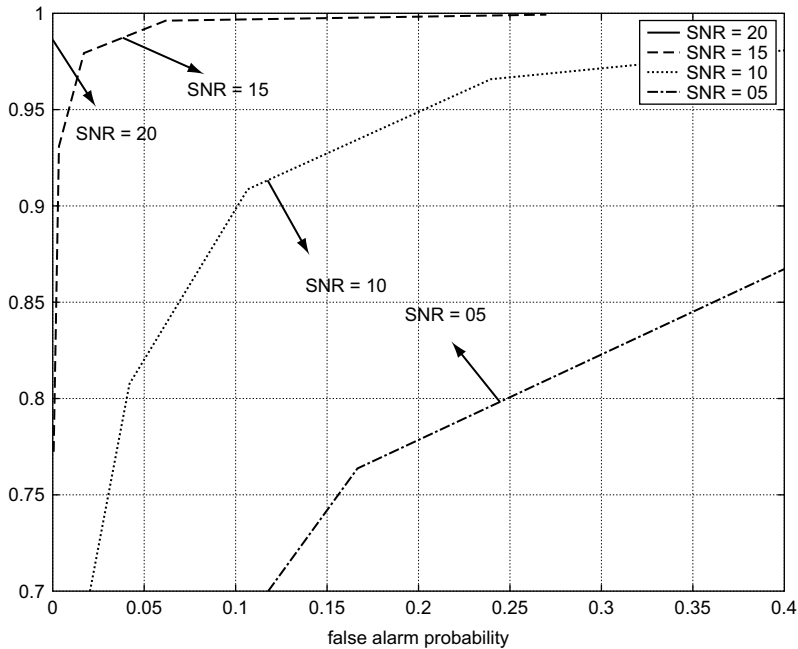
(b) 2-D curve of (R_D,R_F)

Fig. 9.   3-D ROC and 2-D curves produced by OSP from Fig. 4.

(c) 2-D curves of $(R_D, a\%)$

(d) 2-D curves of $(R_F, a\%)$

Fig. 9.　(*Continued*)

(a) 3-D ROC curve

(b) 2-D curve of (R_D,R_F)
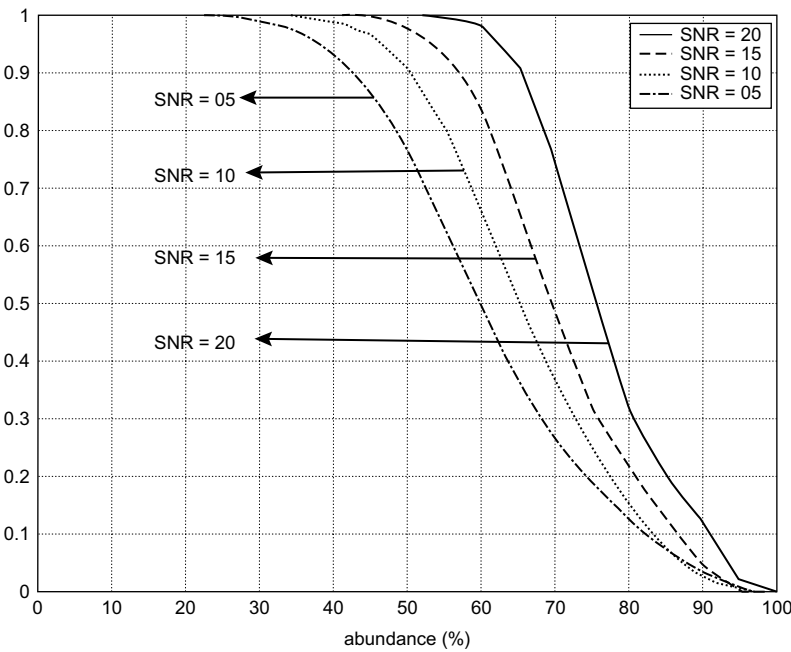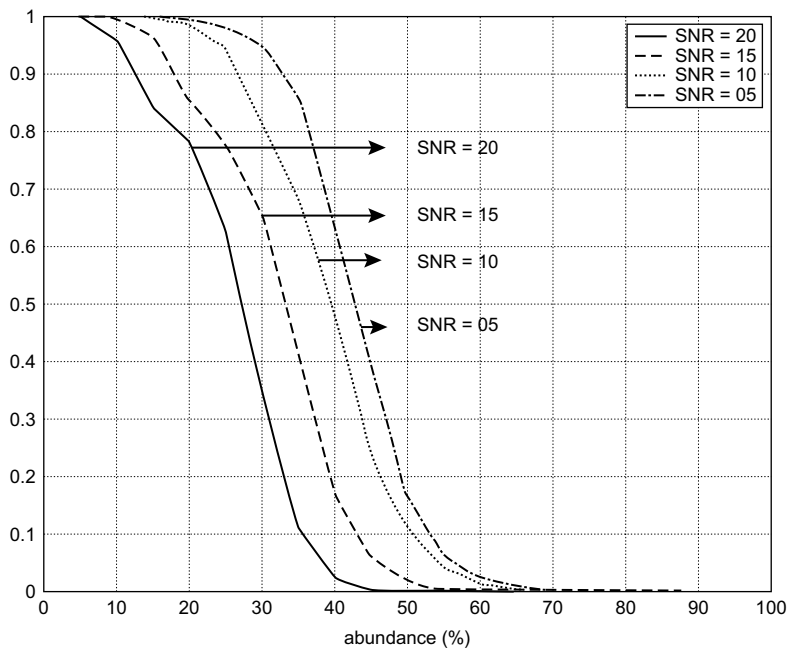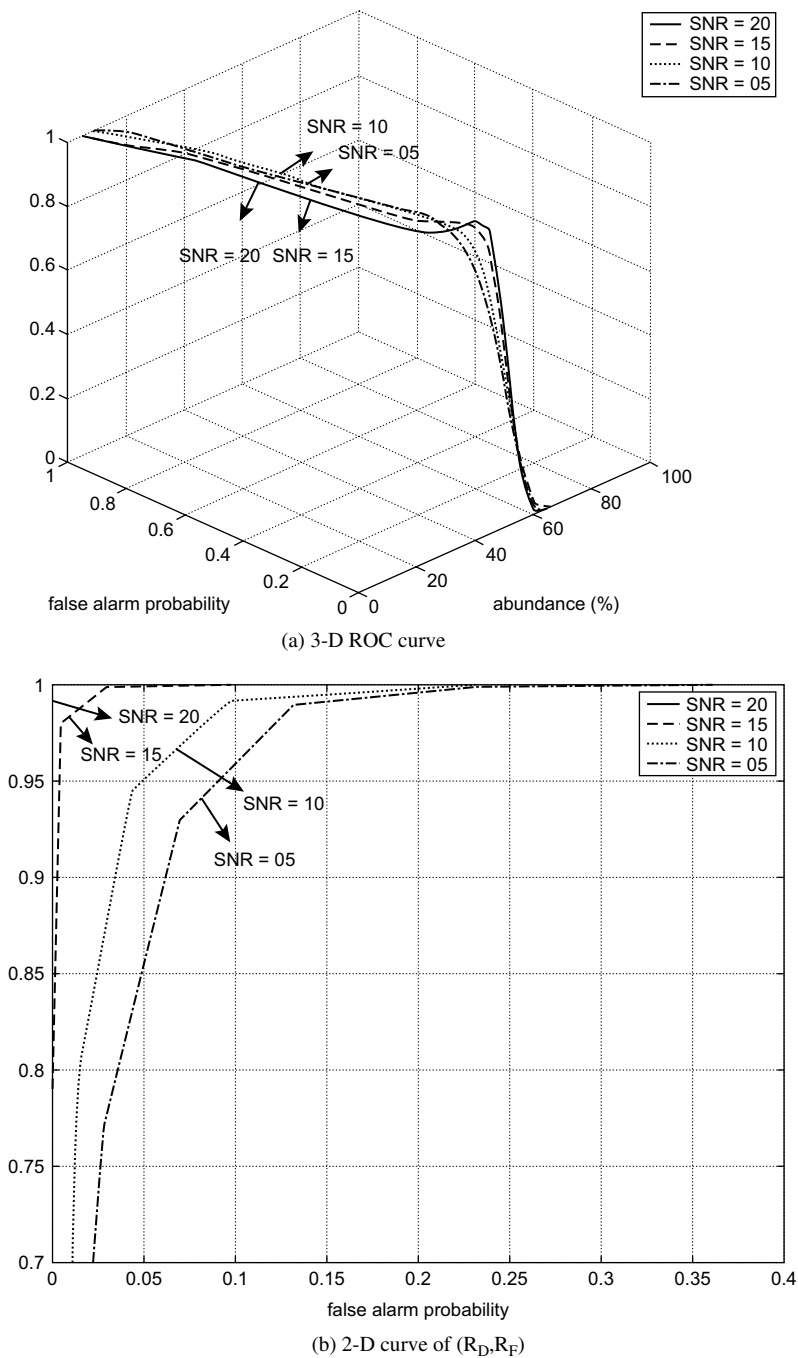
Fig. 10.   3-D ROC and 2-D curves produced by CEM from Fig. 5.

(c) 2-D curves of $(R_D, a\%)$

(d) 2-D curves of $(R_F, a\%)$

Fig. 10.   (*Continued*)

(a) 3-D ROC curve
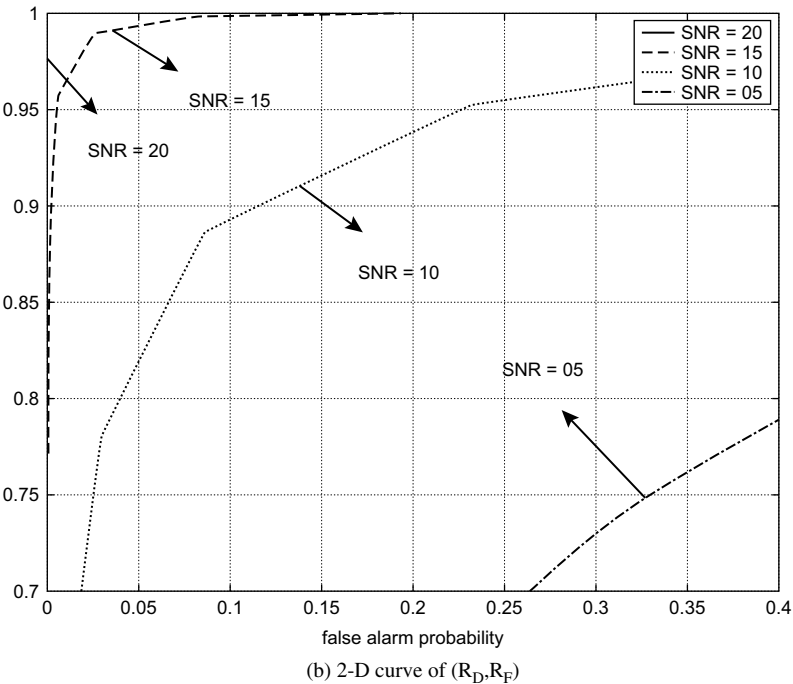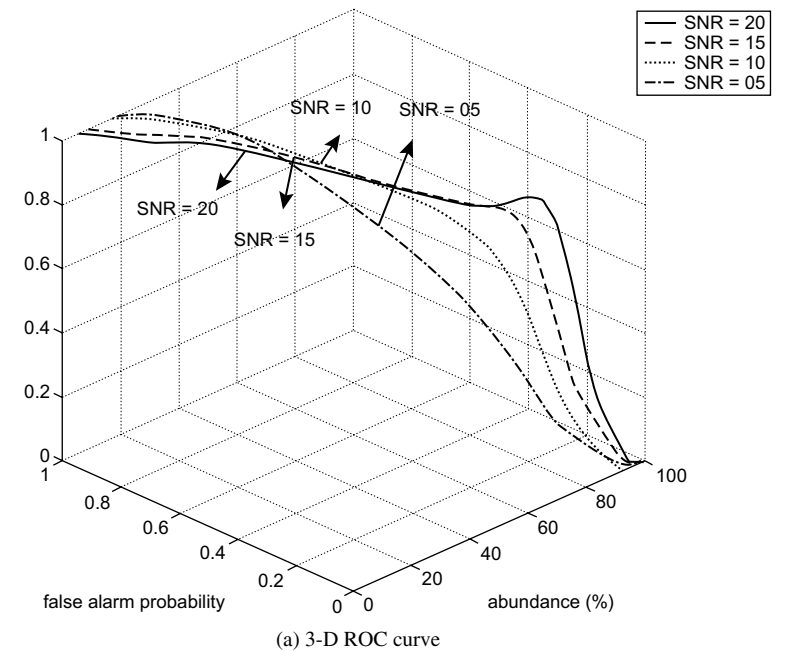
(b) 2-D curve of (R_D,R_F)

Fig. 11. 3-D ROC and 2-D curves produced by UOSP from Fig. 7.
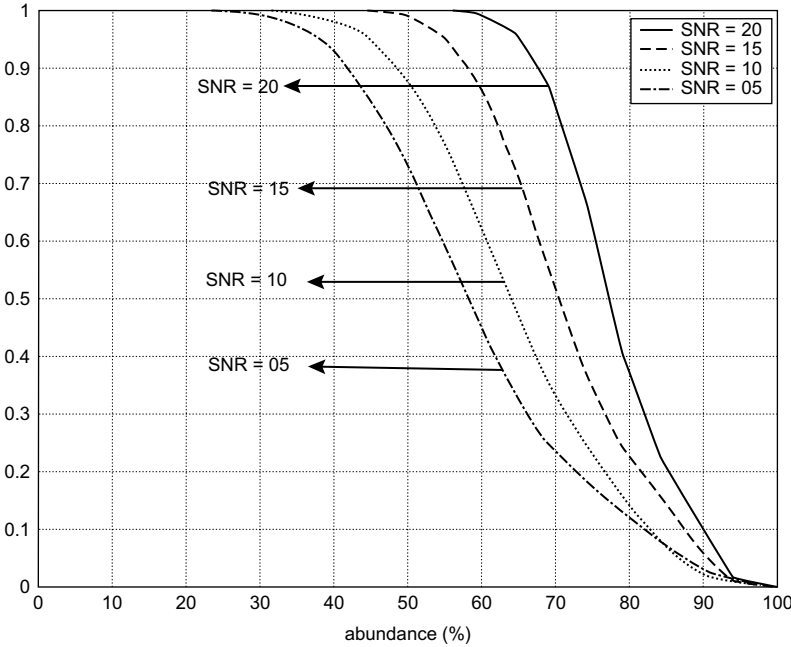
(c) 2-D curves of (R_D,a%)
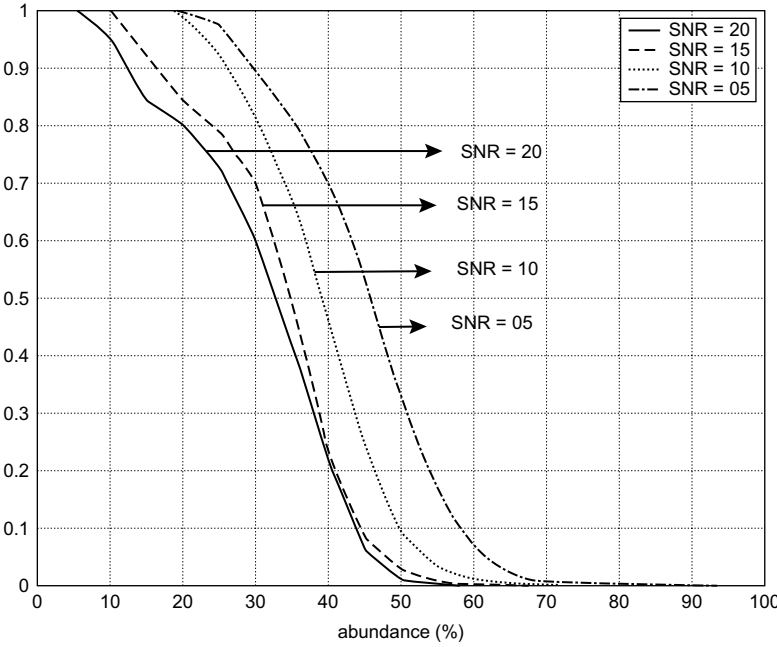
(d) 2-D curves of (R_F,a%)

Fig. 11.   (*Continued*)

Table 2(a). Classification result of GM, WM, and CSF by OSP for the case of SNR = 5 db with $a\%$ chosen to be 5%, 20%, 25%, 30%, 35%, 40%, 45%, and 50%.

OSP	$a(\%)$	N	N_D	N_F	$R_D(\%)$	$R_F(\%)$
GM	5%	9040	9040	56493	100.00	100.00
	20%	9040	9038	56299	99.98	99.65
	25%	9040	9030	55723	99.89	98.63
	30%	9040	9005	53948	99.61	95.49
	35%	9040	8925	49896	98.73	88.32
	40%	9040	8719	42501	96.45	75.23
	45%	9040	8284	30618	91.64	54.20
	50%	9040	7488	12107	82.83	21.43
WM	5%	8745	8745	56786	100.00	99.99
	20%	8745	8738	56389	99.92	99.29
	25%	8745	8696	55675	99.44	98.04
	30%	8745	8571	53834	98.01	94.79
	35%	8745	8267	48614	94.53	85.60
	40%	8745	7654	33820	87.52	59.55
	45%	8745	6694	19431	76.55	34.22
	50%	8745	5336	9098	61.02	16.02
CSF	5%	3282	3282	62253	100.00	100.00
	20%	3282	3282	61768	100.00	99.22
	25%	3282	3282	60551	100.00	97.26
	30%	3282	3282	58067	100.00	93.27
	35%	3282	3282	48732	100.00	78.28
	40%	3282	3282	21340	100.00	34.28
	45%	3282	3282	10128	100.00	16.27
	50%	3282	3281	3532	99.97	5.67

As shown in Figs. 9(a), 10(a) and 11(a), the 3-D ROC curves are the performance of a classifier as a function of three parameters R_D, R_F, $a\%$, whereas the 2-D curves of (R_D,R_F) in Figs. 9(b), 10(b) and 11(b) provide the mean detection rate of a classifier versus the mean false alarm rate. It should be noted that the 2-D curves of (R_D,R_F) in Figs. 9(b), 10(b) and 11(b) were plotted in the ranges of $R_D = [0.7, 1]$ and $R_F = [0, 0.4]$ for visual inspection. According to the 2-D curves in Figs. 9(b), 10(b) and 11(b), the OSP, CEM and UOSP performed extremely well when SNR = 15 dB and 20 dB. Then its performance was degraded when SNR was decreased. Additionally, the 2-D curves of $(R_D,a\%)$ and $(R_F,a\%)$ in Figs. 9(c–d), 10(c–d) and 11(c–d) indicate how a threshold value of $a\%$ affected the performance of the OSP, CEM and UOSP. Figures 9(c), 10(c) and 11(c) show that the 2-D curves of $(R_D,a\%)$ of the OSP and UOSP behaved similarly and their R_D dropped in a very wide range with $a\%$ increased to from 30% to 90% compared to the 2-D curves of $(R_D,a\%)$ of the CEM where the R_D dropped very rapidly with $a\%$ increased to from 30% to 60%. Figures 9(d), 10(d) and 11(d) also demonstrates similar results but the differences among these four SNRs were more visible. It clearly shows that how the R_F of the OSP, CEM and UOSP filter were affected by different SNRs.

Table 2(b).   Classification result of GM, WM, and CSF by CEM
for the case of SNR = 5 db with $a\%$ chosen to be 5%, 20%, 25%,
30%, 35%, 40%, 45%, and 50%.

CEM	$a(\%)$	N	N_D	N_F	$R_D(\%)$	$R_F(\%)$
GM	5%	9040	9040	56470	100.00	99.95
	20%	9040	9040	31480	100.00	55.72
	25%	9040	9040	14664	100.00	25.96
	30%	9040	9026	11249	99.85	19.91
	35%	9040	8896	8598	98.41	15.22
	40%	9040	8150	4806	90.16	8.51
	45%	9040	6328	1905	70.00	3.37
	50%	9040	3551	377	39.28	0.67
WM	5%	8745	8745	56780	100.00	99.98
	20%	8745	8745	54305	100.00	95.62
	25%	8745	8745	45698	100.00	80.47
	30%	8745	8740	20278	99.94	35.71
	35%	8745	8669	9016	99.13	15.88
	40%	8745	8150	4534	93.20	7.98
	45%	8745	6682	1876	76.41	3.30
	50%	8745	4015	440	45.91	0.77
CSF	5%	3282	3282	62221	100.00	99.95
	20%	3282	3282	17074	100.00	27.43
	25%	3282	3282	1907	100.00	3.06
	30%	3282	3282	109	100.00	0.18
	35%	3282	3282	1	100.00	0.00
	40%	3282	3280	0	99.94	0.00
	45%	3282	3236	0	98.60	0.00
	50%	3282	2687	0	81.87	0.00

In order to further provide quantitative analysis, we calculated the classification rates produced by these three techniques, OSP, CEM and UOSP for SNR = 5 dB in Table 2(a–c) and 20 dB in Table 3(a–c) where the cutoff threshold value of $a\%$ was chosen to be 5%, 20%, 25%, 30%, 35%, 40%, 45% and 50%. Since the cases of SNR = 10 db and 15 db could be conducted similarly, their results are not included here. From Table 2(a–c) and Table 3(a–c), we can see that the CEM performed better than the OSP and UOSP when the $a\%$ is small. When the $a\%$ increased, the false alarm rate of the CEM also increased significantly. As a result, its classification was deteriorated rapidly. Interestingly, it was not true for the OSP and UOSP, in which case they actually performed better than the CEM. Nevertheless, the overall mean classification rates produced by the CEM were still better than those produced by the OSP and UOSP as demonstrated in Table 6.

It is interesting to compare the results in Table 2(a–c) and Table 3(a–c) to the results obtained by the traditional classification methods, CM and MCM methods. Table 4 and Table 5 tabulate the classification results of the CM and MCM methods for comparison. As we can see, the MCM method performed much better than the CM method.

Table 2(c). Classification result of GM, WM, and CSF by UOSP for the case of SNR = 5 db with $a\%$ chosen to be 5%, 20%, 25%, 30%, 35%, 40%, 45%, and 50%.

UOSP	$a(\%)$	N	N_D	N_F	$R_D(\%)$	$R_F(\%)$
GM	5%	9040	9040	56495	100.00	100.00
	20%	9040	9039	56371	99.99	99.78
	25%	9040	9020	55957	99.78	99.05
	30%	9040	8982	54519	99.36	96.50
	35%	9040	8861	50305	98.02	89.04
	40%	9040	8532	41021	94.38	72.61
	45%	9040	7929	26699	87.71	47.26
	50%	9040	6856	12073	75.84	21.37
WM	5%	8745	8745	56789	100.00	100.00
	20%	8745	8733	56578	99.86	99.63
	25%	8745	8698	56216	99.46	98.99
	30%	8745	8604	55397	98.39	97.55
	35%	8745	8374	53681	95.76	94.52
	40%	8745	7805	50015	89.25	88.07
	45%	8745	6858	43927	78.42	77.35
	50%	8745	5518	30673	63.10	54.01
CSF	5%	3282	3282	62221	100.00	99.95
	20%	3282	3282	59843	100.00	96.13
	25%	3282	3282	53431	100.00	85.83
	30%	3282	3282	27776	100.00	44.62
	35%	3282	3282	10916	100.00	17.54
	40%	3282	3282	3166	100.00	5.09
	45%	3282	3282	752	100.00	1.21
	50%	3282	3282	124	100.00	0.20

Since 2-D curves of (R_D, R_F) is similar to the 2-D ROC curve commonly used in detection theory, we can calculate the area under the 2-D curve of (R_D, R_F)[39] to quantitatively study the overall performance of the OSP, CEM and UOSP. The first three rows of Table 6 tabulate the mean classification rates calculated from the areas under 2-D curves of (R_D, R_F) in Figs. 9(b), 10(b) and 11(b). The last two rows of Table 6 also tabulate the classification results of the MCM and CM methods for SNR = 5 dB, 10 dB, 15 dB, 20 dB. It should be noted that no ROC curves can be generated by the MCM and CM methods since the both MCM and CM methods are labeling processes and each SNR results in a fixed point specified by one and only one pair (R_D, R_F).

7. Real MR Image Experiments

In this section, real magnetic resonance (MR) images were used for experiments which allow us to assess its utility and effectiveness in medical diagnosis. They are shown in Fig. 3 and were acquired from ten patients with normal physiology and no intensity inhomogeneity correct method was applied to the MR image data. Band 1 in Fig. 3(a) is the PD-weighted spectral image acquired by the pulse

Table 3(a). Classification result of GM, WM, and CSF by OSP for the case of SNR = 20 db with a% chosen to be 5%, 20%, 25%, 30%, 35%, 40%, 45%, and 50%.

OSP	$a(\%)$	N	N_D	N_F	$R_D(\%)$	$R_F(\%)$
GM	5%	9040	9040	56490	100.00	99.99
	20%	9040	9040	51725	100.00	91.56
	25%	9040	9040	41060	100.00	72.68
	30%	9040	9040	21287	100.00	37.68
	35%	9040	9040	5840	100.00	10.34
	40%	9040	9040	1173	100.00	2.08
	45%	9040	9039	239	99.99	0.42
	50%	9040	9038	28	99.98	0.05
WM	5%	8745	8745	56775	100.00	99.97
	20%	8745	8745	53526	100.00	94.25
	25%	8745	8745	43669	100.00	76.89
	30%	8745	8745	23953	100.00	42.18
	35%	8745	8745	8522	100.00	15.01
	40%	8745	8745	1930	100.00	3.40
	45%	8745	8745	283	100.00	0.50
	50%	8745	8740	26	99.94	0.05
CSF	5%	3282	3282	61834	100.00	99.33
	20%	3282	3282	21	100.00	0.03
	25%	3282	3282	0	100.00	0.00
	30%	3282	3282	0	100.00	0.00
	35%	3282	3282	0	100.00	0.00
	40%	3282	3282	0	100.00	0.00
	45%	3282	3282	0	100.00	0.00
	50%	3282	3282	0	100.00	0.00

sequence TR/TE = 2500 ms/25 ms. Bands 2, 3 and 4 in Fig. 3(b–d) are T2-weighted spectral images were acquired by the pulse sequences TR/TE = 2500 ms/50 ms, TR/TE = 2500 ms/75 ms and TR/TE = 2500 ms/100 ms respectively. Band 5 in Fig. 3(e) is the T1-weighted spectral image acquired by the pulse sequence TR/TE = 500 ms/11.9 ms. The tissues surrounding the brain such as bone, fat, skin, were semiautomatically extracted using interactive thresholding and masking.[40] The slice thickness of all the MR images are 6 mm and axial section were taken from GE MR 1.5T scanner. Before acquisition of the MR images the scanner was adjusted to prevent artifacts caused by the magnetic field of static, radio frequency and gradient. All experiments presented in this chapter were performed under supervision of and verified by experienced neuroradiologists.

In many MRI applications, the three cerebral tissues, GM, WM and CSF are of major interest where their knowledge can be generally obtained directly from the images. In the following experiments, the spectral signatures of GM, WM and CSF used for the CEM were extracted directly from the MR images and verified by experienced radiologists. The number of classes was assumed to be 4, i.e. $c = 4$. In analogy with computer simulation experiments in Sec. 6, these signatures were used as prior knowledge for the OSP and CEM. When the OSP was implemented, the

Table 3(b).   Classification result of GM, WM, and CSF by CEM for the case of SNR = 20 db with a% chosen to be 5%, 20%, 25%, 30%, 35%, 40%, 45%, and 50%.

CEM	$a(\%)$	N	N_D	N_F	$R_D(\%)$	$R_F(\%)$
GM	5%	9040	9040	56465	100.00	99.95
	20%	9040	9040	16962	100.00	30.02
	25%	9040	9040	3651	100.00	6.46
	30%	9040	9040	464	100.00	0.82
	35%	9040	9039	15	99.99	0.03
	40%	9040	9019	0	99.77	0.00
	45%	9040	8623	0	95.39	0.00
	50%	9040	6275	0	69.41	0.00
WM	5%	8745	8745	56756	100.00	99.94
	20%	8745	8745	27183	100.00	47.87
	25%	8745	8745	5326	100.00	9.38
	30%	8745	8745	362	100.00	0.64
	35%	8745	8745	10	100.00	0.02
	40%	8745	8719	1	99.70	0.00
	45%	8745	8139	0	93.07	0.00
	50%	8745	5189	0	59.34	0.00
CSF	5%	3282	3282	59741	100.00	95.96
	20%	3282	3282	0	100.00	0.00
	25%	3282	3282	0	100.00	0.00
	30%	3282	3282	0	100.00	0.00
	35%	3282	3282	0	100.00	0.00
	40%	3282	3282	0	100.00	0.00
	45%	3282	3282	0	100.00	0.00
	50%	3282	2411	0	73.46	0.00

signature matrix $\mathbf{M}$ consisted of signatures of GM, WM and CSF plus a signature extracted from the background. When the CEM was implemented, the desired signature $\mathbf{d}$ was selected to be one of the three GM, WM and CSF signatures to be detected. Similarly, the MCM method was implemented, the $p = 1$ and $\boldsymbol{\mu}_1^{(k)}$ was always fixed at $\mathbf{d}$ during iterations, i.e. $\boldsymbol{\mu}_1^{(k)} = \mathbf{d}$.

Following the same manner conducted for MR phantom image experiments, we used five images in Fig. 3(a–e) with the desired object signatures specified in Fig. 2. Figures 12(a–c), 13(a–c), 14(a–c), 15(a–c) and 16(a–c) show the classification results of the OSP, CEM, MCM method, UOSP and CM method for GM, WM and CSF where the images labeled by (a), (b) and (c) were the classification maps for GM, WM and CSF. As noted, the MCM and CM methods were not stable due to its nature in unsupervised learning. When each time both methods were implemented, different classification maps were generated. The results in Figs. 14 and 16 were obtained by averaging five runs of implementation of the MCM and CM methods. Note that when the MCM method proposed in Sec. 5 was implemented, the desired object signature was designated as one specific class and this class was fixed during its unsupervised clustering.

Table 3(c). Classification result of GM, WM, and CSF by UOSP for the case of SNR = 20 db with a% chosen to be 5%, 20%, 25%, 30%, 35%, 40%, 45%, and 50%.

UOSP	$a(\%)$	N	N_D	N_F	$R_D(\%)$	$R_F(\%)$
GM	5%	9040	9040	56490	100.00	99.99
	20%	9040	9040	52675	100.00	93.24
	25%	9040	9040	43626	100.00	77.22
	30%	9040	9040	26005	100.00	46.03
	35%	9040	9040	7400	100.00	13.10
	40%	9040	9040	1347	100.00	2.38
	45%	9040	9039	271	99.99	0.48
	50%	9040	9037	40	99.97	0.07
WM	5%	8745	8745	56760	100.00	99.95
	20%	8745	8745	54450	100.00	95.88
	25%	8745	8745	53699	100.00	94.56
	30%	8745	8745	53063	100.00	93.44
	35%	8745	8745	47535	100.00	83.70
	40%	8745	8745	25960	100.00	45.71
	45%	8745	8745	7809	100.00	13.75
	50%	8745	8745	1005	100.00	1.77
CSF	5%	3282	3282	61859	100.00	99.37
	20%	3282	3282	21	100.00	0.03
	25%	3282	3282	0	100.00	0.00
	30%	3282	3282	0	100.00	0.00
	35%	3282	3282	0	100.00	0.00
	40%	3282	3282	0	100.00	0.00
	45%	3282	3282	0	100.00	0.00
	50%	3282	3282	0	100.00	0.00

Table 4. Classification result of GM, WM, and CSF by CM method for SNR = 5 db and 20 db.

	SNR	N	N_D	N_F	$R_D(\%)$	$R_F(\%)$
GM	5 db	9040	0	4886	0.00	8.65
	20 db	9040	0	0	0.00	0.00
WM	5 db	8745	1376	5	15.74	0.01
	20 db	8745	8745	8504	100.00	14.97
CSF	5 db	3282	3282	16404	100.00	26.35
	20 db	3282	3282	536	100.00	0.86

Table 5. Classification result of GM, WM, and CSF by MCM method for SNR = 5 db and 20 db

	SNR	N	N_D	N_F	$R_D(\%)$	$R_F(\%)$
GM	5 db	9040	8708	6277	96.33	11.11
	20 db	9040	9040	7489	100.00	13.26
WM	5 db	8745	8517	6201	97.39	10.92
	20 db	8745	8745	9285	100.00	16.35
CSF	5 db	3282	2941	4003	89.61	6.43
	20 db	3282	3166	4001	96.47	6.43

Table 6. The mean dclassification rates calculated from the areas under 2-D curves of (R_D, R_F) in Figs. 9(b), 10(b), 11(b).

	SNR = 5 db	SNR = 10 db	SNR = 15 db	SNR = 20 db
OSP	0.8621	0.9638	0.9984	1.0000
CEM	0.9806	0.9895	0.9995	1.0000
UOSP	0.8027	0.9607	0.9987	1.0000
CM	0.5628	0.5926	0.7586	0.7356
MCM	0.9572	0.9830	0.9889	0.9945

(a) GM   (b) WM   (c) CSF
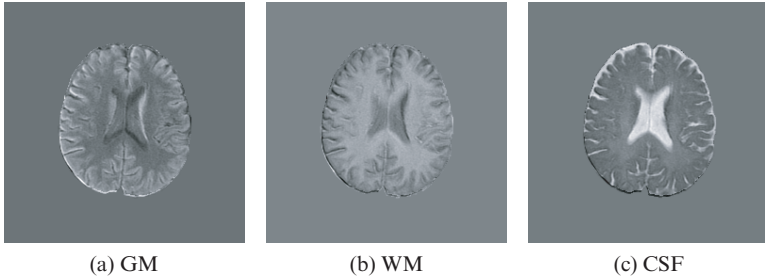
Fig. 12.   Classification results of GM, WM and CSF produced by OSP.

(a) GM   (b) WM   (c) CSF
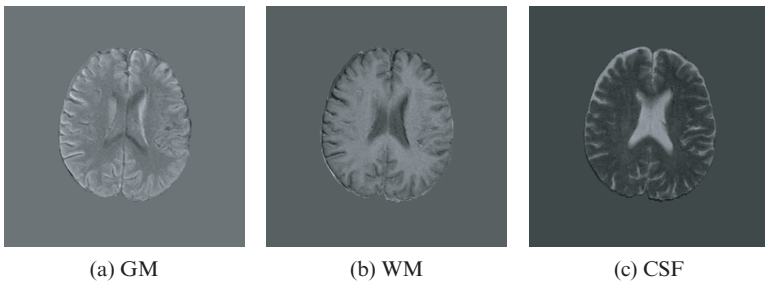
Fig. 13.   Classification results of GM, WM and CSF produced by CEM.

(a) GM   (b) WM   (c) CSF
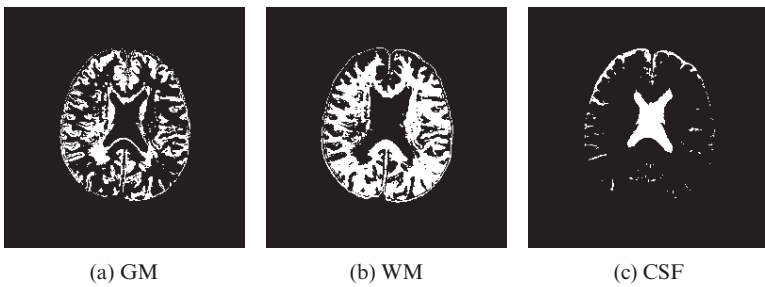
Fig. 14.   Classification results of GM, WM and CSF produced by the MCM method.

(a) GM   (b) WM   (c) CSF
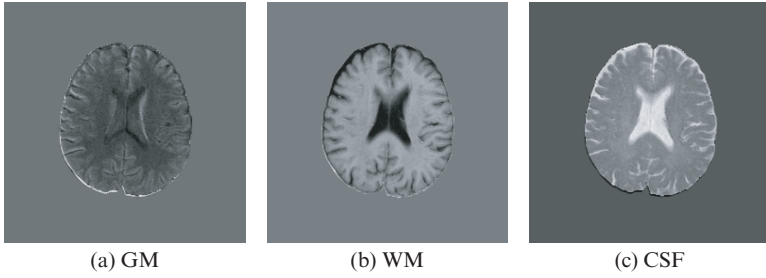
Fig. 15.   Classification results of GM, WM and CSF produced by UOSP.

(a) GM   (b) WM   (c) CSF
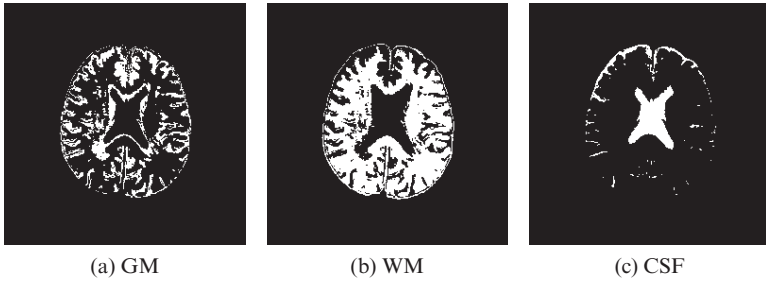
Fig. 16.   Classification results of GM, WM and CSF produced by the CM method.

In the MR phantom image experiments conducted in Sec. 6, Gaussian noise was simulated to achieve various SNR for quantitative analysis. Unfortunately, a quantitative study will be difficult for the above real MR image experiments for the following two reasons. One is that it requires reliable techniques to estimate noise in the MR images. This has been a challenging issue in signal and image processing[41−42] and beyond the scope of this chapter. The OSP, CEM and UOSP generate gray scale abundance fractional images for MR image classification which provide radiologists with gray level information for their visual interpretation. Such qualitative information is useful for medical diagnosis, but will be lost if gray scale images are converted to binary images by thresholding. In addition, it is nearly impossible for radiologists to identify all the pixels in real MR images for quantitative study as the way we did for phantom images where we knew exactly what class to which each pixel belonged. As a consequence, no quantitative analysis was conducted for the real MR image experiments.

8. Conclusions

This chapter presents applications of three hyperspectral image processing techniques, Orthogonal Subspace Projection (OSP), Constrained Energy Minimization (CEM) and Unsupervised OSP (UOSP) to MR image classification. These techniques require different levels of information that can be used for various scenarios in

different applications.[33] Unlike classical image classification techniques which perform classification based on spatial correlation on a pure pixel basis, the proposed techniques are designed to explore spectral characteristics of signatures of interest so as to achieve better classification. Most importantly, the resulting images are not traditional class-map images, but rather fractional abundance images which only show signatures of interest with various gray scales, each of which is proportional to the abundance fraction of a particular signature contained in an image pixel vector. In order to evaluate their performance, the classical 2-D ROC analysis is also extended to a 3-D ROC analysis which is based on three parameters, detection rate, false alarm rate and abundance fraction percentage. This 3-D ROC analysis is a new concept and has not been explored in medical diagnosis. It is particularly useful for MR image classification by including signature abundance fractions as a third dimension in ROC curves where spectral information can be characterized separately from spatial information.

References

1. G. A. Wright, Magnetic resonance image, *IEEE Signal Process. Magazine* (1997) 56–66.
2. G. Sebastiani and P. Barone, Mathematical principles of basic magnetic resonance image in medicine, *Signal Process* **25** (1991) 227–250.
3. B. Johnston, M. S. Atkins, B. Mackiewich and M. Anderson, Segmentation of multiple sclerosis lesions in intensity corrected multispectral MRI, *IEEE Trans. Medical Imaging* **15**(2) (1996) 154–169.
4. A. P. Dhawan, A. Zavaljevski, A. Sarwal and W. S. Ball, A system for MR brain image segmentation, *18th Annual Int. Conf. IEEE Eng. Med. Biol. Soc.*, Amsterdam (1996), pp. 732–733.
5. L. Verard, J. Fadidi, S. Ruan and D. Bloyet, 3D MRI segmentation of brain structures, *18th Annual Int. Conf. IEEE Eng. Med. Bio. Soc.*, Amsterdam (1996), pp. 1081–1082.
6. H. Grahn, N. M. Szeverenyi, N. W. Roggenbuck, F. Delaglio and P. Geladi, Data analysis of multivariate magnetic resonance images I. a principal component analysis approach, *Chemometrics and Intel. Lab. Systems* **5** (1989) 11–322.
7. J. P. Windham, M. A. Abd-Allah, D. A. Reimann, J. W. Froehich and A. M. Haggar, Eigneimage filtering in MR imaging, *J. Computer Assisted Tomography* **12**(1) (1988) 1–9.
8. H. Soltanian-Zadeh and J. P. Windham, Novel and general approach to linear filter designed for contrast-to-noise ratio enhancement of magnetic resonance images with multiple interfering features in the scene, *J. Electronic Imaging* **1**(2) (1992) 171–182.
9. A. M. Haggar, J. P. Windham D. A. Reimann, D. C. Hearshen and J. W. Froehich, Eigneimage filtering in MR imagine: an application in the abnormal chest wall, *Magnetic Resonance in Medicine* **11** (1989) 85–97.
10. H. Soltanian-Zadeh, J. P. Windham, H. Soltanian-Zadeh, J. P. Windham and A. E. Yagle, A comparative analysis of several transformations for enhancement and segmentation of magnetic resonance image scene sequences, *IEEE Trans. Medical Imaging* **11**(3) (1992) 302–316.

11. H. Soltantian-Zadeh, R. Saigal, A. M. Haggar, J. P. Windham, A. E. Yagle and D. C. Hearshen, Optimization of MRI protocola and pulse sequence parameters for eigneimage filtering, *IEEE Trans. Medical Imaging* **13**(1) (1994) 161–175.

12. H. Soltanian-Zadeh, J. P. Windham and D. J. Peck, Optimal linear transformation for MRI feature extraction, *IEEE Trans. Medical Imaging* **15**(6) (1996) 749–767.

13. W. E. Reddick, J. O. Glass, E. N. Cook, T. D. Elkin and R. J. Deaton, Automated segmentation and classification of multispectral magnetic Resonance images of brain using artificial neural networks, *IEEE Trans. Medical Imaging* **16**(6) (1997) 911–918.

14. J. Alirezaie, M. E. Jernigan and C. Nahmias, Automatic segmentation of cerebral MR images using artificial neural networks, *IEEE Trans. Nuclear Science* **45**(4) (1998) 2174–2182.

15. J. S. Lin, R. M. Chen and Y. M. Huang, Medical image sementation using field annealing network, *IEEE Int. Conf. Image Processing* **2** (1997) 885–858.

16. J. S. Lin, K. S. Cheng and C. W. Mao, Multispectral magnetic resonance images segmentation using fuzzy Hopfield neural network, *Int. J. Biomedical Computing* (1996) 205–214.

17. J. S. Lin, K. S. Cheng and C. W. Mao, Modified Hopfied neural network with fuzzy c-means technique for multispectral MR image segmentation, *IEEE Int. Conf. Image Processing* **1** (1996) 327–330.

18. M. C. Clark, L. O. Hall, D. B. Goldgof, L. P. Clarke, R. P. Velthuizen and M. S. Sibiger, MRI segmentation using Fuzzy clustering techniques, *IEEE Eng. Med. Biol.* (1994) 730–742.

19. M. S. Atkins and B. T. Mackiewich, Fully automatic segmentation of brain in MRI, *IEEE Trans. Medical Imaging* **17**(1) (1998) 98–107.

20. M. C. Clark, L. O. Hall, D. B. Goldgof, R. Velthuizen, F. R. Murtagh and S. Silbiger, Automatic tumor segmentation using knowledge-based techniques, *IEEE Trans. Medical Imaging* **17**(2) (1998) 187–201.

21. C.-M. Wang, S.-T. Yang, P.-C. Chung, C. S. Lo, C.-I Chang, C. C. Chen, C.-W. Yang and C. H. Wen, Orthogonal subspace projection-based approaches to classification of MR image sequences, *Computerized Medical Imaging and Graphics* **25**(6) (2001) 465–476.

22. C. M. Wang, C. C. Chen, S.-C. Yang, Y.-N. Chung, P. C. Chung, C. W. Yang and C.-I Chang, An unsupervised Orthogonal subspace projection approach to MR image classification MR images for classification, *Optical Eng.* **41**(7) (2002) 1546–1557.

23. C. M. Wang, C. C. Chen, Y.-N. Chung, S.-C. Yang, P. C. Chung, C. W. Yang and C.-I Chang, Detection of spectral signatures in MR images for classification, *IEEE Trans. Medical Imaging* TMI-**22**(1) (2003) 50–61.

24. C. M Wang, *Applications of Multispectral Imaging Processing Techniques to Brain Magnetic Resonance Image Classification*, PhD dissertation, Department of Electrical Engineering, National Cheng Kung University, Tainan, Taiwan, R.O.C. (June 2002).

25. J. C. Harsanyi, *Detection and Classification of Subpixel Spectral Signatures in Hyperspectral Image Sequences*, PhD dissertation, Department of Electrical Engineering, University of Maryland Baltimore County, Baltimore, MD (1993).

26. J. C. Harsanyi and C.-I Chang, Hyperspectral image classification and dimensionality reduction: an orthogonal subspace projection approach, *IEEE Trans. Geoscience and Remote Sensing* **32**(4) (1994) 779–785.

27. J. C. Harsanyi, W. Farrand and C.-I Chang, Detection of subpixel spectral signatures in hyperspectral image sequences, *Annual Meeting, Proc. Am. Soc. Photogrammetry & Remote Sensing*, Reno (1994), pp. 236–247.

28. R. S. Resmini, M. E. Kappus, W. S. Aldrich, J. C. Harsanyi and M. Anderson, Mineral mapping with HYperspectral Digital Imagery Collection Experiment (HYDICE) sensor data at Cuprite, Nevada, U.S.A., *Int. J. Remote Sensing* **18**(17) (1997) 1553–1570.

29. W. Farrand and J. C. Harsanyi, Mapping the distribution of mine tailing in the coeur d'Alene river valley, Idaho, through the use of constrained energy minimization technique, *Remote Sensing of Environment* **59** (1997) 64–76.

30. C.-I Chang, Target signature-constrained mixed pixel classification for hyperspectral imagery, *IEEE Trans. Geoscience and Remote Sensing* **40**(5) (2002) 1065–1081.

31. C.-I Chang, Further results on relationship between spectral unmixing and subspace projection, *IEEE Trans. Geoscience and Remote Sensing* **36**(3) (1998) 1030–1032.

32. C.-I Chang, X. Zhao, M. L. G. Althouse and J.-J. Pan, Least squares subspace projection approach to mixed pixel classification in hyperspectral images, *IEEE Trans. Geoscience and Remote Sensing* **36**(3) (1998) 898–912.

33. C.-I Chang, *Hyperspectral Imaging: Techniques for Spectral Detection and Classification* (Kluwer Academic/Plenum Publishers, 2003).

34. H. Ren and C.-I Chang, Automatic spectral target recognition in hyperspectral imagery, *IEEE Trans. Aerospace and Electronic Systems* **39**(4) (2003) 1232–1249.

35. R. Schalkoff, *Pattern Recognition: Statistical, Structural and Neural Approaches* (John Wiley and Sons, 1992).

36. R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis* (John Wiley and Sons, 1973).

37. C.-I Chang and H. Ren, An experiment-based quantitative and comparative analysis of hyperspectral target detection and image classification algorithms, *IEEE Trans Geoscience and Remote Sensing* **38** (2000) 1044–1063.

38. C.-I Chang, H. Ren, Q. Du, S.-S. Chiang and A. Ifarraguerri, An ROC analysis for subpixel detection, *IEEE 2001 Int. Geoscience and Remote Sensing Symp.*, Sydney, Australia (July 24–28, 2001).

39. C. E. Metz, ROC methodology in radiological imaging, *Invest. Radiology* **21** (1986) 720–723.

40. H. Suzuki and J. Toriwaki, Automatic segmentation of head MRI images by knowledge guided thresholding, *Computerized Med. Imag., Graphics* **15**(4) (1991) 233–240.

41. A. A. Green, M. Berman, P. Switzer and M. D. Craig, A transformation for ordering multispectral data in terms of image quality with implications for noise removal, *IEEE Trans. Geoscience and Remote Sensing* **26**(1) (1988) 65–74.

42. J. B. Lee, A. S. Woodyatt and M. Berman, Enhancement of high spectral resolution remote sensing data by a noise-adjusted principal components transform, *IEEE Trans. Geoscience and Remote Sensing* **28**(3) (1990) 295–304.

This page intentionally left blank

CHAPTER 10

STUDYING ANATOMY AND DISEASE IN MEDICAL IMAGES USING SHAPE ANALYSIS

DANIEL GOLDBERG-ZIMRING[1], DOMINIK S. MEIER[3],
SYLVAIN BOUIX[4] and SIMON K. WARFIELD[1,2]

Computational Radiology Laboratory, Departments of Radiology
[1]Brigham & Women's Hospital, Harvard Medical School
75 Francis St. Boston, MA, 02115 USA

[2]Children's Hospital, Harvard Medical School
300 Longwood Ave. Boston, MA, 02115 USA

[3]Center for Neurological Imaging, Department of Radiology
Brigham & Women's Hospital, Harvard Medical School
221 Longwood Ave. Boston, MA, 02115 USA

[4]Department of Psychiatry, Harvard Medical School
VA Boston Healthcare System, Brockton, MA, 02301 USA
daniel@bwh.harvard.edu
meier@bwh.harvard.edu
sylvain@bwh.harvard.edu
warfield@bwh.harvard.edu

The recent development of image analysis and visualization tools allowing explicit 3D depiction of both normal anatomy and pathology provides a powerful means to obtain morphological descriptions and characterizations. Shape analysis offers the possibility of improved sensitivity and specificity for detection and characterization of structural differences and is becoming an important tool for the analysis of medical images. It provides information about anatomical structures and disease that is not always available from a volumetric analysis. In this chapter we present our own shape analysis work directed to the study of anatomy and disease as seen on medical images. Section 2 presents a new comprehensive method to establish correspondences between morphologically different 3D objects. The correspondence mapping itself is performed in a geometry- and orientation-independent parameter space in order to establish a continuous mapping between objects. Section 3 presents a method to compute 3D skeletons robustly and show how they can be used to perform statistical analyses describing the shape changes of the human hippocampus. Section 4 presents a method to approximate individual MS lesions' 3D geometry using spherical harmonics and its application for analyzing their changes over time by quantitatively characterizing the lesion's shape and depicting patterns of shape evolution.

Keywords: Shape analysis; shape correspondence; parameterization; spherical harmonics; skeletons; medial representations.

1. Introduction

The advent of fast three-dimensional (3D) magnetic resonance imaging (MRI) has enabled the routine acquisition of high spatial resolution digital representations of

anatomy *in vivo.* The recent development of image analysis and visualization tools allowing explicit 3D depiction of both normal anatomy and pathology provides a powerful means to obtain morphological descriptions and characterizations.

Shape analysis offers the possibility of improved sensitivity and specificity for detection and characterization of structural differences. A number of approaches for characterizing shape and shape changes have been developed. These methods differ in the underlying representation of shape, in the type of shapes that may be modeled, and in their capacity to differentiate shape changes. Different shape analysis methods have been proposed for a variety of clinical applications. A landmark based shape analysis has been applied to study the effect of different diseases on brain structures. For this landmark-based approach, shape is defined as the information about the landmark configuration that remains unchanged under adjustment of position, orientation or scale.[1] Landmarks are manually located on the surface of the studied structures seen on magnetic resonance (MR) images. These landmarks are used to create an average shape of the studied structures, which is later compared to each individual structure deriving displacement maps that are helpful to visualize group differences and to discriminate among clinically meaningful categories.[2,3] This approach has been widely used for the study of brain structures in schizophrenia. Buckley *et al.*[2] studied the ventricular dysmorphology in schizophrenia, DeQuardo *et al.*[3] analyzed ventricular enlargement during first-episode schizophrenia, while Tibbo *et al.*[4] studied the corpus callosum shape in male patients with schizophrenia. In a different study landmarks were used to analyze the shape of the corpus callosum and subcortical structures in the fetal-alcohol-affected brain.[1]

Other shape analysis methods based on MR images have also been applied to the study of different brain structures related studies. All of these shape analysis methods are regularly applied after the corresponding segmentation and registration of the images. Velakoulis *et al.*[5] defined the shape of the hippocampus as the volume of contiguous slices and use it to analyze both, the overall and the behind of the head of the hippocampus volume loss in chronic schizophrenia. Levitt *et al.*[6] studied the correlation between the shape of the caudate nucleus and cognition in neuro-naive schizotypal personality disorder. In that study, they generated a 3D rendering of the caudate nucleus and estimated a shape index using the ratio between the surface area and the volume of the caudate nucleus, which indicate how much a given shape differs from a sphere. Hogan *et al.*[7] studied Hippocampal shape analysis in epilepsy and unilateral mesial temporal sclerosis. They used a color scale showing degrees of outward and inward deviation of each studied hippocampus coregistred to a previously generated average hippocampus. Posener *et al.*[8] also analyzed the shape of the hippocampus but in cases of depression, by superimposing a triangulated mesh of the hippocampus on a hippocampus template and representing the displacement of each graphical point in the mesh by a vector. Sowell *et al.*[9] extracted the cortical surface from the MR images and created a 3D mesh-like render in order to analyze brain surface abnormalities in children, adolescents and young adults with prenatal alcohol exposure. They also assessed the relationship between cortical gray matter density on the brain surface and brain shape. The

shape analysis was point-by-point radial distance estimation between each individual cortical 3D surface and an average 3D cortical surface model. Zilles *et al.*[10] analyzed the inter-subject variability of hemispheric shape and defined gender and inter-ethnic differences between the hemispheric shape of male and female European and Japanese brains. They constructed a mean brain model from all brains and for each sample, and later measured the distance between each voxel at the surface of an individual brain and the corresponding surface voxel of the mean brain. The variability was displayed as a function of the standard deviations of the mean absolute distances between all voxels of the mean brain surface and all the corresponding surface voxels of all individual brains. Studholme *et al.*[11] examined a method for the analysis of Jacobian determinant maps capturing local anatomical size differences. They applied this method to the study of dementia in Alzheimer's disease and aging by capturing shape differences between an individual and a reference anatomy or between repeated scans of an individual.

Shape analysis has also been used as a tool to develop or further improve some image segmentation and registration algorithms. Pizer *et al.*[12] constructed a stable efficiently calculable measure of shape and other geometric object properties and applied it in a uniform method for segmentation and recognition of image objects, object-based registration, and object shape measurement. Tsai *et al.*[13] derived a model-based, implicit parametric representation of the segmenting curve and calculated the parameters of this implicit model to minimize the region-based energy function for medical image segmentation. Goldberg-Zimring *et al.*[14] proposed an algorithm for the two-dimensional (2D) segmentation of multiple sclerosis (MS) lesions based on the assumption that MS lesions have a relatively circular shape. After an initial global thresholding of the images, for each detected segment a shape index and the average intensity inside the region of interest was estimated. This shape index indicating the closeness of each segment to a circular shape was used together with the average intensity value as the input for an artificial neural network in order to discriminate between MS lesions and detected artifacts.

While most of shape analysis techniques have taken the advantage of MRI and its high resolution and quality images, this kind of analysis can be applied to studies based on any kind of medical images. For example, Christodoulou *et al.*[15] developed an algorithm for the automated characterization of carotid plaques recorded from high-resolution ultrasound images. Carotid plaques were characterized based on extracted texture features and shape parameters such as the X and Y coordinates maximum length, area, perimeter, and the relation $perimeter^2/area$.

In this chapter we present our own shape analysis work directed to the study of anatomy and disease as seen on medical images. Section 2 presents a new comprehensive method to establish correspondences between morphologically different 3D objects. Section 3 presents a method based on the 3D skeleton to study the effect of aging on the human Hippocampus, and Sec. 4 presents a method to approximate MS lesions' 3D geometry using spherical harmonics (SH), and to analyze MS lesions changes over time.

2. Shape Correspondence

2.1. *The correspondence problem*

Longitudinal and cross-sectional studies of anatomical shape and structures in clinical images have become a common task in diagnosis, treatment evaluation, surgical planning and exploratory clinical research.

Such morphometric analysis of medical images often requires a one-to-one mapping, or correspondence, between anatomic structures in two or more different but not disparate image sets. So the structures of interest will be similar but morphologically different, such as when a changing structure is imaged over time, e.g., in dynamic cardiac imaging,[16,17] the study of disease progression,[18,19] or the measurement of local changes in response to therapy.[20] Cross-sectional analysis includes the study of anatomic shape variation between normal and pathological structures,[21–23] and non-rigid alignment of a standardized anatomic atlas to individual data sets.[24,25] In many of these cases, the principal challenge of a correspondence search is that a unique solution for matching points on dissimilar objects does not necessarily exist. Thus, the problem also entails the definition of "allowable" or "probable" correspondences according to basic physical constraints or application-specific knowledge.

This section presents a new comprehensive method to establish correspondences between morphologically different 3D objects. Similarities between shape features guide the correspondence and constraints imposed on a global optimization prevent physically unrealizable matches. The correspondence mapping itself is performed in a geometry- and orientation-independent parameter space in order to establish a continuous mapping between objects. We will present this in three parts: the first describes the concept of shape description by parameterization via SH, the second contains the definition of "Parameter Space Warping" and the third shows example applications of this technique.

2.2. *Shape parameterization*

The representation of 3D shapes is a well-recognized challenge in computer vision. Parametric models have received particular attention, because they provide concise and robust representations of shape properties well suited to object comparison.

A rough classification of object descriptors distinguishes between *local* and *global* representations, with respect to composition, and between *implicit* and *explicit* representations with respect to formulation.

Local representations describe the object by a connected series of primitives, such as planar polygons, edges etc., whereas global representations use a single (parametric) function in a global object-oriented coordinate system.

An implicit representation postulates a given relationship, which all points belonging to the object must satisfy. The equation $(x/a)^2 + (y/b)^2 + (z/c)^2 = 1$, for example, is an *implicit* representation of an ellipsoid in Cartesian coordinates

$[x, y, z]$, whereas $[x, y, z]^T = [a\cos(\theta)\cos(\phi), b\cos(\theta)\sin(\phi), c\sin(\theta)]^T$ is an *explicit* representation of the same ellipsoid, using spherical coordinates (θ, ϕ).

An essential advantage of using a parametric description for shape comparison is that it reduces the dimensionality of the domain, i.e. for a 3D object we step from $\Re^3 \to \Re^2$, i.e., a 3D object is represented by a 2D parameter space, a 2D curve is represented by a 1-dimensional (1D) parameter space, etc. Hence we invest extra work in creating an explicit description of the object, but then have the benefit of a much smaller domain, which gives us a reduced search/optimization space.

Obviously, the *mapping* of an object geometry onto such a reduced parameter space critically determines the quality of its description. In other words, how we map the geometry will determine how useful the parameter space will be in representing the object's topology. If we choose our parameter space unwisely, the mapping can become very hard to do and we may have to deal with complicated constraints and boundary conditions (BC). It is advantageous, therefore, to choose a map that is topologically equivalent, i.e., a space that complies with the intrinsic BC of the data. BC not addressed by the geometry of the parameter space will have to be formulated explicitly as external constraints, eventually complicating the description and especially the mutual comparison of parameterized objects. For example, the mapping of closed contours and surfaces is useful in a parameter space with an inherent periodicity, like a circle or sphere, respectively, compared to a planar map like the unit square.

2.2.1. *Shape parameterization: spherical mapping*

Spherical harmonics (SH) are discussed here as basis functions to represent the object surface in parameter space. Because different regions with varying size and shape are to be compared in the correspondence search, it is important that the object representation fulfill the requirement of *equidistant sampling*, which means that sampled points should all be approximately the same distance from each other. In terms of a 3D object mapping, this is formulated as an *equal-area (homologic) mapping* of the object onto the unit sphere. Such a mapping demands that the size of an objects representation on the sphere is proportional to its original, size. For example, a region occupying 20% of the object surface should also cover 20% of the map area, in this case 20% of the unit sphere.

Generally, a homologic spherical mapping will require an optimization of some kind. For the examples shown here, a method developed by Brechbühler *et al.*[26] was used. This method starts from an initial, continuous mapping onto the unit sphere, obtained from solving the Dirichlet problem $\nabla^2\theta = 0$ and $\nabla^2\phi = 0$ for latitude θ (Fig. 1) and longitude ϕ, respectively, with boundary conditions $\theta_{NP} = 0, \theta_{SP} = \pi$, where NP and SP are two arbitrary object points that are mapped to the sphere's north pole and south pole, respectively. Finally, longitude must obey an additional periodicity constraint $\phi(t + 2\pi) = \phi(t)$, which is the equivalent of the geographical "dateline".
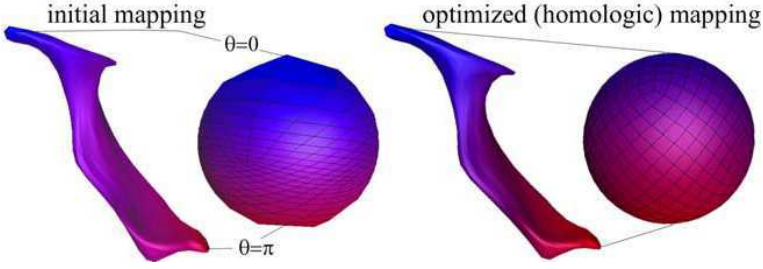
Fig. 1. Example of the spherical mapping in two stages: first, an initial mapping is obtained by solving the Dirichlet problem, with two points selected as north an south-pole as the boundary conditions (left). A second step then runs an optimization that seeks a homologic mapping, i.e. distributes all sample points equally over the sphere without overlaps or excessive distortions (right). The latitude ranges from 0 to π. The object shape is one lateral ventricle obtained from brain MRI. NOTE: A color version of this book chapter is available from the authors.

Thus, there are two sets of partial differential equations to solve, one for latitude θ, and another for longitude ϕ. This initial spherical projection is then subject to an additional optimization to create an isotropic distribution of the object surface area onto the sphere.[26]

2.2.2. *Parameterization: spherical harmonics*

Once a spherical map is obtained, a parameterization of the object surface $\mathbf{x}$ is given as a SH series:

$$\mathbf{x}(\theta, \phi) = [x, y, z]^T = \sum_{l=0}^{H} \sum_{m=-l}^{l} \mathbf{a}_{l,m} Y_l^m(\theta, \phi) \tag{1}$$

$$\mathbf{a}_{l,m} \approx \frac{4\pi}{N} \sum_{i=1}^{N} \mathbf{x}_i Y_l^m(\theta_i, \phi_i). \tag{2}$$

The SH function Y_l^m itself is the solution of the Laplace equation in spherical coordinate system (θ, ϕ), i.e. Y_l^m forms a set of basis functions over the surface of the unit sphere. The integers l and m refer to the angular and azimuthal quantum numbers, respectively, with $l > 0$ and $-l \leq m \leq l$, often also referred to as *order l* and *level m*, respectively.

$$Y_l^m(\theta, \phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos\theta) e^{im\phi}. \tag{3}$$

P_l^m are the associated Legendre polynomials, defined by the differential equation:

$$P_l^m(x) = \frac{(-1)^m}{2^k l!} (1 - x^2)^{m/2} \frac{d^{l+m}}{dx^{l+m}} (x^2 - 1)^l. \tag{4}$$

A recursive algorithm for their computation is given in Ref. 27. A graphical representation of Y_l^m is shown in Fig. 2, for the four lowest harmonics ($l < 4$).
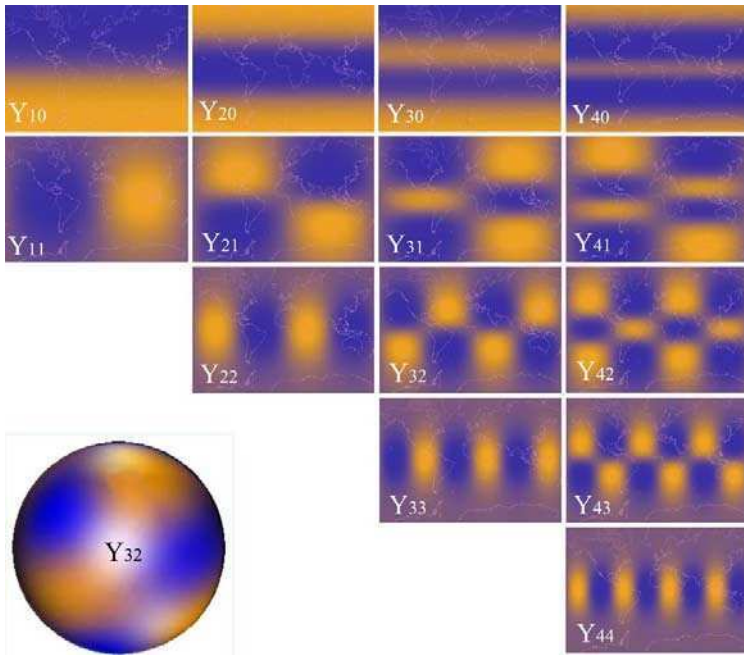
Fig. 2. Graphical representation of the four lowest spherical harmonics, with value of Y_l^m shown in bright and dark as negative and positive extrema, respectively. A world map is superimposed to convey the notion of a spherical map. NOTE: A color version of this book chapter is available from the authors.

2.2.3. *Parameterization: shape characteristics*

Once a parameterization is obtained, we can derive shape characteristics easily from the explicit description. We use these characteristics to build a criterion function that drives the correspondence search. In other words, the criterion function evaluates the quality of a correspondence mapping between two objects, based on measures that capture similarity of shape.

Differential surface features for this shape characterization are obtained directly from the analytical representation.[28] This has the distinct advantage that higher derivatives are not corrupted by noise, as direct measures from the geometry would be. If we denote the partial derivatives of the Cartesian coordinates x (Eq. (1)) with subscripts as $\frac{\partial x}{\partial \theta} \equiv x_\theta, \frac{\partial x}{\partial \phi} \equiv x_\phi, \frac{\partial^2 x}{\partial \theta \partial \phi} \equiv x_{\theta\phi}$ etc. then the surface normal, for example, is given by the cross-product:

$$\mathbf{n} = \frac{x_\theta \times x_\phi}{|x_\theta \times x_\phi|}. \tag{5}$$

The scalar coefficients of the first (E, F, G) and second (e, f, g) fundamental form are calculated as dot-products:

$$\begin{aligned} E &= x_\theta \cdot x_\theta, & F &= x_\theta \cdot x_\phi, & G &= x_\phi \cdot x_\phi, \\ e &= x_{\theta\theta} \cdot \mathbf{n}, & f &= x_{\theta\phi} \cdot \mathbf{n}, & g &= x_{\phi\phi} \cdot \mathbf{n}. \end{aligned} \tag{6}$$

Once these are obtained, the principal curvatures κ (the extrema of normal curvature),[28] are derived from the solutions λ_1, λ_2 to the quadratic equation:

$$(Fg - Gf)\lambda^2 + (Eg - Ge)\lambda + Ef - Fe = 0$$

which yields the principal curvatures:

$$\kappa_i = \frac{e + 2f\lambda_i + g\lambda_i^2}{E + 2F\lambda_i + G\lambda_i^2}, \quad i = 1 \ldots 2. \tag{7}$$

Since we are interested in shape *comparison* rather than descriptive shape analysis, we seek a curvature descriptor with a large and smooth dynamic range. Neither Gaussian nor mean curvature, however, tend to have these properties. In the spirit of polar/spherical transformations, a good and sensitive measure can be obtained by transforming the principal curvatures κ_1, κ_2 into polar space, which yields a term called the "local shape index" c.[29] Unlike the curvatures above, this measure is decoupled from object size, which is what we want for shape comparison, i.e. we seek a relative rather than an absolute curvature measure:

$$c = -\frac{2}{\pi} \cdot \tan^{-1}\left(\frac{\kappa_1 + \kappa_2}{\kappa_1 - \kappa_2}\right). \tag{8}$$

The transformation into polar space yields a size independent curvature measure, and a further normalization yields a "rank-order" measure especially suited for object comparison, i.e. we seek to match the "most convex" region of object 1 with the equivalent region of object 2, without postulating that the actual curvature values necessarily be the same.

$$\hat{c} = \frac{2c - \max(c) - \min(c)}{\max(c) - \min(c)}. \tag{9}$$

2.2.4. *Parameterization: uniform spherical sampling*

When we sample our parametric object description (Eq. 1), it is important that we do so in a uniform manner, i.e. we want sample points equidistantly distributed. This may not be an issue for rendering purposes, but becomes relevant for automated comparison, since we do not know *a priori* which region of an object matches another.

In Fig. 3 we show two examples of a spherical sampling in polar space. If we were to sample the spherical coordinate latitude in constant steps, we would end up with an overrepresentation near the poles.

This can be avoided by using a set of precomputed spherical meshes, with resolutions matching the object size (Fig. 4).

2.2.5. *Parameterization: shape scale space*

Having a frequency-based parametric description also provides us with an elegant way of creating a scale space of shape, i.e. a simple, yet robust way of controlling the level of shape detail. This we will use as part of the optimization. By increasing
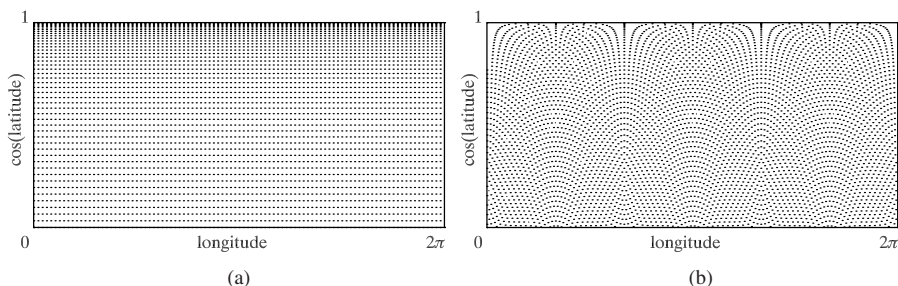
Fig. 3. A uniform sampling of longitude and latitude (a) results in an over-representation of the area near the poles of the sphere. An equidistant sampling is shown in (b). The effect becomes apparent with appropriate scaling of the latitude axis, i.e. plotting cos(latitude).
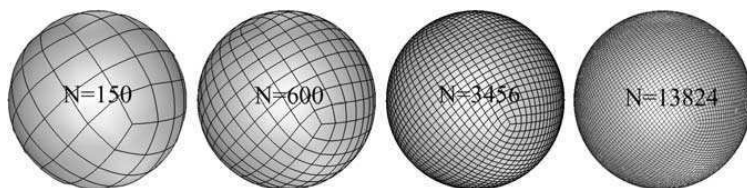
Fig. 4. Example of sphere tessellations providing homogeneous sampling of the spherical parameter space. Such sphere sets can be precomputed and stored to provide sampling of object geometry at multiple resolution levels.
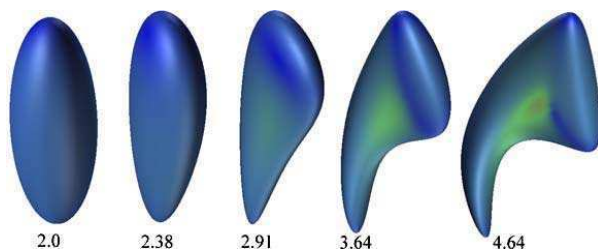
Fig. 5. Example of shape evolution by modulating spherical harmonic frequency. Using the prinicipal two harmonics (2.0) results in an ellipsoid, representing the principal axes of the object. Increasing the number of harmonics provides an elegant way to control the level of shape detail. A continuous and smooth shape space is obtained from a geometric series of fractional harmonics, as indicated by the numbers beneath each rendering. NOTE: A color version of this book chapter is available from the authors.

the level H of the highest harmonic used in the object description, we gradually increase the global level of shape detail. Figure 5 shows the shape development with increasing harmonics H, on a rendering of a lateral ventricle. This continuous scale space provides an ideal platform for a multi-resolution approach to the optimization involved in object matching, i.e. the surface over multiple harmonics H shows a coherent landscape suitable for gradient-descent optimization. The analytical derivation of these differential values assures a high degree of accuracy and robustness. Unlike derivatives computed directly from discrete sample points, they are insensitive to noise and sampling density.

A continuous scale-space is obtained from the discrete frequency descriptions above by linear or exponential interpolation, i.e., by defining "rational harmonics". A "2.38-harmonic object", for example, is created by adding 38% of the 3rd-harmonic to the 2-harmonic ($H = 2$) representation, i.e., $\mathbf{x}_{H=2.3} = (1-0.38)\mathbf{x}_{H=2} + 0.38 \cdot \mathbf{x}_{H=3}$. Note that the changes in geometry by incrementing H grow smaller with larger H, i.e. the change of surface detail encoded by the harmonics is a geometric rather than an algebraic series, as shown in Fig. 5.

The optimization in the Parameter Space Warping (PSW) concept thus executes a form of shape tracking. We begin the optimization with a low-harmonic description, where correspondences are trivial (e.g. comparing two ellipsoids). We then gradually increase the level of detail and track correspondences through this scale space of shape. This provides a powerful and robust mechanism for the automated optimization involved.

2.3. *Correspondence: parameter space warping*

The final goal of a correspondence matching between two parameterized objects, in this context, is to establish a mapping between their respective parameter spaces ψ_1 and ψ_2, rather than between the object geometries $\mathbf{x}_1$, $\mathbf{x}_2$ themselves. In other words, by comparing parametric object representations, we create a *de facto* separation between object geometry and topology, which, as we will see, is an ideal premise to solve the correspondence problem. The desired mapping consists of a transform $\psi_1 \rightarrow \tilde{\psi}_1$, and the quality of the mapping is summarized in a criterion/cost function $J[\mathbf{r_1}(\tilde{\psi}_1), \mathbf{r}_2(\psi_2)]$. The transform $\psi_1 \rightarrow \tilde{\psi}_1$ that minimizes J thus provides a correspondence mapping which has $\mathbf{x}_1$, sampled with parametric coordinates, $\tilde{\psi}_1$, and $\mathbf{x}_2$, sampled with its original parametric coordinates ψ_2. You may visualize the transform $\psi_1 \rightarrow \tilde{\psi}_1$ as a motion of the points $\mathbf{x}_1$ constrained to the surface of object 1. Because this transform is non-rigid and applied to parameter space, we denote this process as "*parameter space warping*" (PSW).[30] This is illustrated in Fig. 6.

Imagine two corresponding subregions of the objects being compared. Those regions will be mapped onto a contiguous area in parameter space. Because of the morphological difference between the two substructures, their parametric representation will differ not only in position, but also in size and shape. The objective of the PSW is to align these corresponding regions by virtue of a spatial transformation applied to the parameter space. Formally, the PSW is modeled as:

$$\tilde{\psi} = \psi + \triangle\psi = \psi + \sum_{k=1}^{K} \mathbf{b}_k \cdot B_k(\psi). \tag{10}$$

In choosing an appropriate basis function $B_k(\psi)$ for this modulation, it is natural to follow the topology of the parameter space itself, i.e. since our parameter space is the unit sphere $\psi = (\theta, \phi)$, we choose again spherical harmonics as basis function for the PSW: $B_k = Y_l^m$. Thus we take advantage of the intrinsic constraints of the spherical basis functions.
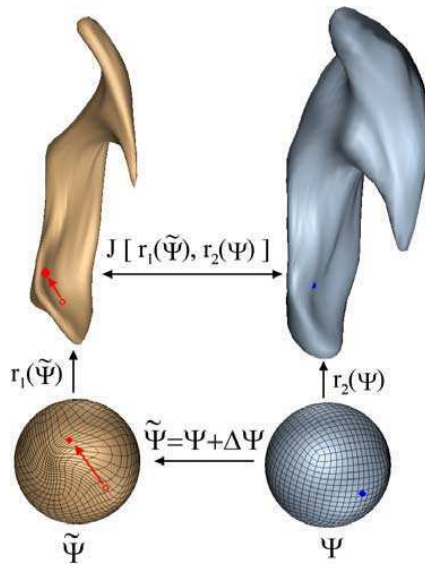
Fig. 6. Parameter Space Warping (PSW) concept: The spherical map of the left ventricle is deformed by the transform $\Delta\psi$ so that the cost function (J) measuring the structural differences between the objects is minimized. The function $r_i(\psi)$ defines the geometry, the function $\Delta\psi$ maps one parameter space onto the other and the cost function J describes the shape similarity between the geometry. After the transform $\Delta\psi$, corresponding points of the two objects should also be in the same position on the sphere. For details see also Ref. 30.

We do not modulate longitude and latitude directly, however, to avoid the singularities at the poles of the spherical coordinate system, where longitude is not defined (recall that when standing at the North Pole, every direction is south and no direction is east). Instead we introduce a *virtual* degree of freedom and formulate the PSW in Cartesian coordinates:

$$\tilde{\boldsymbol{\rho}} = \boldsymbol{\rho} + \triangle\boldsymbol{\rho} = \begin{bmatrix} x \\ y \\ z \end{bmatrix} + \sum_{l=1}^{k} \sum_{m=-l}^{l} \begin{bmatrix} b_{l,m}^x \\ b_{l,m}^y \\ b_{l,m}^z \end{bmatrix} Y_l^m \tag{11}$$

with $b_{l,m}^x$, $b_{l,m}^y$, $b_{l,m}^z$ as the coefficients. Thus, the shifting of sphere surface points is now expressed separately for each Cartesian coordinate. Directional preferences are removed by the additional degree of freedom (DOF) introduced. We compensate this virtual DOF with an additional constraint to keep the points on the sphere, i.e., by forcing $|\tilde{\rho}| = 1$. This can be implemented explicitly within the optimization, or implicitly by re-normalization, i.e. by projecting the points back onto the unit sphere:

$$\tilde{\boldsymbol{\rho}} = \frac{\boldsymbol{\rho} + \triangle\boldsymbol{\rho}}{|\boldsymbol{\rho} + \triangle\boldsymbol{\rho}|}. \tag{12}$$

An example of two different objects, their respective spherical maps and corresponding regions is shown in Fig. 7.
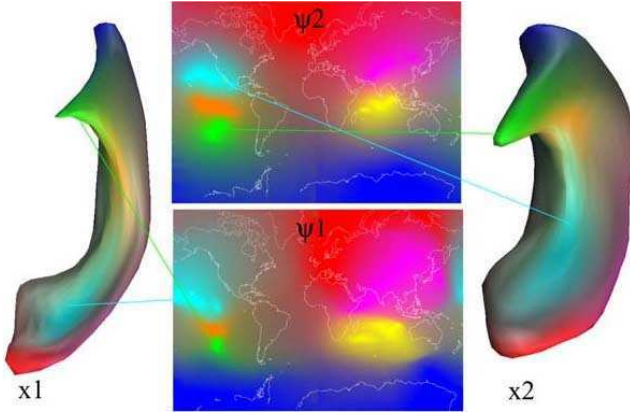
Fig. 7. Example of object correspondence between two different shapes x1 and x2 (two lateral ventricles obtained from brain MRI), and their corresponding spherical maps $\Psi1$ and $\Psi2$ (with world contours superimposed to convey the notion of a spherical map). Note that corresponding regions will differ on the map not only in position, but also in size and shape, which is why a non-rigid parameter space warp $\Delta\Psi$ must be applied to the map to establish congruence. NOTE: A color version of this book chapter is available from the authors.

2.3.1. *Structural criterion function*

The similarity criterion function J should be a cost function that captures our understanding of a good correspondence. Since our emphasis is on morphological equivalence, we choose a J that models instances of structural similarity, but others could also be used, such as instances of deformation energy,[31] purely statistical measures like mutual information,[32] or the distance to explicit landmarks.[23,33]

In following a structural paradigm, we design our cost function as a super-position of three components, representing Euclidian point-to-point distance J_D, differences in surface normal direction J_N, and curvature J_C, summed over the object surface:

$$J_D = \frac{1}{Nd_0^2} \sum_N |\mathbf{x_1} - \mathbf{x_2}|^2$$

$$J_N = \frac{1}{N\pi} \sum_N \cos^{-1}(\mathbf{n}_1 \cdot \mathbf{n}_2) \qquad (13)$$

$$J_C = \frac{1}{4N} \sum_N |\hat{c}_1 - \hat{c}_2|^2.$$

The indices 1 and 2 denote the two objects being compared. N is the number of sample points (Fig. 4) used. Object geometry is represented by the vectors $\mathbf{x}_1$, $\mathbf{x}_2$ from the SH series in Eq. (1), the surface normal vectors $\mathbf{n}_1, \mathbf{n}_2$ from Eq. (5), and the normalized curvature values $\hat{c}_1, \hat{c}_2$ from Eq. (9). Whereas the components J_N and J_C range from [0,1], the distance term J_D, is normalized to [0,1] by an external reference length d_0, for example: $d_0 = \max_N(|\mathbf{x_1} - \mathbf{x_2}|)$. This makes all components of J bounded within the interval [0,1].

Finally, we combine the three components with an empirical weight vector $\mathbf{w} = [w_D, w_N, w_C]^T$, with $\sum w_i = 1$. The final scalar cost is then given by: $J = \mathbf{w} \cdot [J_D, J_N, J_C]^T$.

Different applications are likely to use different weights and thus bias the criterion function toward features that demonstrate the greatest stability for the given morphology. For example, shape comparisons among objects of greatly varying size would emphasize intrinsic measures like curvature and normal direction rather than Euclidian distance.

2.4. *Application: anatomic labeling by atlas warping*

This is an example application of the PSW concept above to the problem of automated atlas-based anatomical parcellation, or anatomical labeling. We take a pre-labeled atlas of subject A and warp it into congruence with subject B, thereby obtaining an estimate of the location, shape and size of the anatomical structures of subject B. In other words, we transfer the anatomical parcellation from A to B, by virtue of extrapolating correspondences between structures delineated in both subjects. This 3D warping is driven by correspondences derived between individual structures, as outlined above. For more details see Ref. 34.

The strategy to deform an atlas image to match a sample geometry also provides estimates for structures that cannot be segmented directly at all, due to lack of image resolution or contrast. The atlas warp applies knowledge we have regarding the shape differences between the two subjects, and extrapolates it in a structurally feasible manner to parts of the anatomy we have no direct segmentation of. In other words, individual segmentation results for two different structures may be useful in finding a third, where the atlas serves as a model to incorporate *a priori* knowledge about the "relational" arrangement of anatomical structures.

Figure 8 shows an example of warping a pre-labeled atlas with major lobe parcellation to a different subject anatomy. This was accomplished by finding explicit
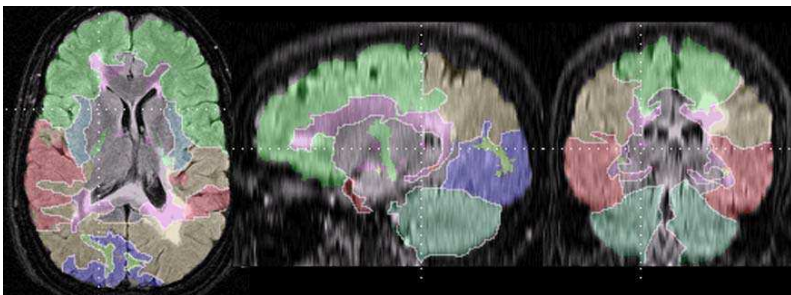
Fig. 8. Example application of correspondence driven atlas warping for automated anatomical parcellation. The lobar anatomy of a subject is automatically delineated into frontal, temporal, parietal, occipital and insular lobe and the cerebellum, by virtue of warping a pre-labeled atlas of another subject (not shown) into congruency.[34] NOTE: A color version of this book chapter is available from the authors.

correspondences between three sub-structures: the brain surface, the lateral ventri-
cles (both left and right) and the inter-hemispheral midline surface, using the PSW
concept.[30] These correspondences were then used to drive an image warp based on
multi-quadrics, aligning the atlas with the subject.[34]

3. Skeletons and Medial Representations

Medial models have been successfully used in medical image analysis in a number of
contexts, see Refs. 12, 35 for some recent applications. Applying these methods in
3D presents an additional challenge because only a small class of computationally
reliable algorithms exists. Nevertheless, Styner and Gerig[36] have recently developed
a framework where spherical harmonics along with a coarse-scale sampled medial
description are used to represent the shape of anatomical structures of the brain in
particular the hippocampus and the lateral ventricles. A key idea is that the medial
primitive can be used to describe the most prominent medial sheet and to provide
an intrinsic frame of reference by which different data sets can be registered.

In this section, we present our own methodology based on the 3D skeleton to
study the effect of aging on the hippocampus.[37] After reviewing the definition and
some of the properties of the skeletons, we briefly present a method to compute
them robustly and show through an example how they can be used to perform
statistical analyses describing the shape changes of the hippocampus.

3.1. *Definition and properties*

There are several definitions of skeletons in the literature, some of which are equiv-
alent. We focus on the most common, popularized by Blum,[38] which is based on
the notion of grassfire.

Definition 1. The skeleton (axis in 2D, surface in 3D) of X, Sk(X), is the set of
points of X simultaneously reached by grassfires initiated from at least two different
points of the boundary of X. This is equivalent to the set of points of X for which
there are at least two closest points of the boundary of X, in the sense of Euclidean
distance.

Definition 2. The skeleton transform of an object X, ST(X), is the skeleton Sk(X)
together with the distance function, defining for each point in Sk(X) the shortest
distance to the boundary of X.

An example 2D object and its skeleton are shown in Fig. 9. We now review some
properties of the skeleton.

Property 1. *Invariance under translation and rotation*: Sk(X) is invariant under
translation and rotation. Given a translation or a rotation g Sk(g(X)) = g(Sk(X)).

This is a crucial property, as most shape analysis applications require at least
invariance to translations and rotations.
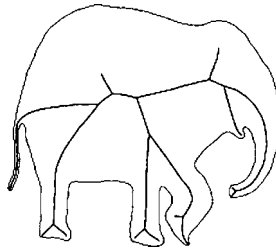
Fig. 9.   The 2D skeleton of an elephant shape.

Property 2. *Reversibility*: given Sk(X) and the minimum distance r at each point of Sk(X), it is possible to recover the object. X is the union of the balls centered at Sk(X) of radius r.

In other words, the set Sk(X) together with its associated distance function is equivalent to the set X. The skeleton and its radius information is thus a complete representation of the object.

Property 3. *Thickness*: The interior of Sk(X) is empty.

The skeleton is at least one dimension lower than the object. It consists of curves and points in 2D and of surfaces, curves and points in 3D. Thus it provides a compact representation of the object.

Property 4. *Homotopy*: It is important that the "essential structure" of the object be preserved in the skeleton. For example in 2D, the skeleton and the object should have the same number of connected components and the same number of holes. This idea of "topological equivalence" is called *Homotopy*. Intuitively, it would seem natural that the skeleton and the object are homotopy equivalent; however, proving this is a rather difficult task. Recently, Lieutier[39] claimed to have a proof that any open bounded subset of R^n and its skeleton are of the same homotopy type. It is possible, as we will later on to ensure digital homotopy equivalence between a digital 3D object and its digital skeleton.

3.2. *Computational technique*

Computing skeletons is a challenging problem, as there is no simple technique that can detect the local maxima of the Euclidean distance transform and ensure all the mathematical properties of the skeleton (reversibility, compactness and homotopy equivalence). A number of algorithms have been developed in order to achieve this task, but only a few have been successful.[40-42]

In this section, we review the work of Siddiqi *et al.*,[43] which we believe is robust and accurate, and can thus be used as the basis for our hippocampus study.

3.2.1. *Hamilton Jacobi skeletons*

We now review the algorithm on which the subsequent analysis of hippocampus shape is based. For details we refer the reader to Ref. 43. Consider the grassfire flow

$$\frac{\partial S}{\partial t} = N \tag{14}$$

acting on a closed 3D surface S, such that each point on its boundary is moving with unit speed in the direction of the inward normal N. The shocks of this equation correspond to the location of the skeleton. They also correspond to the local maxima in the Euclidean distance function of the initial surface.

Let D be the Euclidean distance function to the initial surface S_0.[44] The discrimination of medial from non-medial surface points can be approached by computing the average outward flux of the vector field ∇D at a point q. This quantity is given by

$$AOF(q) = \frac{\displaystyle\int_{\partial R} \langle \nabla D, N_o \rangle \, dS}{area(\partial R)} \tag{15}$$

where dS is a surface area element of the bounding surface ∂R of a volume R and N_o is the outward normal at each point on the surface. It can be shown that as the volume shrinks to a point not on the medial surface, the average outward flux approaches zero. In contrast, when the volume over which it is computed shrinks to a medial surface point, the average outward flux approaches a strictly negative number.[45] Thus, it is an effective way for distinguishing between these two cases. This calculation is used to guide a thinning process in a cubic lattice, while taking care to preserve the object's homotopy type.

3.2.2. *Preserving homotopy*

A point is *simple* if its removal does not disconnect the object, create a hole, or create a cavity. We are interested in detecting those points, as their removal will not change the homotopy type of the object. Malandain *et al.*[40] have introduced a topological classification of a 3D point in a cubic lattice by computing two numbers:

- C: the number of 26-connected components of the object 26-adjacent to x.
- C*: the number of 6-connected components of the background 6-adjacent to x.

Further, they have shown that if C $= 1$ and C* $= 1$, the point is simple, hence removable. Our basic strategy now is to guide the thinning of the object by the average outward flux measure computed over a very small neighborhood. Points with the most negative average outward flux are the strongest medial surface points. The process is stopped when all surviving points are not simple, or have an average outward flux below some chosen (negative) value, or both. Unfortunately, the result is not guaranteed to be a thin set, i.e., one without an interior. This last constraint can be satisfied by defining an appropriate notion of an endpoint in a
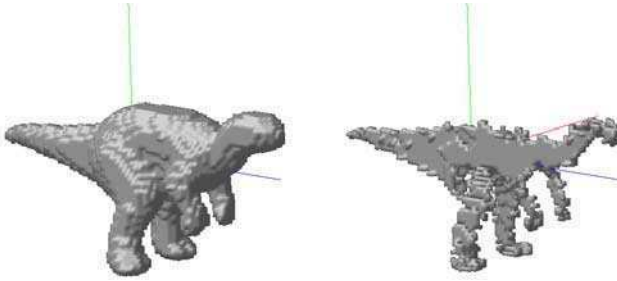
Fig. 10.   A 3D model of a dinosaur and its skeleton.

cubic lattice. In R^3, if there exists a plane that passes through a point x such that the intersection of the plane with the object includes an open curve which ends at x, then x is an end point of a 3D curve, or is on the rim or corner of a 3D surface. This criterion can be discretized easily to 26-connected digital objects by examining 9 digital planes in the 26-neighborhood of x.[46] The thinning process proceeds as before, but the threshold criterion for removal is applied only to endpoints. An example 3D object and its skeleton are shown in Fig. 10.

3.3. *Hippocampus and aging: an example study*

3.3.1. *Hippocampus data*

In a previously published analysis,[47] the left and right hippocampus were manually segmented from T_1-weighted MR images from 80 normal healthy subjects. These subjects included 39 healthy men and 41 healthy women in the age range of 18 to 42 years (mean age $25.4 + / - 5.6$ years). The MRI data for each subject was first corrected for image intensity non-uniformity[48] and linearly registered to a standard brain-based coordinate system known as stereotaxic space.[49]

Quantitative analysis of the original label volumes derived from the MR images revealed an overall hemispheric difference with the right hippocampus being bigger than the left (4300 vs. $4072 \, \text{mm}^3$, $p < 0.001$). Furthermore, there was a correlation with age in the group of men for both left and right hippocampus (Spearman's
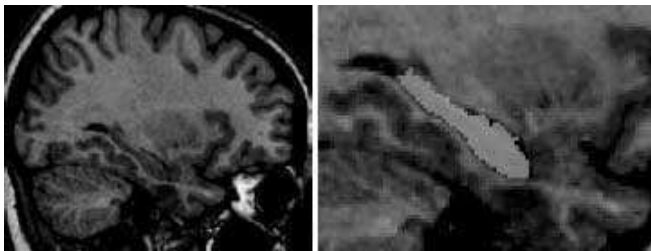
Fig. 11.   Left: a sagittal view of the hippocampus within the brain. Right: 3D model of the corresponding hippocampus.

$r = -0.44$, p < 0.005) that was not apparent in the group of women ($r = 0.01$, p > 0.20). However, it was not possible to determine whether the observed differences were due only to a volume reduction or were associated with the overall shape of the hippocampus as well. The medial surface algorithm presented here extends the possibilities for analysis, allowing for the investigation of shape differences in the hippocampus between men and women. We chose to focus this analysis on the radius function of the medial surface of the hippocampus, which indicates the local width of this structure at different locations in the medial temporal lobe.

3.3.2. *Methodology*

First, we observed that across the entire data set of hippocampuses, only one main medial sheet in the skeleton was necessary to accurately represent the object. In fact, reconstructing the object with the medial sheet alone accounts for 99% of the overall volume of the original object. We thus chose to focus only on this sheet. Our approach is to use the medial sheet with its associated distance (also called radius) function as a measure of local thickness of the hippocampus. Hence, by comparing these sheets with respect to gender or age, one can get important information about possible shape changes in the hippocampus.

In order to perform a meaningful statistical analysis, a correspondence between the medial sheet of the different individual hippocampuses must be established. One approach is to associate each medial sheet with a common coordinate frame by fitting a coarsely sampled medial model to it, as suggested by Styner and Gerig.[36] Whereas this leads to a one to one correspondence between the sheets, the deformation each medial manifold has to undergo to fit the common frame may itself introduce shape changes, particularly when there is a variation in medial manifold shapes across the data set. Our approach is to first flatten the 3D sheet onto a 2D plane and then try to establish a correspondence through linear alignment on the 2D plane.

The flattening is inspired by early work on the map making problem.[50,51] Let there be N points in l-D that are to be flattened so that they lie in m-D, with $l > m$. Let dd_{ij} be the geodesic distance between points $x_i = (x_{i1}, x_{i2}, \ldots, x_{il})$ and x_j in l-D and d_{ij} be the geodesic distance between the corresponding points $y_i = (y_{i1}, y_{i2}, \ldots, y_{im})$ and y_j in m-D. Sammon[50] defines an error function E_r based on dd_{ij} and d_{ij}, which represents how well the m-space configuration fits the points in l-space:

$$E_r = \frac{1}{c} \sum_i \sum_j \frac{(dd_{ij} - d_{ij})^2}{dd_{ij}} \tag{16}$$

where $c = \sum_i \sum_j dd_{ij}$. One is essentially left with the task of moving the points y_i in order to minimize E_r, which can be solved using a standard gradient descent strategy.

The next challenge of the registration process and probably the most difficult is to establish a correspondence between each individual flattened sheet. When a

(a) Hippocampus   (b) Medial Manifold

(c) Flattened Surface   (d) Radius Map
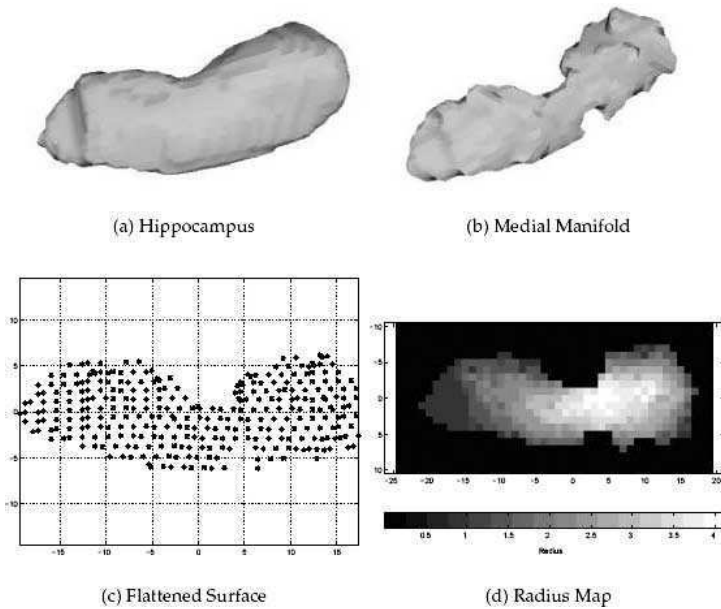
Fig. 12. The different stages of our radius map generation. The medial sheet (b) of the hippocampus (a) is first extracted then flattened (c), realigned and the radius function mapped onto it (d).

one to one mapping is not required, one can align the 2D shapes by translating and rotating each individual in a common coordinate frame using their centroid and principal directions. This method is not optimal and can lead to serious artifacts if the 2D sheets do not have similar shape. Nevertheless, our data set is very consistent and there are little differences in shape between the different 2D manifolds, which is why we feel principal directions are sufficient to register the 2D flattened manifolds.

Figure 12 presents the output of the different stages of the method: medial surface computation, flattening of the medial sheet and mapping of the radius function on the realigned flattened medial manifold.

3.3.3. *Statistics*

The flattening of the medial surfaces derived from the original hippocampus labels resulted in two dimensional maps, with each point on the map referring to the medial surface radius as a height function (Fig. 12(d)). These maps could then be employed for further statistical analysis keeping in mind that the signal-intensity information contained in the 2D maps at each point refers to the local width of the original hippocampus. A pixel-based regression was used to localize statistically significant shape differences that are correlated with age. This regression was estimated at each pixel location in the 2D map. In order to detect regions that are significantly different (i.e., where the differences are above chance), it was necessary to
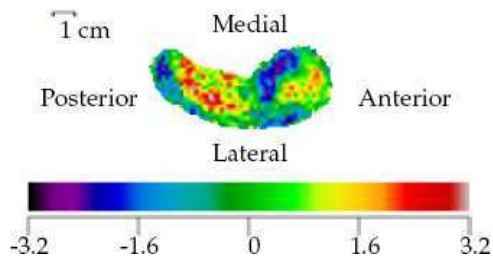
Fig. 13. An example regression map, the bar represent the t-value, $p < 0.05$ significance is reached for $|t| > 3.2$. NOTE: A color version of this book chapter is available from the authors.
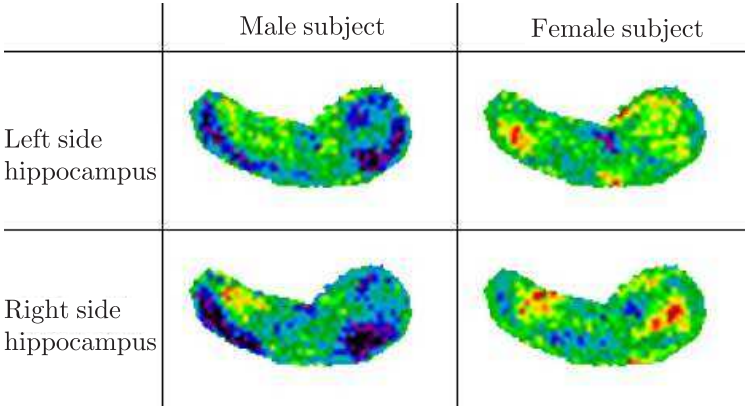
Fig. 14. Linear regressions to examine the effects of age on the radius maps. NOTE: A color version of this book chapter is available from the authors.

correct for multiple comparisons and local spatial correlation. We used the random field theory of Worsley *et al.*[52] to determine the t-value threshold that indicates statistically significant differences, given a selected level of alpha-error. Given the number of subjects, data resolution, pixel size and area of search region, the critical t-value for a 2-tailed alpha significance level of $p < 0.05$ after correcting for multiple comparisons is $|t| > 3.2$. Figure 13 shows an example regression map.

3.3.4. *Results and discussion*

Figure 14 illustrates the effects of age on the hippocampus radius maps in both the left and the right hemispheres. Within the group of women, a positive regression with age occurs in the middle of the structure, rather than towards the rim as was observed with the support maps. This positive regression appears as two clusters in the areas adjacent to the head and tail of the hippocampus, almost exclusively in the right hemisphere. In the left hemisphere, a small region of positive regression appears in the lateral part of the hippocampus tail. Due to the conservative t-value threshold, these results also fail to reach significance. Within the group of men,

an almost opposite picture emerges. Here, the results show a consistent pattern of negative regression with age in the right hemisphere, towards the lateral rim of the structure in the area adjacent to the head of the hippocampus. Also, a streak of negative regression can be observed in the posterior portion along the lateral rim, in the area adjacent to the tail of the hippocampus. A similar pattern, although less strong, can be observed in the left hemisphere.

Using quantified hippocampus volumes from a recently developed segmentation protocol,[53] we observed a gender-specific age-related volume decline in a healthy population of men and women in early adulthood.[47] The current study implementing the method of hippocampus shape analysis based on medial surfaces is able to complement and extend the previously observed findings. First, a decrease in radius in the group of men could be observed, which was most prominent in the lateral rim of the hippocampus structure. This finding provides evidence for the previously stated hypothesis that the observed volume decline of the hippocampus with age in the group of men is due to shrinkage of the hippocampus in areas adjacent to the ventricles.

We suggest that this new methodology complements current approaches in structural MRI research, allowing investigations of thus far neglected aspects. It is a prerequisite of most investigations in MRI to first try to minimize shape differences between target regions via linear or nonlinear transformations to maximize comparability between them. By doing so, the researcher is then able to compare quantitative aspects of the target, like thickness of the cortical mantle, gray-matter density of specified regions of the brain, etc. However, pre-existing quantitative differences between subjects in target regions, such as the radius of the hippocampus as described here, will in most cases not be detected using these standard approaches. The results of the current study suggest that it is worthwhile to investigate these aspects, since shape differences might occur before differences in volume start to appear.

4. Shape Analysis Over Time in Multiple Sclerosis Lesions

Multiple Sclerosis (MS) is a demyelinating disease affecting principally young adults' brain and spinal cord, and MRI is the most reliable imaging technique for evaluation and assessment of MS lesions. MS is a dynamic disease, presenting both active and chronic lesions. Disease activity monitored by MRI was reported to be 5–10 times higher than the clinical neurological changes,[54–57] with MRI evidence suggesting that disease activity can be identified in as many as 30% of patients who appear to be in clinical remission.[58]

The geometry and pattern of evolution of individual MS lesions reflect the complex pathological chain of events that often culminate in the structural damage of axons and their myelin sheath. MS lesions present a considerable variability in shape and size between patients and even for the same patient. Active lesions are prone to changes in size and shape over time. While the volumetric changes of MS lesions over time have been documented to some extent,[59–61] almost no attention has been devoted to understanding the nature of changes in MS lesion shapes.

Analysis of temporal changes in MS patients by brain MRI includes counting the number of lesions, and assessment of lesion area and/or volume. An assessment of disease activity can be obtained from measurements of the number and volume of new or active lesions, and changes in older lesions.[62–64] Shape analysis of MS lesion and their changes over time could lead to better understanding of the pathogenesis of MS. However, quantitative characterization of MS lesion geometry is a complicated task. Although most of MS lesions present a rather spherical or ellipsoidal shape, complex shapes can also be found.[65]

In this section, first we present our own work on the approximation of MS lesions' 3D geometry using spherical harmonics, based on the information taken from the contour points of the lesions, which were segmented from MR images.[66] Second, we present a straight application of our method for analyzing the changes in MS lesions over time by quantitatively characterizing the lesion's shape and depicting patterns of shape evolution in individual lesions.[67]

4.1. *Characterization of the individual MS lesions' 3D geometry using SH*

The characterization of the individual MS lesions' 3D geometry is based on the use of sets of MS lesion contours segmented from MR images. In order to obtain the same number of points in all traced contours belonging to each individual lesion, the contour lines are interpolated and resampled using the Akima interpolation method.[68] Later, the points of each contour are arranged in such a way that the first point of each contour is the most extreme point on the right-hand side. In order to have ϕ, the azimuthal coordinate starting in 0 and ending at 2π for all the contours, the rest of the points are then registered in a counterclockwise direction. Thus, at the end of this stage each traced lesion is represented by a set of digitized contours with a predefined standard structure. The application and optimization procedures of the interpolation and arrangement are explained in more detail in Ref. 66.

When MS patients are studied using axial MR images using a 3 mm slice thickness with no gap, lesions commonly extend only into 1–3 consecutive slices; therefore an additional preparatory stage is needed in order to constrain the solution along the direction perpendicular to the axial plane. In this stage, we estimate a sagittal and a coronal contour and added them to the available axial contours to obtain the SH approximation of the lesions' 3D surface and its corresponding coefficients. The sagittal and coronal contours are obtained by applying the Akima interpolation method[68] to data points (from the axial contours) located on these planes and two additional points extrapolated from the contour's center of gravity. These two points are needed to ensure that the lesion size will not exceed the distance between the last contoured image plane and the image plane that followed it. The final sets of contours obtained for each lesion are arranged according to their respective spatial orientation, yielding a 3D cluster of data points.

Using these cluster of data points expressed in spherical coordinates, we estimated the corresponding $Y_l^m(\theta, \phi)$ followed by the coefficients $\mathbf{a}_{l,m}$, which define the estimated surface $\mathbf{x}(\theta, \phi)$ of Eq. (1). Then we obtained a new set of data points that cover the entire 3D space (i.e., θ in the range $[0, \pi]$, and ϕ in the range $[0, 2\pi]$) and the corresponding $Y_l^m(\theta, \phi)$ are again estimated. Finally, $\mathbf{x}(\theta, \phi)$ is recalculated using the obtained coefficients $\mathbf{a}_{l,m}$. Consequently, a complete 3D lesion surface is obtained, with known coordinates expressed in terms of R, θ, and ϕ at each point.

We analyze the changes in MS lesions' size/volume and the changes in the individual lesion shape by using the polynomial coefficients that describe the estimated surfaces. For the analysis, a space rotation invariant set of indices "I_l"[69] is estimated using:

$$I_l = \sum_{m=-l}^{l} |\mathbf{a}_{l,m}|^2 \qquad (17)$$

where $l = 0, 1, 2, \ldots, H$. The obtained I_l indices are then normalized to the value of I_0 of the baseline position (i.e., $I_l / |\mathbf{a}_{0,0}|^2$). The index I_0 comprises only $\mathbf{a}_{0,0}$, which is the general average radius of the approximated surface. Thus, $I_0^{(1/2)}$ is proportional to the average radius of the shape and $I_0^{(3/2)}$ is proportional to the volume defined by the average radius. Thus, the variations in I_0 represent the changes in size/volume. Unlike I_0, which express global changes, the $I_{l>0}$ indices express more local changes in the 3D SH surface shape. As l increases, the I_l indices characterize more localized changes.

4.2. *Detection of shape changes over time using SH*

4.2.1. *MS imaging data*

Brain MRI images were acquired from a 38-year-old man presenting a relapsing-remitting course of MS. The subject completed a 12-month imaging study consisting in 24 scans, the first eight weeks on a weekly basis, biweekly for the next 16 weeks, and then once monthly until the completion of one year.

MR image acquisition was done on a 1.5-T Signa System (GE Medical Systems, Milwaukee, WI). The acquired sets of images included proton density (PD) and T_2-weighted images obtained using two interleaved dual-echo (Time to Echo = 30 and 80 msec) long time to repeat (3000 msec) spin-echo sequences (VEMP = variable-echo multi-planar). Contiguous 3 mm thick axial sections covered the whole brain from the foramen magnum to the higher convexity with an in-plane voxel size of 0.94 mm × 0.94 mm (24-cm field of view with a 256 × 256 acquisition matrix). Signal-to-noise ratio (SNR) in the second echo was improved using a variable receiver bandwidth technique. Standard flow compensation was achieved using first-order gradient moment nulling. Spins were saturated in an 8-mm thick slab inferior to the imaging volume in slice-selective direction. Scan duration was kept at 11 minutes and 36 seconds using the half-Fourier technique.

4.2.2. *Image registration and segmentation*

As the acquired MRI scans provided comprehensive volumetric coverage of the head, it was not necessary to achieve precise repositioning of the patient in the MR scanner at subsequent examinations. Rather, we registered the 23 subsequent sets of images to the set of images obtained at the first scan by applying the maximization of mutual information algorithm proposed by Viola and Wells.[70]

We applied an automatic multistep segmentation procedure to the registered PD and T_2-weighted images. First, a region of interest is defined for all the scans by outlining the intercranial cavity (ICC) mask, which contains the brain parenchyma and cerebrospinal fluid (CSF). Later, each pixel in the images is assigned to one of the following four categories: white matter, grey matter, CSF, or lesion. This classification is performed by combining the self-adaptive expectation-maximization algorithm described by Wells *et al.*[71] with a template driven segmentation,[72,73] where anatomic knowledge is provided by nonlinear projection of a brain model (template) onto the images of a patient's brain. This technique has shown to have high reproducibility and accuracy of lesion volume measurements with respect to the outlines provided by expert radiologists.[74]

4.2.3. *Analysis of changes in the lesion's 3D geometry*

We approximated the 3D geometry of the segmented MS lesions using SH. In order to allow a systematic analysis of the changes in the SH coefficients, each individual lesion has to be approximated using the same order (l) at all time points. We observed that the choice of l must be set according to the ratio of lesion size to image resolution in order to avoid distortions in the estimated 3D surface. Because of the changes in size, the lesions do not extend into the same number of planes over time. Therefore, l is selected to be equal to the minimal number of axial slices in which the lesion extended at all time points.

After approximating the 3D geometry of the lesions, we estimate the I_l indices (Eq. (17)) for each lesion in each time point. The mean amplitude of the oscillations of the changes in the normalized I_l indices is estimated using the mean discrete total variation[75] ($MDTV$), which is calculated by approximating the signal derivative by a finite difference over the sampling distance:

$$MDTV_l = \frac{\sum_{t=1}^{N} |I_l(t) - I_l(t+1)|}{N} \tag{18}$$

where t is the time index and $N+1$ is the total number of time points. The $MDTV$ does not vary linearly with respect to changes in lesions size/volume.

4.3. *Results*

From the sets of registered and segmented images, 10 individual lesions were identified and analyzed. For a complete reference to this study see Ref. 67. A general

Table 1. Average $MDTV$ (%) for the normalized I_l indices for all the lesions over time.

I_l	$MDTV$ (%)
I_0	24.8 ± 22.1
I_1	12.0 ± 8.2
I_2	25.1 ± 19.5
I_3	7.5 ± 3.6
I_4	3.3 ± 0.4

analysis of the obtained $MDTV$ values reveals that in four of the analyzed lesions the $MDTV$ values for I_0 were higher than for the other normalized I_l indices (i.e., $I_1 - I_3$), indicating larger changes in lesions' size/volume than in the lesions' shape. However, for the remaining six analyzed lesions, the $MDTV$ for at least one of the normalized I_l indices that represent changes in shape (i.e., $I_1 - I_4$) were higher than for I_0. Half of those six lesions had an $MDTV$ for I_2 between 1.4 and 2.6 times higher than for I_0. The fourth lesion, presented an $MDTV$ for I_1 1.8 times higher than I_0, while for the fifth lesion, the $MDTV$ for $I_1 - I_3$ were between 2.3 and 8.0 times higher than I_0. Finally for the sixth lesion, the $MDTV$ for $I_1 - I_2$ were 2.1 and 3.1 times higher than I_0, respectively. The average of the estimated $MDTV$ (in %) for the normalized I_l indices for all the lesions over time are summarized in Table 1.

In Fig. 15 we depict the 3D reconstructions of one of the lesions in all time points. The corresponding values of its shape indices are depicted in Fig. 16. This lesion appeared about six weeks after the beginning of the study. The size of the lesion increased, reaching its maximum about three weeks later. This finding is also quantitatively indicated in Fig. 16 by the large value of I_0 on day 64. Following this climax, the lesion shrank, initially at a rapid rate and then at a more moderate rate. The changes in size were associated with some changes in shape.

The 3D reconstructions at the 24 time points of an additional lesion are depicted in Fig. 17. This lesion seems to engulf two overlapping lesions, one extending in a vertically and the second diagonally. Over time, the upper part of the lesion decreases, while the lower part has almost no variation. Accordingly, it can be graphically noted in Fig. 18 that for this lesion, the $MDTV$ values were higher for I_1 and I_2 than for I_0.

4.4. Discussion

The results obtained from the analysis of the in vivo MS lesions show oscillating changes in both size/volume and shape even within a short period of time. Figures 15 and 16, provide a clear example where a lesion grows substantially within a short period of time with the volumetric increase accompanied by changes in shape. However, there are also cases (e.g., Figs. 17 and 18) in which changes
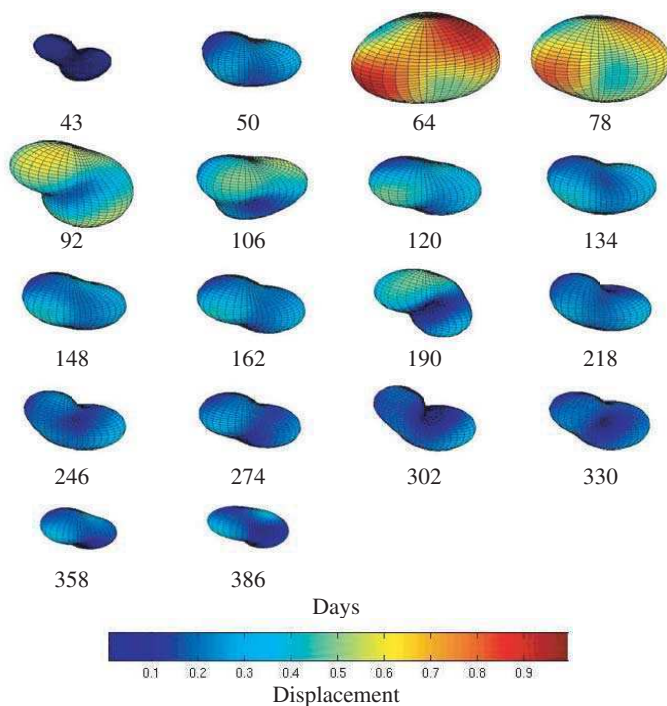
Fig. 15. 3D reconstructions of one of the analyzed lesions in all time points. This lesion appeared in the set of images at time point number 7. Note the rapid increase in its size, which peaked at day 64, and then the slow decrease in size that was associated with some changes in the lesion's shape. The bar of the 3D reconstructions represents the displacement (in voxels) for each time point from the first time point in which the lesion appeared. NOTE: A color version of this book chapter is available from the authors.
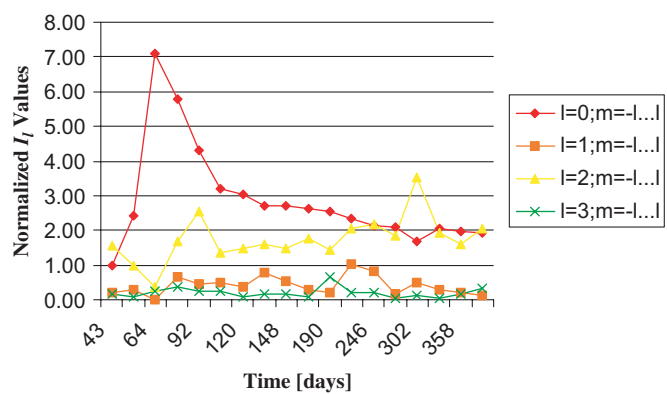
Fig. 16. The temporal patterns for the shape indices corresponding to lesion depicted in Fig. 15. The pattern is dominated by the rapid change in the lesion's size/volume indicated by the large values for I_0 at the early time points. Marked changes in I_2 are noted as well. NOTE: A color version of this book chapter is available from the authors.
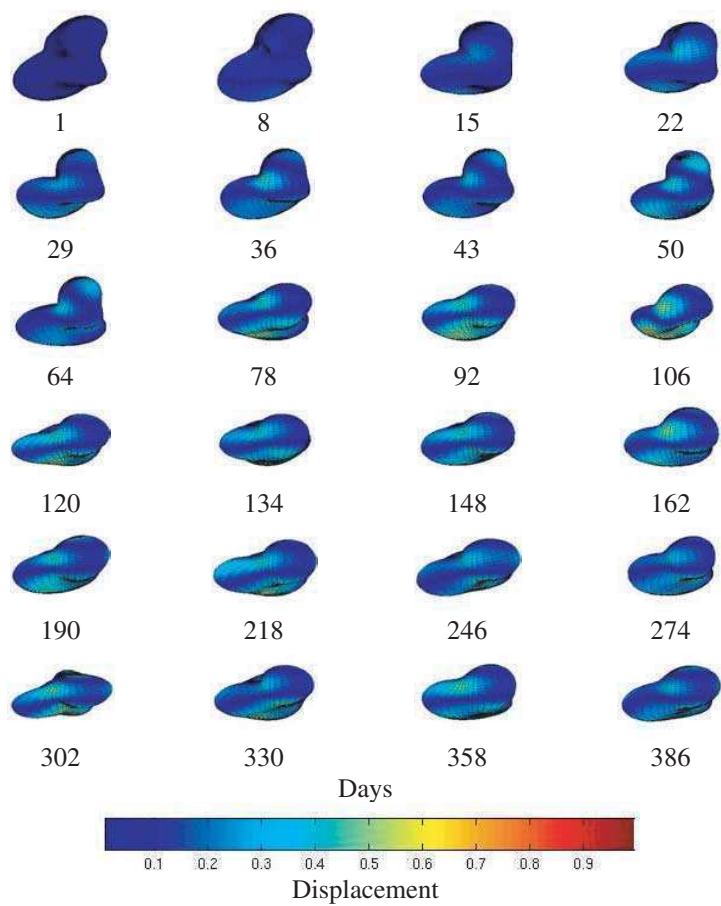
Fig. 17. The 3D reconstructions at the 24 time points of another MS lesion. This lesion seems to engulf two overlapping lesions, one extending vertically and the other diagonally. Over time, the upper part of the lesion decreases, while the lower part presents smaller variations, causing larger changes in the lesion's shape than in its size/volume. The bar of the 3D reconstructions represent the displacement (in voxels) for each time point from the first time point in which the lesion appeared. NOTE: A color version of this book chapter is available from the authors.

in the lesions' shape are larger than the changes in size/volume. Moreover, the analysis of the *MDTV* as indicated in Table 1, shows that the changes in shape, especially those represented by the *MDTV* for I_2 (related to ellipsoidal shapes) are comparable to the changes in size/volume. By studying these lesions, it can be suggested that changes in shape are indicative of the lesion's activity and that shape changes may be utilized as a tool for exploration of characteristic patterns of lesion development.

In summary, a method for studying temporal changes in individual MS lesion geometry is presented. The method provides a quantitative analysis for
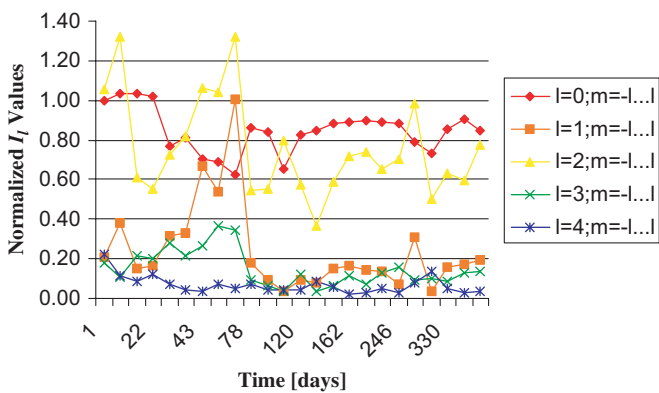
Fig. 18.   The temporal patterns for the shape indices corresponding to the lesion depicted in Fig. 17. The large changes in the lesion's shape are manifested mainly by the changes in the values for I_2, which in this case had the largest $MDTV$ values. NOTE: A color version of this book chapter is available from the authors.

characterizing MS lesion size/volume and shape over time. This method may be applied for monitoring lesion's activity and for a systematic study of lesion evolution in pursuit of characteristic mechanisms during the life cycle of MS lesions.

5. Summary

Shape analysis is becoming an important tool for the analysis of medical imaging data providing information about anatomical structures and disease that is not always available from a volumetric analysis. Future developments and improvements of shape analysis techniques will allow a more accurate quantitative estimation, representation and visualization of the changes in anatomy and diseases.

An accurate analysis of the changes in shape will eventually become particularly useful in diverse applications. For example studies for the discrimination between groups of subjects (e.g., healthy vs. unhealthy, unhealthy presenting different symptoms, etc) based on the shape and/or evolution of anatomical structures, analysis of lesion evolution by tracking temporal changes associated with certain diseases. We also believe that shape analysis will also become an important tool in the decision-making process for image-guided surgery.

Acknowledgments

Section 2 is based on work done in collaboration with Dr. E. Fisher at the Department of Biomedical Engineering in the Cleveland Clinic Foundation, Cleveland, Ohio.

Section 3 is based on work done in collaboration with K. Siddiqi, J. Pruessner and D. L. Collins of McGill University. The medical data was aquired at Montreal Neurological Institute. We are grateful to P. Dimitrov and J. Zhang for providing the examples of 2D and 3D skeletons.

Section 4 is based on a Ph.D. thesis supervised by Dr. H. Azhari from the Department of Biomedical Engineering at the Technion-Israel Institute of Technology, Haifa, Israel and Dr. A. Achiron from the Multiple Sclerosis Center at Sheba Medical Center, Tel Hashomer, Israel.

References

1. F. L. Bookstein, P. D. Sampson, A. P. Streissguth and P. D. Connor, Geometric morphometrics of corpus callosum and subcortical structures in the fetal-alcohol-affected brain, *Teratology* **64** (2001) 4–32.
2. P. F. Buckley, D. Dean, F. L. Bookstein, L. Friedman, D. Kwon, J. S. Lewin, J. Kamath and C. Lys, Three-dimensional magnetic resonance-based morphometrics and ventricular dysmorphology in schizophrenia, *Biol. Psychiatry* **45** (1999) 62–67.
3. J. R. DeQuardo, M. S. Keshavan, F. L. Bookstein, W. W. Bagwell, W. D. Green, J. A. Sweeney, G. L. Haas, R. Tandon, N. R. Schooler and J. W. Pettegrew, Landmark-based morphometric analysis of first-episode schizophrenia, *Biol. Psychiatry* **45** (1999) 1321–1328.
4. P. Tibbo, P. Nopoulos, S. Arndt and N. C. Andreasen, Corpus callosum shape and size in male patients with schizophrenia, *Biol. Psychiatry* **44** (1998) 405–412.
5. D. Velakoulis, G. W. Stuart, S. J. Wood, D. J. Smith, W. J. Brewer, P. Desmond, B. Singh, D. Copolov and C. Pantelis, Selective bilateral hippocampal volume loss in chronic schizophrenia, *Biol. Psychiatry* **50** (2001) 531–539.
6. J. J. Levitt, C. F. Westin, P. G. Nestor, R. S. Estepar, C. C. Dickey, M. M. Voglmaier, L. J. Seidman, R. Kikinis, F. A. Jolesz, R. W. McCarley and M. E. Shenton, Shape of caudate nucleus and its cognitive correlates in neuroleptic-naive schizotypal personality disorder, *Biol. Psychiatry* **55** (2004) 177–184.
7. R. E. Hogan, R. D. Bucholz and S. Joshi, Hippocampal deformation-based shape analysis in epilepsy and unilateral mesial temporal sclerosis, *Epilepsia* **44** (2003) 800–806.
8. J. A. Posener, L. Wang, J. L. Price, M. H. Gado, M. A. Province, M. I. Miller, C. M. Babb and J. G. Csernansky, High-dimensional mapping of the hippocampus in depression, *Am. J. Psychiatry* **160** (2003) 83–89.
9. E. R. Sowell, P. M. Thompson, S. N. Mattson, K. D. Tessner, T. L. Jernigan, E. P. Riley and A. W. Toga, Regional brain shape abnormalities persist into adolescence after heavy prenatal alcohol exposure, *Cereb. Cortex* **12** (2002) 856–865.
10. K. Zilles, R. Kawashima, A. Dabringhaus, H. Fukuda and T. Schormann, Hemispheric shape of European and Japanese brains: 3-D MRI analysis of intersubject variability, ethnical, and gender differences, *Neuroimage* **13** (2001) 262–271.
11. C. Studholme, V. Cardenas, A. Maudsley and M. Weiner, An intensity consistent filtering approach to the analysis of deformation tensor derived maps of brain shape, *Neuroimage* **19** (2003) 1638–1649.

12. S. M. Pizer, D. S. Fritsch, P. A. Yushkevich, V. E. Johnson and E. L. Chaney, Segmentation, registration, and measurement of shape variation via image object shape, *IEEE Trans. Med. Imaging* **18** (1999) 851–865.

13. A. Tsai, A. Yezzi, Jr., W. Wells, C. Tempany, D. Tucker, A. Fan, W. E. Grimson and A. Willsky, A shape-based approach to the segmentation of medical imagery using level sets, *IEEE Trans. Med. Imaging* **22** (2003) 137–154.

14. D. Goldberg-Zimring, A. Achiron, S. Miron, M. Faibel and H. Azhari, Automated detection and characterization of multiple sclerosis lesions in brain MR images, *Magn. Reson. Imaging* **16** (1998) 311–318.

15. C. I. Christodoulou, C. S. Pattichis, M. Pantziaris and A. Nicolaides, Texture-based classification of atherosclerotic carotid plaques, *IEEE Trans. Med. Imaging* **22** (2003) 902–912.

16. P. Clarysse, D. Friboulet and I. E. Magnin, Tracking geometrical descriptors on 3-D deformable surfaces: Application to the left-ventricular surface of the heart, *IEEE Trans. Med. Imaging* **16** (1997) 392–404.

17. A. Matheny and D. B. Goldgof, Use of three- and four-dimensional surface harmonics for rigid and nonrigid shape recovery and representation, *IEEE Trans. Pattern Anal. Machine Intel.* **17** (1995) 967–981.

18. M. S. Keshavan, G. L. Haas, C. E. Kahn, E. Aguilar, E. L. Dick, N. R. Schooler, J. A. Sweeney and J. W. Pettegrew, Superior temporal gyrus and the course of early schizophrenia: Progressive, static, or reversible? *J. Psychiatr. Res.* **32** (1998) 161–167.

19. E. H. Aylward, Q. Li, O. C. Stine, N. Ranen, M. Sherr, P. E. Barta, F. W. Bylsma, G. D. Pearlson and C. A. Ross, Longitudinal change in basal ganglia volume in patients with Huntington's disease, *Neurology* **48** (1997) 394–399.

20. R. Kikinis, C. R. Guttmann, D. Metcalf, W. M. I. Wells, G. J. Ettinger, H. L. Weiner and F. A. Jolesz, Quantitative follow-up of patients with multiple sclerosis using MRI: Technical aspects., *J. Magn. Reson. Imaging* **9** (1999) 519–530.

21. M. Miller, A. Banerjee, G. Christensen, S. Joshi, N. Khaneja, U. Grenander and L. Matejic, Statistical methods in computational anatomy, *Stat. Meth. Med. Res.* **6** (1997) 267–299.

22. G. Subsol, N. Roberts, M. Doran, J. P. Thirion and G. H. Whitehouse, Automatic analysis of cerebral atrophy, *Magnetic Resonance Imaging* **15** (1997) 917–927.

23. P. M. Thompson, C. Schwartz, R. T. Lin, A. A. Khan and A. W. Toga, Three-dimensional statistical analysis of sulcal variability in the human brain., *J. Neurosci.* **16** (1996) 4261–4274.

24. A. C. Evans, D. L. Collins and C. J. Holmes, Computational approaches to quantifying human neuroanatomical variability, in *Brain Mapping — The Methods*, eds. A. W. Toga and J. C. Maziotta (Academic Press, San Diego, CA, 1996), pp. 343–361.

25. J. C. Gee, On matching brain volumes, *Pattern Recognition* **32** (1999) 99–111.

26. C. Brechbühler, G. Gerig and O. Kubler, Parametrization of closed surfaces for 3-D shape description, *Computer Vision and Image Understanding* **61** (1995) 154–170.

27. W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery, Spherical Harmonics, in *Numerical Recipes in C — the art of scientific computing*, Vol. 2nd edn. (Cambridge University Press, Cambridge, MA, 1992), pp. 252–254.

28. W. C. Graustein, *Differential Geometry* (Dover Publications Inc., New York, 1966).

29. J. J. Koenderink, The local shape index, in *Solid Shape* (MIT Press, Cambridge, MA, 1990), pp. 319–324.

30. D. Meier and E. Fisher, Parameter space warping: Shape-based correspondence between morphologically different objects, *IEEE Trans. Med. Imaging* **21** (2002) 31–47.

31. F. L. Bookstein, Landmark methods for forms without landmarks: Morphometrics of group differences in outline shape, *Med. Image Anal.* **1** (1996) 225–243.
32. A. Rangarajan, H. Chui and J. S. Duncan, Rigid point feature registration using mutual information, *Med. Image Anal.* **3** (1999) 425–440.
33. P. Thompson and A. W. Toga, Surface-based technique for warping three-dimensional images of the brain, *IEEE Trans. Med. Imaging* **15** (1996) 402–417.
34. D. S. Meier and E. Fisher, Atlas-based anatomic labeling in neurodegenerative disease via structure-driven atlas warping, *J. Neuroimaging* **15** (2005) 16–26.
35. P. Golland, W. E. L. Grimson and R. Kikinis, Statistical shape analysis using fixed topology skeletons: Corpus callosum study, presented at IPMI, 1999.
36. M. Styner and G. Gerig, Medial models incorporating object variability for 3D shape analysis, presented at IPMI, Davis, CA, USA, 2001.
37. S. Bouix, J. Pruessner, D. L. Collins and K. Siddiqi, Hippocampal shape analysis using medial surfaces, presented at MICCAI, Utrecht, The Netherlands, 2001.
38. H. Blum, Biological shape and visual science. I, *J. Theor. Biol.* **38** (1973) 205–287.
39. A. Lieutier, Any open bounded subset of R^n has the same homotopy type than its medial axis, presented at Eighth ACM symposium on Solid modeling and applications, Seattle, Washington, USA, 2003.
40. G. Malandain, G. Bertrand and N. Ayache, Topological segmentation of discrete surfaces, *Int. J. Computer Vision* **10** (1993) 183–197.
41. M. Naf, O. Kubler, R. Kikinis, M. E. Shenton and G. Szekely, Characterization and recognition of 3D organ shape in medical image using analysis using skeletonization, presented at IEEE Workshop on Mathematical Methods in Biomedical Image Analysis, San Francisco, CA, USA, 1996.
42. D. Attali, G. S. di Baja and E. Thiel, Skeleton simplifications through non-significant branch removal, *Image Processing and Commun.* **3** (1997) 63–72.
43. K. Siddiqi, S. Bouix, A. Tannenbaum and S. W. Zucker, Hamilton–Jacobi skeletons, *Int. J. Computer Vision* **48** (2002) 215–231.
44. G. Borgefors, Distance transformation in arbitrary dimensions, *CVGIP* **27** (1984) 321–345.
45. P. Dimitrov, J. N. Damon and K. Siddiqi, Flux invariants for shape, presented at CVPR, Madision, Wisconsin, USA, 2003.
46. C. Pudney, Distance-ordered homotopic thinning: A skeletonization algorithm for 3D digital images, *Computer Vision and Image Understanding* **72** (1998) 404–413.
47. J. C. Pruessner, D. L. Collins, M. Pruessner and A. C. Evans, Age and gender predict volume decline in the anterior and posterior hippocampus in early adulthood, *J. Neurosci.* **21** (2001) 194–200.
48. J. G. Sled, A. P. Zijdenbos and A. C. Evans, A nonparametric method for automatic correction of intensity nonuniformity in MRI data, *IEEE Trans. Med. Imaging* **17** (1998) 87–97.
49. D. L. Collins, P. Neelin, T. M. Peters and A. C. Evans, Automatic 3D intersubject registration of MR volumetric data in standardized Talairach space, *J. Comput. Assist. Tomogr.* **18** (1994) 192–205.
50. J. W. Sammon, A nonlinear mapping for data structures analysis, *IEEE Trans. Comp.* C-**18** (1969) 401–409.
51. E. L. Schwartz, A. Shaw and E. Wolfson, A numerical solution to the generalized mapmakers problem: Flattening nonconvex polyhedral surfaces, *IEEE Trans. Pattern Anal. Machine Intel.* **11** (1989) 1005–1008.
52. K. J. Worsley, J. B. Poline and A. C. Evans, Characterization to response of PET and fMRI data using multivariate linear models, *Neuroimage* **6** (1998) 305–319.

53. J. C. Pruessner, L. M. Li, W. Serles, M. Pruessner, D. L. Collins, N. Kabani, S. Lupien and A. C. Evans, Volumetry of hippocampus and amygdala with high-resolution MRI and three-dimensional analysis software: Minimizing the discrepancies between laboratories, *Cereb Cortex* **10** (2000) 433–442.

54. L. Truyen, Magnetic resonance imaging in multiple sclerosis. A review, *Acta Neurol. Belg.* **94** (1994) 98–102.

55. C. J. Wallace, T. P. Seland and T. C. Fong, Multiple sclerosis: The impact of MR imaging, *AJR Am. J. Roentgenol* **158** (1992) 849–857.

56. M. Colosimo, A. Amatruda and R. P. Cioffi, Magnetic resonance imaging in multiple sclerosis: An overview, *Ital. J. Neurol. Sci.* **13** (1992) 113–123.

57. D. W. Paty, Magnetic resonance imaging in the diagnosis and follow-up of patients with multiple sclerosis, *Ital. J. Neurol Sci.* **13** (1992) 125–131.

58. S. Wiebe, D. H. Lee, S. J. Karlik, M. Hopkins, M. K. Vandervoort, C. J. Wong, L. Hewitt, G. P. Rice, G. C. Ebers and J. H. Noseworthy, Serial cranial and spinal cord magnetic resonance imaging in multiple sclerosis, *Ann. Neurol.* **32** (1992) 643–650.

59. M. Filippi, A. Campi, V. Martinelli, B. Colombo, G. Scotti and G. Comi, Brain and spinal cord MR in benign multiple sclerosis: A follow-up study, *J. Neurol. Sci.* **143** (1996) 143–149.

60. C. R. Guttmann, R. Kikinis, M. C. Anderson, M. Jakab, S. K. Warfield, R. J. Killiany, H. L. Weiner and F. A. Jolesz, Quantitative follow-up of patients with multiple sclerosis using MRI: Reproducibility, *J. Magn. Reson. Imaging* **9** (1999) 509–518.

61. H. L. Weiner, C. R. Guttmann, S. J. Khoury, E. J. Orav, M. J. Hohol, R. Kikinis and F. A. Jolesz, Serial magnetic resonance imaging in multiple sclerosis: Correlation with attacks, disability, and disease stage, *J. Neuroimmunol* **104** (2000) 164–173.

62. G. Edan, D. Miller, M. Clanet, C. Confavreux, O. Lyon-Caen, C. Lubetzki, B. Brochet, I. Berry, Y. Rolland, J. C. Froment, E. Cabanis, M. T. Iba-Zizen, J. M. Gandon, H. M. Lai, I. Moseley and O. Sabouraud, Therapeutic effect of mitoxantrone combined with methylprednisolone in multiple sclerosis: A randomised multicentre study of active disease using MRI and clinical criteria, *J. Neurol. Neurosurg. Psychiatry* **62** (1997) 112–118.

63. D. W. Paty and D. K. Li, Interferon beta-1b is effective in relapsing-remitting multiple sclerosis. II. MRI analysis results of a multicenter, randomized, double-blind, placebo-controlled trial. UBC MS/MRI Study Group and the IFNB Multiple Sclerosis Study Group, *Neurology* **43** (1993) 662–667.

64. M. Rovaris, G. Comi, M. A. Rocca, J. S. Wolinsky and M. Filippi, Short-term brain volume change in relapsing-remitting multiple sclerosis: Effect of glatiramer acetate and implications, *Brain* **124** (2001) 1803–1812.

65. D. Mc Alpine, N. D. Compston and C. E. Lumsden, *Multiple Sclerosis* (E & S Livingston London, 1998).

66. D. Goldberg-Zimring, H. Azhari, S. Miron and A. Achiron, 3-D surface reconstruction of multiple sclerosis lesions using spherical harmonics, *Magn. Reson. Med.* **46** (2001) 756–766.

67. D. Goldberg-Zimring, A. Achiron, C. R. Guttmann and H. Azhari, Three-dimensional analysis of the geometry of individual multiple sclerosis lesions: Detection of shape changes over time using spherical harmonics, *J. Magn. Reson. Imaging* **18** (2003) 291–301.

68. A. Akima, A new method of interpolating and smooth curve fitting based on local procedures, *Assoc. Comput. Mach.* **17** (1970) 589–602.

69. D. Goldberg-Zimring, A. Achiron, S. K. Warfield, C. R. Guttmann and H. Azhari, Application of spherical harmonics derived space rotation invariant indices to the

analysis of multiple sclerosis lesions' geometry by MRI, *Magn. Reson. Imaging* **22** (2004) 815–825.

70. P. Viola and W. M. Wells, Alignment by maximization of mutual information, *Int. J. Computer Vision* **24** (1997) 137–154.

71. W. M. Wells, III, W. E. L. Grimson, R. Kikinis and F. A. Jolesz, Adaptive Segmentation of MRI data, *IEEE Trans. Med. Imaging* **15** (1996) 429–442.

72. S. Warfield, J. Dengler, J. Zaers, C. R. Guttmann, W. M. Wells, 3rd, G. J. Ettinger, J. Hiller and R. Kikinis, Automatic identification of gray matter structures from MRI to improve the segmentation of white matter lesions, *J. Image Guid. Surg.* **1** (1995) 326–338.

73. S. K. Warfield, A. Robatino, J. Dengler, F. A. Jolesz and R. Kikinis, Nonlinear registration and template driven segmentation, in *Brain Warping*, ed. A. W. Toga, (Academic Press, San Diego, USA, 1999), pp. 67–84.

74. X. Wei, S. K. Warfield, K. H. Zou, Y. Wu, X. Li, A. Guimond, J. P. Mugler, 3rd, R. R. Benson, L. Wolfson, H. L. Weiner and C. R. Guttmann, Quantitative analysis of MRI signal abnormalities of brain white matter with high reproducibility and accuracy, *J. Magn. Reson. Imaging* **15** (2002) 203–209.

75. S. A. Mallat, *A Wavelet Tour of Signal Processing* (Academic Press, New York, 1999).

This page intentionally left blank

INDEX