

Multiresolution-Based Segmentation of Calcifications for the Early Detection of Breast Cancer

Clusters of microcalcifications in a mammogram may be an early indication of breast cancer. Unfortunately, due to size, shape and limited contrast from surrounding normal tissue, microcalcifications can occasionally be hard to detect in computer-aided detection (CAD) systems. These CAD systems can also be slow compared to a radiologist's performance when reviewing film-screen mammography. The research described here investigates a rapid, multiresolution-based approach combined with wavelet analysis to provide an accurate segmentation of potential calcifications. An initial multiresolution approach to fuzzy *c*-means (FCM) segmentation is employed to rapidly distinguish medically significant tissues. Tissue areas chosen for high-resolution analysis are broken into multiple windows. Within each window, wavelet analysis is used to generate a contrast image, and a local FCM segmentation generates an estimate of local intensity. A simple two-rule fuzzy system then combines intensity and contrast information to derive fuzzy memberships of pixels in the high-contrast, bright pixel class. A double threshold is finally applied to this fuzzy membership to detect and segment calcifications. This sequence of steps is shown to approach detection rates of conventional classifier designs and may therefore be useful as a pre-processing module for these systems to improve speed. Results are reported for 25 images obtained from the *Digital Database for Screening Mammography (DDSM)*.

© 2002 Published by Elsevier Science Ltd.

S. Sentelle C. Sentelle and M.A. Sutton

*Department of Computer Science, 11000 University Parkway,
University of West Florida, Pensacola,
FL 32514, USA
E-mail: msutton@uwf.edu*

Introduction

At the turn of the century, breast cancer remains the second most lethal form of cancer in women, and the most non-preventable form [1]. Early detection

provides the best chance of surviving breast cancer, while manual radiologist review of film-screen mammography remains the most cost effective and efficient way to do this. However, even with advances in film-screen technology, radiologists, due to poor

contrast in the mammogram, occasionally miss small abnormalities.

Currently, there are several companies seeking governmental approval to convert to digital systems and implement computer-aided detection (CAD) systems [2], but these systems are expensive and slow compared to film-screen reading techniques. One of the greatest concerns of practicing radiologists, as presented at the 5th International Workshop on Digital Mammography [3], is the large amount of computational time that most detection algorithms require, which can limit throughput in clinical settings unless processing can be performed off-line or coordinated by trained technical assistants.

The focus of this research is detection and enhancement of digital mammograms so that abnormalities, specifically microcalcifications, may be brought to the attention of a radiologist. Microcalcifications are small calcium deposits usually less than 0.5 mm in diameter, which appear on mammograms as small bright spots (see Figure 1) and may be an early indicator of cancer [4]. Microcalcifications that appear more in one breast than the other, that are clustered together, or that cover only one quadrant of a breast need additional clinical evaluation. However, most reevaluations or biopsies result in a benign diagnosis.

Prior Research

In the last 5 years, there has been substantial work in the area of computer-aided microcalcification detection and diagnosis. A summary of a subset of this work is

presented in Table 1. For example, Wang and Karyiannis [5] have employed a simple scheme of decomposing images with Daubechies wavelets, DAUB 4 and DAUB 20. The wavelet decomposed images were reconstructed after zeroing the lowest resolution approximation coefficients, and the inverted image was thresholded to segment the calcifications. Performance when using the DAUB 4 vs. DAUB 20 wavelet was compared. Preliminary testing in this research indicated that sub-band analysis of a wavelet decomposition could be employed to detect calcifications.

Chen *et al.* [6] used a morphological whitening filter and wavelet analysis to remove background clutter from mammograms. They sub-sampled the images, originally 2048×2048 pixels and 12 bits of gray-scale resolution, to 512×512 pixels and 8 bits of gray-scale resolution for computational efficiency. An opening relation was applied to the mammogram resulting in a residual image consisting of objects and noise. A biorthogonal wavelet was then applied for its matched filter properties. Region growing was employed with multiresolution fusion to make a final detection.

Rezaee *et al.* [7] used a multiresolution pyramidal segmentation technique. This research indicated methods in which multiresolution pyramidal segmentation could be combined with fuzzy *c*-means (FCM) to provide a final segmentation of the image. Starting with the top layer (lowest resolution), a pyramidal segmentation algorithm was applied based on both spatial relationships and similarity of parent-child pixels. This resulted in the lowest image (highest resolution) being classified into regions, usually more than desired. These regions were then merged into the desired number of




Figure 1. Image case0087LMLO from the *Digital Database for Screening Mammography (DDSM)*. The white rectangular area on the left is enlarged on the right to highlight radiologist-marked ground-truth locations of existing microcalcifications.

Table 1. Summary of related research in calcification detection algorithms, indicating authors, overview of the technique, and image database source (when available)

Authors	Relevant techniques	Image database sources
Wang and Karayiannis [5]	Daubechies wavelet decomposition with image reconstruction without the low-resolution sub-bands removing all background information. Resulting binary detection image of microcalcifications and noise.	Unreported database
Chen <i>et al.</i> [6]	Used a morphological whitening filter and wavelet analysis to remove background clutter from mammograms. An opening filter resulted in an image consisting of objects and noise. A biorthogonal wavelet was then used to simulate a matched filter.	Nijmegen database; mammogram images sampled to 512×512 pixels with 8-bit gray-scale resolution
Rezaee <i>et al.</i> [7]	Employed a multiresolution pyramidal segmentation technique classifying the highest spatial resolution image into a large number of regions. Segmentation based on spatial relationships and similarity of pixel values. These regions were then merged into an appropriate number of regions using FCM.	2 lung CT images from an emphysema study, 40 MRI images; 256×256 pixels with 12-bit gray-scale resolution
Strickland and Hahn [8]	Wavelet transformation acted as a bank of matched filters for microcalcification detection. Used thresholding and dilation in sub-bands to enhance abnormalities. Sub-bands then received different weights for reconstruction.	Nijmegen database; 40 mammogram images
Aghdasi <i>et al.</i> (1994) [9]	Studied physics of noise origins. Applied a locally adaptive smoothing filter to reduce the signal-dependant noise. Finally, a Weiner filter was applied to deblur the image. Results were sharper and contained more detail.	British Columbia Cancer clinic database; 400 mammograms images with 10-bit gray-scale resolution and $100 \mu\text{m}$ spatial resolution

classes by the FCM algorithm. This system was tested with two synthetic images containing a disk object with Gaussian noise, two CT images of patients in an emphysema study, and 40 MRIs obtained from 10 subjects with abnormalities and 10 subjects with normal images of the heart. The MR images had a size of 256×256 pixels at 12 bits of gray-scale resolution.

Strickland and Hahn [8] used a wavelet transformation that acted as a bank of matched filters for the detection of microcalcifications. They thresholded the resulting sub-bands to find the microcalcifications and dilated the microcalcifications to enhance them. They then weighted the sub-bands containing microcalcifications higher for the inverse wavelet transform to create a new output image. Their algorithm was tested with 40 mammogram images obtained from the Nijmegen database located at the University Hospital Nijmegen, The Netherlands. These 40 images represented 21 different patients, each containing at least one verified calcification cluster. The authors reported a 55% detection rate at less than 1 false positive cluster per image and nearly 100% detection rates at greater than 8 false positives per image.

Aghdasi *et al.* [9] approached the issues of noise and blurring by studying the physics of the noise source: the physical imaging system, the image digitizer, and the image display unit. They then applied a locally adaptive smoothing filter to reduce the signal-dependant noise. A Weiner filter was then used to deblur the image. As a result, all images were sharper and contained a greater amount of detail and information. Four hundred mammograms from patients who had breast biopsies at the British Columbia Cancer Agency's Vancouver clinic were obtained. Seventy of these images containing clusters of microcalcifications were digitized to a spatial resolution of $100 \mu\text{m}$ and 10-bit gray-scale resolution for testing.

Finally, Karssemeijer [10] implemented an adaptive noise equalization and statistics-based model for the detection of microcalcifications. As a pre-processing step, a robust rescaling was applied as a filter for adaptive noise equalization, which estimated noise characteristics from the image at hand. This system also took into account the impact of the variability in the tissues across the image. Microcalcifications were detected in these rescaled images using an initial segmentation and an iterative process used to update

the pixel labels. The labeling process used filtered versions of the mammogram, each of which depicted a local image feature thought to be important in distinguishing a microcalcification. This work used 25 training images and another 40 for testing, from the Nijmegen database. These images all contained one or more known clusters of microcalcifications labeled by an expert radiologist. The images were digitized at 12 bits of gray-scale resolution to a size of 2048×2048 pixels.

As Table 1 indicates, there has been a broad range of work in the area of calcification detection. In almost all cases, a contrast image is created and thresholded to generate detections [11], or a wavelet decomposition is performed and the coefficients are analyzed to provide detections of or to enhance the contrast of microcalcifications [5,9]. Contrast measures can be expensive to perform, and wavelet analysis techniques, although excellent at removing noise, enhancing calcification contrast, and detecting locations of calcifications, do not provide a direct segmentation of the calcifications without additional steps. A precise segmentation of the calcification can be critical to classifier performance as a poor segmentation can result in inaccurate feature measurements. This research presents a methodology for quickly pre-processing an image to provide a segmentation of possible calcifications with a minimal number of false positive segmentations to be provided to a classifier. The research described here differs from existing approaches by integrating segmentation and wavelet analysis into an enhancement system that: (1) applies FCM in a multiresolution approach for rapid labeling and determination of regions of interest (ROIs), (2) employs a wavelet decomposition to rapidly create a contrast image tuned to the shape and size of calcifications, (3) employs FCM to estimate local intensity properties within a ROI, and (4) combines FCM

intensity estimation and contrast through a simple fuzzy rule to detect and provide accurate segmentation of calcifications.

The remainder of the paper is organized as follows. The following two subsections provide brief overviews of the foundations for the proposed wavelet, multi-resolution, and segmentation analysis. The next subsection concludes this discussion with an overview of the database selected for this work. The subsequent section provides a description and flowchart of the overall processing and system design. The last two sections provide experimental results and suggest directions for future research.

Wavelets and multiresolution analysis

Wavelets developed in the 1970s and 1980s out of the limitations of the Fourier transform to view signals in both the temporal and frequency domains. The wavelet transform solves this problem by employing a basis filter having a compact representation in the time domain as well as in the frequency domain. This basis filter, known as a mother-wavelet, is dilated and translated to provide frequency information as a function of time or location. As shown in Figure 2, for low-frequency analysis this mother-wavelet is dilated and then translated across the signal. To encode high-frequency signals the same initial wavelet is compressed and translated. A continuous wavelet transform (CWT) allows for arbitrary dilations/contractions and translations while dilations/contractions and translations occur in factors of 2 for the dyadic wavelet transform. Eqn (1) shows how the wavelet coefficients are obtained given a wavelet basis filter that is scaled and translated by factors of 2. In this equation, a dot product is formed between $f(x)$, representing the function to encode, and Ψ representing the mother wavelet. The parameter j is the scaling factor, k is the




Figure 2. Example wavelet functions. Original mother wavelet (left), mother wavelet stretched to encode lower frequency signals (center), mother wavelet compressed to encode higher frequency signals (right).

translation factor, and the multiplication factor normalizes the dot product.

$$W(j, k) = \frac{1}{2^j} \left\langle f(x), \Psi \left(\frac{x-k}{2^j} \right) \right\rangle \quad (1)$$

Frequency, as applied to images, is equivalent to spatial resolution. A particular wavelet that is dilated or contracted encodes a specific spatial resolution. The discrete wavelet transform (DWT) is computed differently within the image processing domain. (This description pertains to wavelet processing used in MATLAB while not necessarily inclusive of all manners in which wavelet processing can be applied to the image domain.) For each wavelet there is a corresponding high- and low-pass filter, wavelet and scaling function, respectively, employed to encode frequency information. Within the image processing domain, a one-dimensional (1D) wavelet must be applied to a two-dimensional (2D) data set. In order to process the image, each of either the scaling function or wavelet filter is applied in several different orientations. These orientations are constructed by multiplying a 1D scaling function Φ and the corresponding wavelet Ψ [11]. As provided below, Eqn (2) responds to variations in the vertical direction. Eqn (3) responds to variations in the horizontal direction. Eqn (4) responds to variations along diagonals. Finally, Eqn (5) describes the scaling function, which is used to generate the approximation coefficients.

$$\Psi^1(x, y) = \Phi(x)\Psi(y) \quad (2)$$

$$\Psi^2(x, y) = \Psi(x)\Phi(y) \quad (3)$$

$$\Psi^3(x, y) = \Psi(x)\Psi(y) \quad (4)$$

$$\Phi(x, y) = \Phi(x)\Phi(y) \quad (5)$$

After each filtering process, the resulting coefficient image is decimated by $\frac{1}{2}$ in each coordinate. As a result, four different images containing the mentioned coefficients are generated each being $\frac{1}{4}$ the size of the original image. To generate multiple levels of decomposition, the wavelet decomposition is reapplied recursively to the approximation coefficients to generate the next set of coefficients. The four images generated from the approximation coefficients can be recombined to regenerate the approximation coefficients, and, therefore, the new four images can replace the approximation coefficients from the previous layer. This process can continue until too few data points are available to continue the process.

It is easily observed that the process of dilating the wavelet for each level of decomposition effectively occurs automatically at each level due to the decimation of the images. Additionally, it is interesting to note that the final decomposition is equal in size to the original image with the exception of the fact that the number of zero coefficients in the decomposed image can lead to substantial compression. Depending upon the wavelet employed for decomposition, the original image can be reconstructed at any time with a reverse decomposition.

Within this application, the wavelet coefficients can be modified prior to performing a reverse decomposition. The wavelet coefficients can also be employed directly to detect calcifications, and, once the calcifications are detected, a certain number of coefficients near the abnormality can be scaled to provide increased contrast for the calcification upon performing the reverse decomposition. In some cases, the contrast is enhanced significantly enough that a simple threshold is sufficient to create a binary detection of the calcifications in the image [9]. In this research, the wavelet coefficients can be modified to remove the lowest resolution approximation coefficients, thereby converting a resultant image into a contrast image similar to [5]. In addition, the wavelet coefficients can be scaled or zeroed out on other levels to remove sensitivity to noise and adjust sensitivity to certain sizes. For a more comprehensive discussion on wavelets, see [12].

Segmentation methods for digital mammography

Numerous supervised and unsupervised, fuzzy, knowledge-based, and neural network-based techniques have shown promising results for medical image segmentation (see e.g. [13–16]). This research employs the FCM segmentation algorithm, an unsupervised least-squares clustering approach [17]. The FCM algorithm is similar to the hard c -means (HCM) approach but employs fuzzy membership when labeling data points rather than assigning crisp labels as in HCM. Within the FCM algorithm, a set of initial prototype centers are specified, and an iterative process begins to alternately adjust prototype centers and calculate the fuzzy membership of all data points to the corresponding prototype centers. After completion, the prototype centers mark the centroids of the discovered data clusters. The final fuzzy membership for each data point can then be “hardened” so that each data point can be labeled according to which prototype it best matches. Maximal membership is employed in this research for the determination of the crisp labeling from the fuzzy membership.

Parameters to FCM include an initialization methodology, the number of classes or regions to partition the data into (c), a termination threshold to determine the amount of error allowed between partitions, a weighting exponent (m), and a maximum number of iterations allowed so that an infinite computation can be halted. FCM is initiated by assigning an initial class label to each data point determined from the initialization routine. (In this work, centers are initialized with the average of equivalent ranges across the feature space hypercube.) This assignment is stored in what is called the membership or \mathbf{U} matrix (see Eqn (6)). In this equation, k is an index into the image for data point x_k , j and i are the class indices for the prototypes v , c is the number of prototypes, and m is the fuzzy index, with a default value of 2.0.

$$u_{ik} = \left[\sum_{j=1}^c \left(\frac{\|x_k - v_i\|^2}{\|x_k - v_j\|^2} \right)^{\frac{2}{(m-1)}} \right]^{-1} \quad \forall i, k \quad (6)$$

This \mathbf{U} matrix is then used to calculate point prototypes (also called cluster centers or signatures) of the data points for each class label, which are stored in the prototype matrix \mathbf{V} (see Eqn (7)). In this equation, i is the prototype number, n is the number of pixels, x_k is the pixel value, and m is the fuzzy index.

$$v_i = \left(\sum_{k=1}^n u_{ik}^m x_k / \sum_{k=1}^n u_{ik}^m \right) \forall i \quad (7)$$

From these new prototypes, the \mathbf{U} membership matrix is then recalculated. Once the \mathbf{U} matrix has been recalculated and updated, the prototypes are then recalculated. This process continues until the measure of the distance between old prototypes and each new set of prototypes meets a threshold criterion. In this research, the Euclidean distance between the old and the new prototypes for each label is calculated and if all distances are below the termination threshold, the classification will halt. At termination, the prototypes for the newly discovered data clusters as well as labels for the data points can be retrieved. An example segmentation is shown in Figure 3.

Database selection

Several digital mammography databases are available for use in the development and formal assessment of algorithms, and the increasing use of these will greatly facilitate the creation of reliable, real-time systems feasible for clinical settings [18]. Given this, our research is based on the use of data sets derived from the recently completed *Digital Database for Screening Mammography (DDSM)*, located at the University of South Florida [19]. The *DDSM* contains approximately 2500 cases of fully annotated mammographic images. Each case contains two views of each breast along with patient information such as age at the time of the study, American College of Radiology (ACR) breast density rating, subtlety rating for abnormalities, and ACR descriptions of each abnormality. Information about how the mammograms were digitized is also included, such as the relevant spatial resolution and scanner used. Also provided with the database is software for




Figure 3. Rapid FCM segmentation with seven classes. Original image is case 0087LMLO from the *DDSM* database.

uncompressing the images, creating the ground-truth images, and performance algorithms for standardized evaluation.

This research used cases submitted to this database from a local mammography clinic, the Ann L. Baraco Center for Women's Health, located at Sacred Heart Hospital in Pensacola, Florida. As described in Table 2, 20 images with malignant microcalcifications were divided into four sets according to density and subtlety ratings: (1) high-density breast tissue with obvious abnormalities, (2) high-density breast tissue with subtle

abnormalities, (3) low-density breast tissue with obvious abnormalities and (4) low-density breast tissue with subtle abnormalities. A summary of the corresponding *DDSM* ground-truth information is also provided in Table 2.

In Table 2, column 1 indicates the case number and view. Column 2 indicates the Breast Imaging Reporting and Data System (BI-RADSTM) density ratings for these cases, using a scale ranging from 1 to 4, as follows [4]: 1 = The breast is almost entirely fat. 2 = There are scattered fibroglandular densities that could obscure a

Table 2. Summary of groundtruth information from *DDSM* database. Training images are highlighted in bold. (*R*=right, *L*=left, *MLO*= mediolateral oblique view, *CC*= craniocaudal view)

	Case #/view	Density	Assessment	Subtlety	Pathology	Lesion type/description
High density/subtle	c0060LCC	3	4	2	Malignant	1 Calcification/amorphous distribution: clustered
	c0060LMLO	3	4	2	Malignant	1 Calcification/amorphous distribution: clustered
	c150LCC	3	4	1	Malignant	2 Calcifications/amorphous distribution: clustered
	c0150LMLO	3	4	1	Malignant	2 Calcifications/amorphous distribution: clustered
	c0198RCC	3	4	2	Malignant	1 Calcification/amorphous distribution: segmental
	c0198RMLO	3	4	2	Malignant	1 Calcification/amorphous distribution: segmental
Low density/subtle	c0002LCC	2	4	2	Malignant	1 Calcification/pleomorphic distribution: segmental
	c0002LMLO	2	4	1	Malignant	1 Calcification/pleomorphic distribution: segmental
	c0020LMLO	2	4	2	Malignant	1 Calcification/pleomorphic distribution: regional
	c0169RCC	2	4	2	Malignant	1 Calcification/amorphous distribution: clustered
	c0169RMLO	2	4	2	Malignant	1 Calcification/amorphous distribution: clustered
High density/obvious	c0036RCC	4	4	4	Malignant	1 Calcification/punctate distribution: clustered
	c0036RMLO	4	4	4	Malignant	1 Calcification/punctate distribution: clustered
	c0120LCC	4	5	5	Malignant	1 Calcification/amorphous-pleomorphic distribution: segmental
	c0120LMLO	4	5	5	Malignant	1 Calcification/amorphous-pleomorphic distribution: segmental
Low density/obvious	c0045RMLO	2	4	5	Malignant	1 Calcification/pleomorphic distribution: clustered
	c0087LCC	2	5	5	Malignant	3 Calcifications/pleomorphic distribution: clustered
	c0087LMLO	3	5	5	Malignant	3 Calcifications/pleomorphic distribution: clustered
	c0214RCC	2	4	4	Malignant	1 Calcification/pleomorphic distribution: segmental
	c0214RMLO	2	4	4	Malignant	1 Calcification/pleomorphic distribution: segmental

lesion on a mammogram. 3 = The breast is heterogeneously dense. This may lower the sensitivity of mammogram. 4 = The breast is extremely dense which lowers the sensitivity of mammogram.

Thus for this metric, a lower density rating implies that the abnormalities may be more easily seen by the radiologist (and by the computer). Column 3 provides a BI-RADS™ assessment rating for each view containing an abnormality, using a scale ranging from 1 to 5, as follows: 1 = Negative. 2 = Benign finding. 3 = Probably benign finding — short interval follow-up suggested. 4 = Suspicious abnormality — biopsy should be considered. 5 = Highly suggestive of malignancy — appropriate action should be taken.

Column 4 contains a subtlety rating utilized in the DDSM data set. This measure has a value in the range 1–5, where 1 means “subtle” and 5 means “obvious”. In this case, a higher subtlety rating implies the interpretation task should be easier for both the mammographer and the computer (relative to a lower rating). Columns 5 and 6 provide the pathology and lesion type and indicate that all the microcalcifications in the views containing abnormalities were in fact determined to be malignant microcalcifications.

As a control group, five images with no abnormalities (normal) were also selected. One of these five was used for training purposes. These images were also taken from cancer volumes, but were views that were determined by a practicing radiologist to not have an abnormality. These images were selected randomly from the subset of all possible normal views of abnormal cases selected. For example, the left views of case0087 contained diagnosed abnormalities; however the right breast views of the same case contained none.

All images selected for this work had a spatial resolution of 50 μm and a 12-bit gray-scale resolution. Due to memory limitations and computational overhead, images were reduced to 8-bit grayscale. The system then analyzed the entire images without any cropping.

Experimental Methodology and System Design

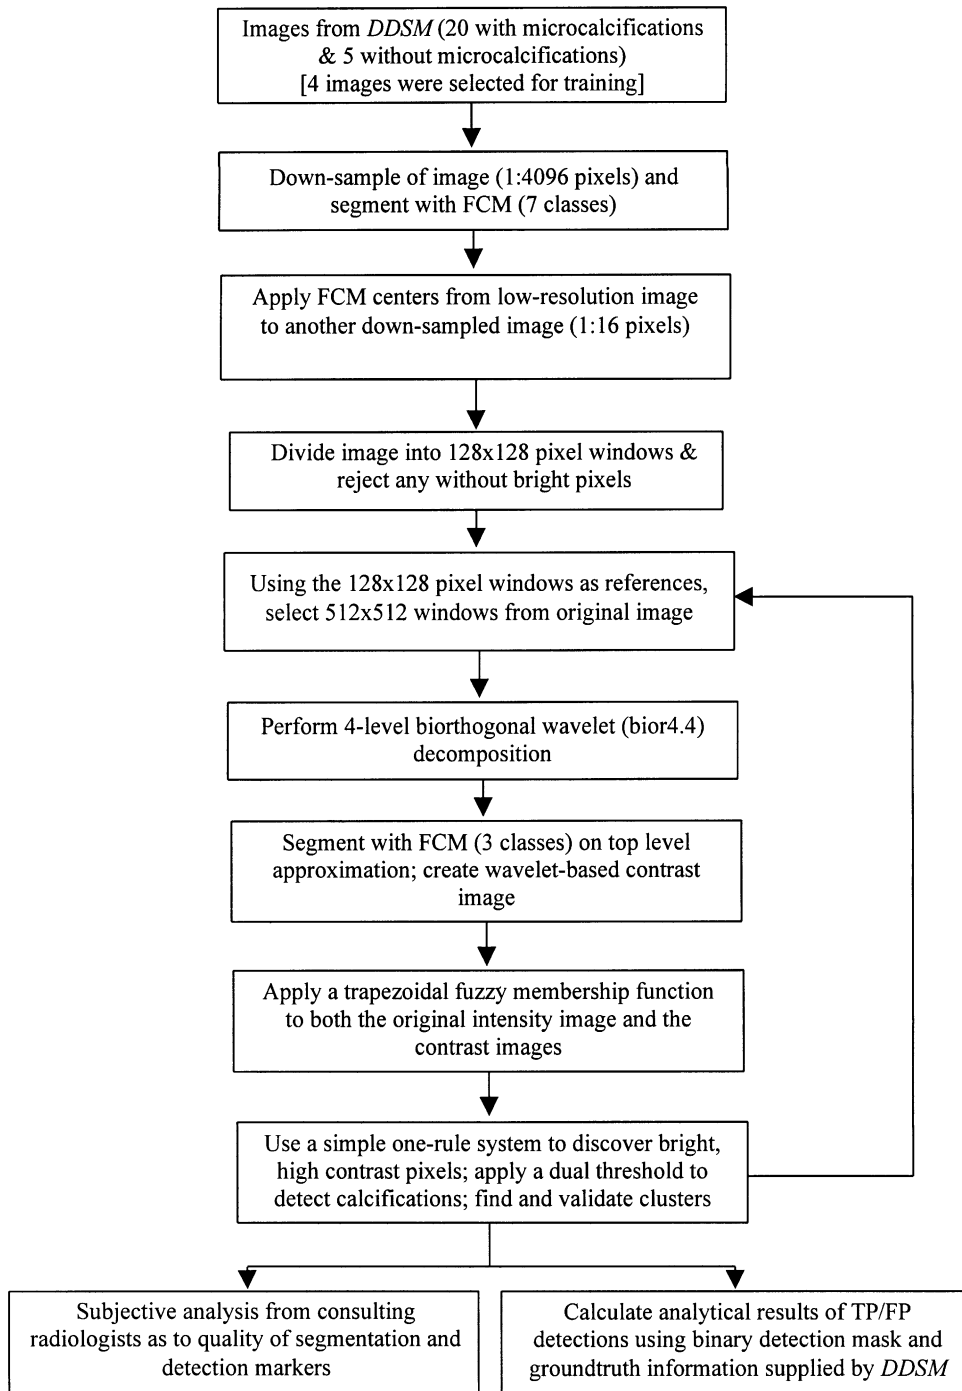
Utilizing the globally accessible database and algorithms described above, a system was developed, which begins by separating mammographic tissue from non-tissue

regions in the image. Previous work in this area focused on using the FCM algorithm for segmentation (see [20]). Used traditionally, however, this algorithm is computationally expensive due to its memory requirements, iterative nature and slow convergence for large data sets. The procedure proposed here employs a multiresolution approach to minimize data points processed by the FCM algorithm for segmentation. Once areas of the image requiring analysis are identified, FCM is employed to estimate local image properties, and wavelet analysis is employed to provide a contrast measure. Two images are then created containing fuzzy memberships for high-contrast pixels and high-intensity pixels. A fuzzy AND operation is performed to generate fuzzy membership for high-contrast and high-intensity pixels. A dual threshold applied to the final fuzzy memberships allows for detection and segmentation or “hardening” of the membership values. The steps for the algorithm are summarized in Table 3. The following will describe each of these steps in more detail.

The first step of the calcification detection algorithm is to determine likely areas in the image that may contain calcifications. Previous work on sub-sampling [20,21] has indicated that the FCM algorithm can successfully discover prototypes by clustering a subset of pixels, which are statistically representative of the original image. FCM is an averaging algorithm, and therefore, classification is based on a majority of points, such that sparse outlier data points will be averaged out. These “outlier” points tend to represent the details in the high-resolution image, such as microcalcifications, suggesting that FCM can be applied to a reduced resolution image to discover the prototypes, since lower resolutions will contain more homogenous regions. In this implementation, FCM is applied to a down-sampled version of the original image for the initial analysis. A pixel is obtained from every 4096 pixels in the image (64 in the x direction and 64 in the y direction), yielding a total of 3000–4000 data points to analyze.

As seen in Figure 3, FCM is applied to obtain a set of seven prototypes or clusters of intensity values from these points. Heuristically, it was discovered that the following seven classes could be roughly attained from the image: (1) pectoral muscle and bright structures, (2) fibroglandular tissue, (3) adipose tissue, (4) one class representing image background or the off-breast part of the image and (5) three classes devoted to the air–skin interface.

Table 3. Overall flow of the developed microcalcification detection algorithm. An input image is processed at low resolution to determine windows of interest for additional high-resolution microcalcification analysis and detection



Next, the system applies the prototypes discovered at the lowest resolution to a higher resolution version of the image. For this work, that image is another down-

sampled image created by taking 1 pixel for every 16 pixels (4 in the x direction and 4 in the y direction) in the original image. The prototypes are applied to this

resolution instead of the original image because it is of sufficient detail to detect relevant tissue regions, while remaining a small enough data set as to result in rapid labeling. This image is labeled using the prototypes discovered earlier by performing a single iteration of the FCM membership function (Eqn (6)), generating a U matrix from the existing prototypes.

This image is then divided into approximately 128×128 pixel windows, which are processed to reject areas that do not contain tissue pixels. This window size was found experimentally to provide optimal performance in detection rates and processing times. The number of pixels belonging to either of the two brightest intensity classes as well as to the darkest intensity class (background), discovered from the previous FCM application, is computed. Those 128×128 windows containing background pixels are rejected from future analysis. The number of allowable background pixels is, of course, dependent upon the window size employed for analysis. Within the window size used here, this generally prevents analysis of the air–skin interface. Generally, the air–skin interface was not observed to contain calcifications within the images analyzed during this research. Windows surviving rejection are then used to derive pointers to full resolution 512×512 pixel windows in the original image, which will require additional analysis for microcalcification detection.

The second half of the processing algorithm begins with a 4-coefficient biorthogonal wavelet filter (bior4.4 in MATLAB) applied to each 512×512 window to produce a four-level decomposition structure. The fourth-level approximation coefficients are segmented into three classes with the FCM algorithm. Any zeros within the approximation coefficients generated by zero padding are excluded before segmentation. This segmentation step provides an estimate of local intensity properties, which can vary substantially across the image. Prototypes from this segmentation are compared with the highest three prototypes from the first, global FCM segmentation, using the Euclidean measure, to ensure that prototypes in the local image roughly match those from the global image. These prototypes correlate to three of the seven classes representing roughly fibroglandular tissue, adipose tissue, and bright structures. If the difference exceeds a threshold (40 was used here), the current sub-image is rejected from further analysis. Typically, this will remove areas near the air–skin interface which were not excluded in the first part of the algorithm due to the fact that no background pixels were in this image.

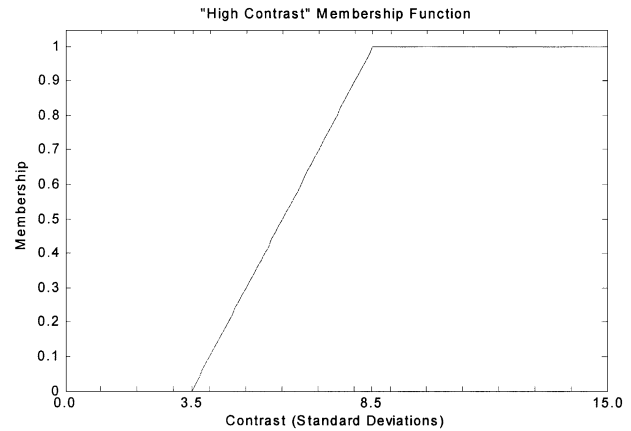


Figure 4. Trapezoid membership function based on standard deviation of the image’s contrast.

As referenced in [5], calcifications can be detected during wavelet processing by simply removing the lowest resolution approximation coefficients and performing a reverse wavelet transform. The reverse wavelet transform output is then thresholded appropriately to provide binary detections of calcifications. This methodology is performed in this research with the following modifications:

- (1) The first level wavelet detail coefficients are dominated by noise and are zeroed out to reduce sensitivity to lower contrast, single-pixel fluctuations.
- (2) The fourth-level wavelet detail coefficients are scaled by a factor of $\frac{1}{2}$ to reduce sensitivity to structures larger than the expected calcification sizes.

After modifying the wavelet coefficients, a reverse wavelet transform is performed to produce a “tuned” contrast image [22].

Instead of performing a “hard” threshold on the contrast image, a trapezoidal membership function is next applied to the wavelet coefficients to identify “high-contrast” pixels. This membership function is created dynamically by estimating the amount of noise in the contrast image. To do this, an image is created using the inverse of the first level detail wavelet coefficients. Standard deviations are calculated from these data and are used to create the trapezoidal membership function (see Figure 4). The trapezoidal membership function observed in Figure 4 is generated based upon experimentation and appears optimal for the detection of microcalcifications. The trapezoidal membership function is applied to the contrast image to create a




Figure 5. Original hand-cropped image with microcalcifications (top left), output image after fuzzification of contrast with trapezoidal membership equation (top right), output image after fuzzification of intensity with trapezoidal membership equation (bottom left), and final output image after simple fuzzy rule and thresholding (bottom right).

membership of each pixel for belonging to the high-contrast class (see top-right image of Figure 5).

Intensity estimation is used as a second feature in this implementation. To do this, the prototypes generated from the local FCM segmentation are used to dynamically construct an intensity trapezoidal membership function allowing description of bright objects in the local image. This trapezoidal membership function scales intensities between the second highest prototype multiplied by a scale factor of 0.9 and the third highest prototype to fuzzy values between 0 and 1. Applying this membership function to the 512×512 window results in a fuzzy intensity image (see lower left image of Figure 5).

Estimation of pixel intensity is important for a couple of reasons. First, in order to detect subtle calcifications, the wavelet contrast image is thresholded at a low level. This results in possible sensitivity to dim objects that

should not be classified as calcifications. Second, existing literature suggests that noise energy in the mammogram image is a function of background intensity, noise being lower in bright areas of an image and higher in regions containing areas of lower gray-scale intensity [11]. The combination of intensity with the wavelet-based contrast image appears to create an equalization effect that significantly reduces the false positive detections in “darker” areas of the image where noise is observed to be greater.

A simple one-rule fuzzy logical AND function is applied to the fuzzy contrast and fuzzy intensity image windows resulting in a fuzzy detection output identifying areas as being both high contrast and high intensity (see lower right image of Figure 5). A dual threshold is employed on this image to segment calcifications. A first threshold is applied to this output to “harden” and provide microcalcification detections. A second

threshold acts to segment the calcifications, providing a more accurate description of the calcification boundary after the initial detection step. For this research, the first threshold is varied to alter the detection rate of the algorithm for ROC curve generation, while the second threshold is set to 0.3. Finally, a morphological operator is applied to clean solitary pixel detections.

After detection of potential calcifications has occurred, an algorithm is applied to analyze clusters of calcifications. The calcification centers are estimated and a Euclidean distance matrix is generated indicating the distance, in number of pixels, between every calcification in the image. The matrix is thresholded to find distances less than a certain threshold (90 pixels or 450 μm in this instance). A recursive algorithm then searches the thresholded matrix for groups of calcifications in the image window. If a prospective cluster contains three or more calcifications, density is measured by dividing the total number of calcifications by the area of a bounding box surrounding the region. If the density meets or exceeds a density threshold, a region of interest (ROI) marker is placed in the centroid of the cluster density. The density threshold, here set at 5 calcifications/ cm^2 , is based upon detecting 5 calcifications within a 1 cm^2 area.

Results and Discussion

After creating the algorithm based on its performance on the four training images, the algorithm was tested

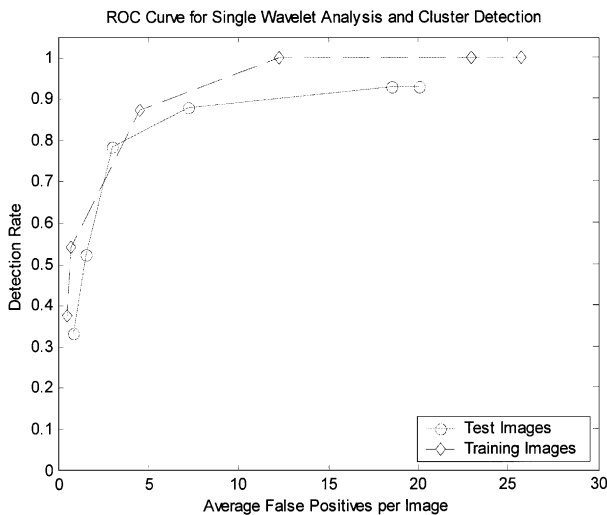


Figure 6. ROC curve for single wavelet analysis algorithm demonstrating performance on training and testing datasets derived from *DDSM* database.

with 21 additional images. The ROC curves shown in Figure 6 were generated by varying the detection threshold to “harden” the fuzzy membership for high-contrast, bright pixels. The images were reduced to an 8-bit gray-scale resolution before being initially analyzed. However, the full 50 μm resolution inherent in the original image was utilized at specific points of analysis. All processing was performed on a Sun Ultra 60 workstation running SunOS 5.7 with 512 MB of RAM.

The ROC curve represents detection rates for clusters as opposed to individual calcifications. Results indicate an average of 3.0 false positives per image is attained at the 75% detection rate and an average of 17 false positives per image is attained for the 94% detection rate. As can be seen in Figure 6, this algorithm’s detection performance does not exceed some of the results reported in other research. For example, Karssemeijer [10] reported detection rates of 90% with 1 false positive per image using an alternative dataset provided by the Nijmegen Hospital. In other research,

Table 4. Tabulated results for single wavelet implementation. Training images are highlighted in bold

	Case	Number of true positive regions	True positive rate	False positives	Fuzzy hardening threshold
High density/subtle	c0060LCC	1	1	4	0.5
	c0060LMLO	1	1	13	0.4
	c0150LCC	2	0.5	2	0.7
	c0150LMLO	2	0.5	1	0.7
	c0198RMLO	1	1	1	0.7
Low density/subtle	c0198RCC	1	1	1	0.7
	c0002LCC	1	1	2	0.7
	c0002LMLO	1	1	1	0.7
	c0020LMLO	1	1	1	0.7
	c0169RCC	1	0	14	0.4
High density/obvious	c0169RMLO	1	1	6	0.4
	c0036RCC	1	1	0	0.7
	c0036RMLO	1	1	1	0.7
	c0120LCC	1	1	1	0.4
	c0120LMLO	1	1	0	0.4
Low density/obvious	c0045RMLO	1	1	2	0.8
	c0087LCC	2	1	1	0.7
	c0087LMLO	3	0.7	2	0.7
	c0214RCC	1	1	2	0.9
	c0214RMLO	1	1	3	0.9
Normals	c0087RMLO	0	N/A	2	0.7
	c0120RMLO	0	N/A	3	0.4
	c0169LMLO	0	N/A	4	0.4
	c0150RMLO	0	N/A	0	0.7
	c0198LMLO	0	N/A	2	0.7

Woods [23] reported 1 false positive per image with a 79% detection rate, and approximately 10 false positives at a 95% detection rate. The performance of the algorithm described in this research is surprising in that this algorithm performs nearly as well as algorithms possessing a classifier or neural network, such as the work described in [23] and other sources. It is expected that a classifier would be employed to further reduce false positives and provide classification of the potential calcifications segmented by this algorithm.

While the overall detection performance of this algorithm appears good, a more detailed analysis of individual cases reveals additional insight into the performance. As described in the section titled “Database selection” the 25 images selected for this work were classified as four tissue types: (1) low-density tissue with obvious calcifications, (2) low-density tissue with subtle calcifications, (3) high-density tissue with obvious calcifications, and (4) high-density tissue with subtle calcifications. The algorithm has a different response to these four tissue classes, as can be seen in Table 4. The data in Table 4 were generated by reporting the threshold and false positive rate for the highest threshold allowing for maximal probability of detection.

In general, this algorithm performs well in detecting obvious and subtle calcifications, in dense tissue. In fact, subtle calcifications were detected in dense tissue at an 80% detection rate with less than 4 false positives per image. Several characteristic features were noted with mammograms marked as being high density. First, high-contrast linear structures are somewhat obscured in these images. These linear structures are often a source of false alarms. Secondly, noise levels tend to be lower in dense tissue regions as noted during observation of the wavelet coefficients and as suggested by Karssemeijer [10].

Mammograms marked as low density tended to present a more significant challenge for this algorithm in that a large number of false positives resulted from the analysis of these images. Two primary sources were observed for these false alarms. First, noise appears higher in low-density tissue regions. Secondly, numerous false alarms appeared to be triggered by the bright, thin, linear structures in the images. In many cases, the algorithm was observed to mark several false detections along the linear structures, as shown in Figure 7. These findings support the need for detecting linear structures and for compensating for noise levels that vary as a function of tissue density or image intensity.

Two specific cases, case0169 and case0150, represented a significant challenge for calcification detection. These images contained subtle calcifications in low-to-medium density tissue. Subjective review of the case0169RMLO highlights the possibility that image quality was a contributing factor to poor performance. Observation of the histogram for this image, when compared to other images, revealed a narrow-intensity histogram suggesting the overall contrast of the images was poor.

As described earlier, the ROC curve is generated by adjusting the threshold employed for binary detection of calcifications in the fuzzy “high-contrast, high-intensity” images. Although the ROC curve is generated using the same threshold across all the images, it is noted that an optimal threshold can be found for each image and, more generally, for each class of images as described previously. Density and overall contrast of an image appear to be the largest factors in determining the optimal threshold.

Table 5 reflects the average amount of time required to process a typical mammogram image. On average, an image was processed in just over 3 min. The 21 cases used for testing consisted of two different image sizes.




Figure 7. The bounding box shown on the leftmost image highlights two false-positive markers for case0214RMLO. The upper right image shows the high-contrast linear structures apparent in this region. The bottom right image shows the false-positive calcification detections generated by the linear structures.

Table 5. Average time required to process images on Sun Ultra60 running SunOS5.7 with 512 MB of RAM

Image case	SUN workstation processing time (s)	Image size (Pixels)
c0002LCC	265	5928 × 3776
c0002LMLO	362	5824 × 4104
c0020LMLO	307	5832 × 3664
c0036RCC	64	4672 × 2848
c0036RMLO	83	4640 × 2776
c0045RMLO	399	5888 × 3832
c0060LCC	133	4704 × 2624
c0060LMLO	148	4704 × 2512
c0087LMLO	199	4824 × 2936
c0087RMLO	125	4688 × 2528
c0087LCC	167	4696 × 2704
c0120LCC	96	4824 × 2632
c0120LMLO	117	4808 × 2632
c0120RMLO	89	4760 × 2464
c0169RCC	437	5920 × 3800
c0169RMLO	298	5912 × 3760
c0169LMLO	292	5856 × 3360
c0150LCC	168	4696 × 2688
c0150RMLO	116	4688 × 2528
c0150LMLO	159	4736 × 2464
c0198RMLO	93	4704 × 2448
c0198LMLO	112	4752 × 2480
c0198RCC	57	4752 × 2288
c0214RCC	375	5888 × 3624
c0214RMLO	312	5912 × 3776
Average time	198	

Images were either of a size of approximately 4800×3800 pixels or approximately 6000×4000 pixels. This results in two drastically different processing times. The smaller images were processed in less than 200 s while the larger images were processed in 300–400 s. In the literature, for comparison, Karssemeijer maintains outstanding detection curves while processing a 2048×2048 image in 3 min on an HP 710 workstation [10]. These images were, however, pre-cropped from the original images and reduced to an 12-bit gray-scale resolution.

The overall performance of this algorithm is significant considering the platform and environment used to test the algorithm. For the purposes of developing a rapid prototype, the algorithm was developed in a MATLAB environment. In addition to being an interpretive language, many of the algorithms supplied with MATLAB, such as the FCM algorithm, are not optimal for time or memory consumption. It is expected that implementing more efficient algorithms within a compiled environment would result in a several-fold increase in performance.

Several additional factors are worth noting regarding the performance of this algorithm. First, multiresolution analysis enables processing of smaller units of data. This can prevent virtual page swapping as well as allow a smaller number of data points to be provided to parts of the algorithm maintaining a computational complexity much higher than $O(N)$. This was especially important for the FCM algorithm, which is iterative in nature. The FCM algorithm was crucial in determining the areas of the image requiring further analysis, often reducing the total percentage of to-be-processed pixels by 50%. Secondly, wavelet analysis replaces computationally expensive convolutions normally required for measuring contrast. Finally, if the false alarm rate of the pre-processing step is low enough, it is possible to perform cluster analysis reducing the number of false positives requiring detailed analysis.

Conclusion and Future Work

This research shows that a multiresolution, FCM-based segmentation combined with wavelet processing can be employed to quickly detect and segment calcifications as a possible pre-processing step for a classifier. A multi-resolution segmentation employing FCM can be employed to dynamically “crop” the image, leaving high-resolution analysis for as a little as 50% of the pixels in the entire image. The combination of intensity and the wavelet contrast image appears to provide reduced false alarms while also providing enhanced segmentation of individual calcifications over that of using either feature alone. Wavelet analysis can be employed to replace the conventional contrast measures, reducing the overall computational complexity of the calcification detection algorithm. Finally, if the pre-processing step maintains a minimal false positive rate, it is possible to perform cluster analysis prior to classifying individual calcifications to further reduce false positives and the work of a subsequent classifier.

This algorithm exhibits difficulty with the detection of some subtle calcifications, and a large number of false positives can occur in images containing low-density tissue as well as strong contrast. Overall, it is expected that the addition of linear structure detections, finer tuned fuzzy rules, local measurement of standard deviation, and enhanced use of the FCM algorithm output during the pre-processing stage could result in significant improvements in performance.

Continuing work on this project includes obtaining input from consulting radiologists to provide subjective feedback as to the quality and medical relevancy of the image enhancements. Also a side-by-side comparison of the output images to output from the R2 ImageChecker system [24] is under way. Several other areas of this research are also being considered for future work. First concern is to determine precisely and quantitatively the image qualities that are driving the significant changes in performance of the algorithm. It is suspected that the overall image quality and/or image contrast affects the output of the algorithm and, if this measure can be found, the algorithm could be enabled to automatically adjust thresholds based upon the estimated image properties. Also, it has been generally observed that individual calcifications in a cluster of calcifications tend to have similar properties in terms of size, contrast, and intensity. As a result, it is possible that a final validation step of cluster detection might include an estimation of the similarity of the calcifications in the cluster. Finally, future work will consist of moving this algorithm to a compiled environment to provide a better estimate of the achievable computational speeds.

Acknowledgements

Funding for this research project was provided to Dr. Melanie A. Sutton by The Whitaker Foundation, under a Biomedical Engineering Research Grant (RG980414). We would like to thank Suzanne Wooldridge, M.D. and Patricia Griffith, M.D. of the Ann L. Baroco Center for Women's Health at Sacred Heart Hospital for feedback on initial images produced as part of this research. We would also like to express our appreciation to Dr. Michael Heath at the University of South Florida for assisting in using the *DDSM* data set. Researchers interested in obtaining a copy of the code for their own research may download it from our www site: <http://uwf.edu/ria>.

References

1. American Cancer Society (2000) *Cancer Facts & Figures 2000*. Atlanta, GA: ACS, Inc.
2. Moore, S.K. (2001) Better breast cancer detection. *IEEE Spectrum*. New York, NY: IEEE, Inc., pp. 50–54.
3. Gale, A.G., Muggleston, M.D., Cowley, H.C. & Wooding, D. (2001) Human factors considerations for CAD implementation in breast screening. In: Yaffe, M.J. (ed.), *IWDM 2000 5th International Workshop on Digital Mammography*. Madison, WI: Medical Physics Publishing, pp. 461–467.
4. American College of Radiology (ACR) (1995) *Breast Imaging Reporting and Data System (BIRADS)* (2nd edn). Reston, VA: American College of Radiology.
5. Wang, T.C. & Karayiannis, N.B. (1998) Detection of microcalcifications in digital mammograms using wavelets. *IEEE Transactions on Medical Imaging* **17**: 498–509.
6. Chen, L., Chang, W.C. & Parker, K.J. (1997) Small object detection using morphological filtering and multiresolution analysis with application to microcalcification detection in mammograms. *SPIE Proceedings 2962*. Bellingham, WA: SPIE Press, pp. 14–25.
7. Rezaee, M.R., van der Zwet, P.M.J., Lelieveldt, B.P.F., van der Geest, R.J. & Reiber, J.H.C. (2000) A multi-resolution image segmentation technique based on pyramidal segmentation and fuzzy clustering. *IEEE Transactions on Image Processing* **9**: 1238–1248.
8. Strickland, R.N. & Hahn, H.I. (1996) Wavelet transforms for detecting microcalcifications in mammograms. *IEEE Transactions on Medical Imaging* **15**: 218–229.
9. Aghdasi, F., Ward, R.K. & Palcic, B. (1994) Restoration of mammographic images in the presence of signal-dependant noise. In: Bowyer, K.W. & Astley, S. (eds), *State of the Art in Digital Mammographic Image Analysis*. Singapore: World Scientific, pp. 42–63.
10. Karssemeijer, N. (1994) Adaptive noise equalization and recognition of microcalcification clusters in mammograms. In: Bowyer, K.W. & Astley, S. (eds), *State of the Art in Digital Mammographic Image Analysis*. Singapore: World Scientific, pp. 148–166.
11. Hubbard, B.B. (1998) *The World According to Wavelets: The Story of a Mathematical Technique in the Making* (2nd edn). Natick, MA: A K Peters, Ltd.
12. Daubechies, I. (1992) *Ten Lectures on Wavelets (CBMS-NSF Regional Conference Series in Applied Mathematics, No. 61)*. Philadelphia, PA: SIAM.
13. Fletcher-Heath, L.M., Hall, L.O., Goldgof, D.B. & Murtagh, F.R. (2001) Automatic segmentation of non-enhancing brain tumors in magnetic resonance images. *Artificial Intelligence in Medicine* **21**: 43–63.
14. Clark, M., Hall, L., Goldgof, D., Velthuisen, R., Murtagh, F.R. & Silbiger, M. (1998) Automatic tumor segmentation using knowledge-based techniques. *IEEE Transactions on Medical Imaging* **17**: 187–201.
15. Sutton, M.A., Bezdek, J.C. & Cahoon, T.C. (2000) Image segmentation by fuzzy clustering: methods and issues. In: Bankman, I.N. (ed.), *Handbook of Medical Imaging Processing and Analysis*. San Diego, CA: Academic Press, Inc., pp. 87–106.
16. Hall, L.O., Bensaid, A.M., Clarke, L.P., Velthuisen, R.P., Silbiger, M.S. & Bezdek, J.C. (1992) A comparison of neural network and fuzzy clustering techniques in segmenting magnetic resonance images of the brain. *IEEE Transactions on Neural Networks* **3**: 672–682.
17. Bezdek, J.C., Keller, J., Krisnapuram, R. & Pal, N.R. (1999) *Fuzzy Models and Algorithms for Pattern Recognition and Image Processing*. Norwell, MA: Kluwer Academic Publishers.
18. Sutton, M.A. & Bezdek, J.C. (1997) Enhancement and analysis of digital mammograms using fuzzy models. *SPIE*

- Proceedings 3240 on Applied Imagery and Pattern Recognition: Exploiting New Image Sources and Sensors*. Bellingham, WA: SPIE Press, pp. 179–190.
19. Heath, M., Bowyer, K., Kopans, D., Moore, R. & Kegelmeyer, P. (2001) The digital database for screening mammography. In: Yaffe, M.J. (ed.), *IWDM 2000 5th International Workshop on Digital Mammography*. Madison, WI: Medical Physics Publishing, pp. 212–218. Available at: <http://marathon.csee.usf.edu/Mammography/Database.html>. Accessed December 19, 2001.
 20. Sutton, M.A., Marin, A., Sentelle, S., Wooldridge, S., Griffith, P. & Dell, C. (2001) Development and assessment of protocols for efficient utilization of large-scale digital mammography databases. In: Yaffe, M.J. (ed.), *IWDM 2000 5th International Workshop on Digital Mammography*. Madison, WI: Medical Physics Publishing, pp. 777–784.
 21. Shankar, B.U. & Pal, N.R. (1994) FFCM: an effective approach for large data sets. *Proceedings of the 3rd International Conference on Fuzzy Logic, Neural Nets and Soft Computing*, pp. 331–332.
 22. Laine, A., Fan, J. & Yang, W. (1995) Wavelets for contrast enhancement of digital mammography. *IEEE Engineering in Medicine and Biology Magazine* **14**: 536–550.
 23. Woods, K.S. (1994). Automated image analysis techniques for digital mammography, PhD dissertation, University of South Florida, Tampa, FL.
 24. R2 Technology, Inc. (1998) *R2 ImageChecker M1000 System*. Los Altos, CA: R2 Tech., Inc.