

CERIAS Tech Report 2001-141
Multiresolution Detection of Spiculated Lesions in Digital Mammograms
by S Liu, C Babbs, E Delp
Center for Education and Research
Information Assurance and Security
Purdue University, West Lafayette, IN 47907-2086

MULTIRESOLUTION DETECTION OF SPICULATED LESIONS IN DIGITAL MAMMOGRAMS

Sheng Liu †, Charles F. Babbs ‡, and Edward J. Delp †

†Video and Image Processing Laboratory (*VIPER*)

School of Electrical and Computer Engineering

‡Department of Basic Medical Sciences

School of Veterinary Medicine

Purdue University

West Lafayette, Indiana 47907

Corresponding Author:

Professor Edward J. Delp

School of Electrical and Computer Engineering

1285 Electrical Engineering Building

Purdue University

West Lafayette, IN 47907-1285

USA

Telephone: +1 765 494 1740

Fax: +1 765 494 0880

Email: ace@ecn.purdue.edu

Abstract

In this paper we present a novel multiresolution scheme for the detection of spiculated lesions in digital mammograms. First, a multiresolution representation of the original mammogram is obtained using a linear phase nonseparable 2-D wavelet transform. A set of features is then extracted at each resolution in the wavelet pyramid for every pixel. This approach addresses the difficulty of predetermining the neighborhood size for feature extraction to characterize objects that may appear in different sizes. Detection is performed from the coarsest resolution to the finest resolution using a binary tree classifier. This top-down approach requires less computation by starting with the least amount of data and propagating detection results to finer resolutions. Experimental results using the MIAS image database have shown that this algorithm is capable of detecting spiculated lesions of very different sizes at low false positive rates.

Index Terms: digital mammogram, spiculated lesion, multiresolution, feature analysis, binary classification tree.

EDICS: IP 1.6

This work was supported by a grant from the National Institutes of Health, Grant Number CA62243, and a fellowship from the Purdue Cancer Center. Address all correspondence to E. J. Delp, ace@ecn.purdue.edu, <http://www.ece.purdue.edu/~ace>, or +1 765 494 1740.

1 Introduction

Breast cancer is the most common form of cancer in American women and the second major cause of death after lung cancer [1]. It is a serious health problem in the United States the incidence of which continues to rise [2]. Mammography is currently the best technique for reliable detection of early, non-palpable, potentially curable breast cancer [3]. The mortality rate from this disease decreased for the first time in 1995, due in part to the increasing utilization of screening mammography [2]. However, radiologists vary in their interpretation of mammograms. In addition, the interpretation is a repetitive task that requires much attention to minute detail. Therefore, in the past several years there has been great interest in the use of image processing and analysis techniques [4, 5, 6, 7, 8, 9] for automated tumor detection in digital mammograms. The goal has been to increase diagnostic accuracy as well as the reproducibility of mammographic interpretation.

Among breast abnormalities, spiculated masses having a stellate appearance in mammograms are highly suspicious indicators of breast cancer [1]. Their central masses are usually irregular with ill-defined borders. Their sizes vary from a few millimeters to several centimeters in diameter. Computer aided diagnosis of digital mammograms generally consists of feature extraction followed by classification [4]. For spiculated lesions, Kegelmeyer, *et al* [8, 10] extracted a five-dimensional feature vector for each pixel which included the standard deviation of the edge orientation histogram (ALOE) and the output of four spatial filters. Each feature vector was then classified using a binary decision tree. They chose the window size for obtaining the ALOE feature to be 4 cm so that it was large enough to encompass all of the spiculated lesions in their data. Karssemeijer and Brake [11] investigated a method based on statistical analysis of a map of edge orientations. Two features were derived from the edge orientation map that were used to classify suspicious regions. To

cover the range of sizes of the spicules in their data set, edge orientations were derived at 3 spatial scales and the one with the maximum absolute value was used. In general, it has proved difficult to determine the size of the neighborhood that should be used to extract local features of spiculated lesions. If the neighborhood is too large, small lesions may be missed; while if the neighborhood is too small, parts of large lesions may be missed.

In this paper, we present a multiresolution scheme for the detection of spiculated lesions in mammograms that specifically addresses this difficulty of predetermining the neighborhood size for feature extraction. The system diagram is shown in Figure 1. We first obtain a multiresolution representation of the original mammogram using a linear phase nonseparable 2-D wavelet transform. Features that are designed to differentiate spiculated lesions from normal structures are then extracted for every pixel at each resolution. Detection is performed from the coarsest resolution to the finest resolution using a binary tree classifier. There is a fundamental difference in our method than other approaches, which is also the novelty of our algorithm, in that we extract and classify features at multiple resolutions, hence overcoming the difficulty of choosing a neighborhood size *a priori* to capture tumors of varying sizes. Furthermore, the top-down classification we use requires less computation by starting with the coarsest resolution image (least amount of data) and propagating detection results to finer resolutions.

In Section 2, we describe the multiresolution decomposition method. In Section 3, we discuss the mammographic features used to characterize spiculated lesions of different sizes at various resolutions. The top-down classification approach is discussed in Section 4. Finally we present experimental results and discussion in Section 5.

2 Multiresolution Decomposition

As indicated above, the spiculated lesions we are interested in can occur in very different sizes. Hence, it is impossible to define *a priori* an optimal resolution for analyzing a mammogram. A multiresolution representation reorganizes image information into a set of details appearing at different spatial resolutions, and thus provides a hierarchical framework for image analysis [12]. Recently the wavelet transform has become a popular technique for multiresolution representation and analysis in a wide variety of image processing applications, including computer aided diagnosis of mammograms [13, 14, 15, 16, 17].

We choose the linear phase nonseparable 2-D perfect reconstruction wavelet transform described in [18] to obtain a multiresolution representation [12] of the original mammogram. This transform does not introduce phase distortions in the decomposed images. In addition, no bias is introduced in the horizontal and vertical directions as would occur with a separable transform. The impulse response of the analysis low pass filter is

$$h(n_1, n_2) = \begin{pmatrix} 0 & 0.125 & 0 \\ 0.125 & 0.5 & 0.125 \\ 0 & 0.125 & 0 \end{pmatrix}$$

The dilation matrix used to represent the subsampling lattice is

$$D = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$$

which corresponds to the 2-D quincunx sublattice [18], as shown in Figure 2.

Let the original mammogram have the finest resolution $N \times N$ pixels. Since D expands the sampling lattice by $\sqrt{2}$ in each direction, image resolution decreases by a factor of

$1/\sqrt{2}$ after each decomposition. For example, the image at the second finest resolution has $N/\sqrt{2} \times N/\sqrt{2}$ pixels. We will use only images of spatial resolutions $(N \times N), (N/2 \times N/2), \dots$ for feature extraction and classification.

3 Multiresolution Feature Analysis

Spiculated lesions range in size from a few millimeters to several centimeters. Usually the larger the tumor center, the longer are the spicules or “arms” [19]. Hence a large lesion preserves a stellate appearance at several more coarse resolutions. This can be seen in Figure 3, in which a spiculated lesion of 1 cm in radius is shown at multiple resolutions.

Also note that an $M \times M$ region at a coarser spatial resolution $N/n \times N/n$ corresponds to an $nM \times nM$ region in the original mammogram with spatial resolution $N \times N$. For example, if a set of features extracted within an 8×8 window in the original mammogram with spatial resolution $N \times N$ differentiates spiculated lesions of size ≈ 1 mm from normal tissue, then the same set of features extracted at the coarser resolution $N/4 \times N/4$, using the same sized 8×8 window, would be able to detect similar spiculated lesions of size ≈ 4 mm. This enables us to choose a fairly small neighborhood for feature extraction at the finest resolution and detect the smallest possible spiculated lesions. Larger lesions will be detected using the same set of features extracted at a coarser resolution. This strategy circumvents the difficulty of choosing a neighborhood size *a priori* to capture features of objects of varying sizes.

We extract four features from the low frequency wavelet coefficients for every pixel at each resolution. Let (i, j) be the spatial location in the mammogram at row i and column j ; $f(i, j)$ be the pixel brightness at (i, j) ; ∂S_{ij} be some neighborhood of (i, j) , and K be the number of pixels within ∂S_{ij} . The first two features are then defined as the following:

- mean pixel brightness in ∂S_{ij} :

$$\bar{f}(i, j) = \frac{1}{K} \sum_{(m,n) \in \partial S_{ij}} f(m, n)$$

- standard deviation of pixel brightnesses in ∂S_{ij} :

$$\sigma_f(i, j) = \sqrt{\frac{1}{K-1} \sum_{(m,n) \in \partial S_{ij}} (f(m, n) - \bar{f}(i, j))^2}$$

Let $D_y(i, j)$ and $D_x(i, j)$ be an estimate of the vertical and horizontal spatial derivatives of f at (i, j) , respectively. Let $\theta(i, j) = \tan^{-1}\{D_y(i, j)/D_x(i, j)\}$ be an estimate of the gradient orientation at (i, j) with value $\in (-\pi/2, \pi/2]$. Also let $hist_{ij}$ be the histogram of θ within ∂S_{ij} using 256 bins, hence $hist_{ij}(n) = \#$ of pixels in ∂S_{ij} that have gradient orientations $\in (-\pi/2+n\pi/256, -\pi/2+(n+1)\pi/256]$, where $n = 0, 1, \dots, 255$. Let $\overline{hist}(i, j) = \frac{1}{256} \sum_{n=0}^{255} hist_{ij}(n)$ be the average bin height of $hist_{ij}$.

The third feature is defined as:

- standard deviation of gradient orientation histogram in ∂S_{ij} :

$$\sigma_{hist}(i, j) = \sqrt{\frac{1}{255} \sum_{n=0}^{255} (hist_{ij}(n) - \overline{hist}(i, j))^2}$$

Note that this feature is the same as the ALOE feature described in [8].

Let $K_{\theta_+}(i, j)$ and $K_{\theta_-}(i, j)$ be the number of positive and negative gradient orientations within ∂S_{ij} , respectively; $\bar{\theta}_+(i, j) = \frac{1}{K_{\theta_+}(i, j)} \sum_{\theta(m,n) \geq 0, (m,n) \in \partial S_{ij}} \theta(m, n)$ and $\bar{\theta}_-(i, j) = \frac{1}{K_{\theta_-}(i, j)} \sum_{\theta(m,n) < 0, (m,n) \in \partial S_{ij}} \theta(m, n)$ be the average positive and negative gradient orientations within ∂S_{ij} , respectively.

The folded gradient orientation $\theta'(i, j)$ at (i, j) is defined as:

$$\theta'(i, j) = \begin{cases} \theta(i, j) + \pi & \text{if } \bar{\theta}_+(i, j) - \theta(i, j) > \frac{\pi}{2} \text{ and } K_{\theta_+}(i, j) \geq K_{\theta_-}(i, j) \\ \theta(i, j) - \pi & \text{if } \theta(i, j) - \bar{\theta}_-(i, j) > \frac{\pi}{2} \text{ and } K_{\theta_+}(i, j) < K_{\theta_-}(i, j) \\ \theta(i, j) & \text{otherwise} \end{cases}$$

And the fourth feature is then defined as:

- standard deviation of the folded gradient orientations in ∂S_{ij} :

$$\sigma_{\theta'}(i, j) = \sqrt{\frac{1}{K-1} \sum_{(m,n) \in \partial S_{ij}} (\theta'(m, n) - \bar{\theta}'(i, j))^2}$$

This modification of θ to θ' in generating $\sigma_{\theta'}$ is necessary for this feature not to be sensitive to the nominal value of θ , but to the actual gradient orientation variances. As one can see from Figure 4, the gradient orientation distance between $\pi/2$ and $-\pi/4$ is the same as that between $\pi/2$ and $\pi/4$. For our purpose of detecting spicules, we would like to obtain similar feature values for these two cases. However, if the θ s are used directly, a much larger variance for the former pair would result (2.8 vs 0.3); while using the folded gradient orientations θ' s, $-\pi/4$ will be folded to $3\pi/4$ and so we obtain the same $\sigma_{\theta'}$ for both cases. This feature is different than a similar one we presented in [14] and has proved to better characterize spiculated lesions. More information describing the feature set can be found in [20]¹.

The features σ_{hist} and $\sigma_{\theta'}$ are chosen to be especially responsive to spiculated lesions. As illustrated in Figure 5, a spiculated lesion differs from the normal ducts and connective tissue septae of the breast that tend to radiate from the nipple to the chest wall. It appears as an irregular tumor center surrounded by spicules that radiate in all directions.

¹This reference is available at <http://www.ece.purdue.edu/~ace>

Accordingly, pixels in normal areas have rather similar gradient orientations, while pixels near spiculated lesions tend to have gradient orientations in many directions. Therefore, the standard deviation of gradient orientations in the neighborhood of a lesion pixel will be larger than that in the neighborhood of a normal pixel. For the same reason, the gradient orientation histogram will be flat near a lesion pixel, but will have a peak corresponding to normal structure orientations near a normal pixel [8]. This result is shown in Figure 6 which compares the gradient orientation histograms within a spiculated lesion region and within a normal region, respectively. The mean \bar{f} and standard deviation σ_f of the pixel brightnesses help further to differentiate lesions from normal tissue regions, because lesions usually have higher density, or appear brighter, in mammograms.

Figure 7 shows a test pattern at multiple resolutions, which consists of an idealized spiculated lesion and parallel linear markings, embedded in uncorrelated Gaussian distributed noise with PSNR (Peak Signal to Noise Ratio) of 7db. The lesion has a radius of 128 pixels in the original image. Figure 8, Figure 9, Figure 10, and Figure 11 show the features $\sigma_{\theta'}$, σ_{hist} , \bar{f} , and σ_f , respectively, for this test pattern. Each feature at different resolutions is extracted within same sized circular neighborhoods. For this test pattern, the neighborhood radius is 30 pixels for features $\sigma_{\theta'}$ and σ_{hist} , and 20 pixels for features \bar{f} and σ_f . A larger neighborhood size is used for features $\sigma_{\theta'}$ and σ_{hist} because they respond stronger to a spiculated lesion if the entire halo of spicules is included in the feature extraction window, while features \bar{f} and σ_f respond better to the central mass of the lesion. These features are able to differentiate a spiculated lesion from a complex background. From Figure 8 and Figure 9, we can see that the features best discriminate the lesion from its background at resolution $(N/4 \times N/4)$ when the size of the feature extraction neighborhood matches that of the lesion. In this case, the brightest spot in the feature image in Figure 8 (c) and the darkest spot in the feature image in Figure 9 (c) directly overlay both the central mass and the spicules of

the lesion.

4 Classification Algorithm

A sequential hierarchical decision scheme has been shown to achieve better performance than employing a single “best” set of features in a one-step decision [21, 22]. A Binary Classification Tree (BCT) is simple, fast, and efficient type of hierarchical classifier. Figure 12 shows a simple BCT with 3 features and 2 classes where circular nodes are binary decision nodes and square nodes are terminal nodes with class labels assigned. When an unknown feature vector is submitted for classification, it will first go to the root node, which is always a binary decision node, and then take one of the two branches based on the outcome of testing one of its features against the threshold at that binary decision node. This process continues until the feature vector reaches one of the terminal nodes where it is assigned a class label. This tree structured classification approach has several advantages over more traditional nonparametric methods such as the nearest neighbor method [21]: Firstly, BCT does automatic stepwise feature selection and complexity reduction; Secondly, BCT is robust with respect to outliers and misclassified points in the training set; Thirdly, the final classifier can be compactly stored; Fourthly, BCT efficiently classifies new data; And finally, BCT provides easily understood and interpreted information regarding the predictive structure of the data. We choose the iterative growing and pruning algorithm proposed in [22] for our classification tree design because it not only produces trees with higher classification accuracy, but also requires less computation than other widely used tree design algorithms, such as CART [21].

Considering that there is redundancy in mapping the feature space by spatially adjacent samples, only a subsampling of the training set is used to grow BCT. After the BCT is generated, we associate with each terminal node a suspicious probability which is the percentage

of lesion pixels in the training images that falls in this terminal node. The suspicious probability is then recomputed using the entire set of training samples. This partly compensates for the lost classification accuracy due to the fact that not all the training samples are used to design the tree [8].

At each resolution, five features are used for classification: the four features extracted at that resolution plus the feature $\sigma_{hist_{ij}}$ extracted at the next coarser resolution. Our experiments have shown that using features across resolutions simultaneously helps capture spiculated lesions of varying sizes. Detection starts from the second coarsest resolution and then goes to the next finer resolution until the original mammogram is reached. A positive detection at a coarser resolution propagates to finer resolutions, which effectively reduces the number of pixels to be classified. A negative result at a coarser resolution will be combined with those at finer resolutions via a weighted sum.

We obtain a probability image from the BCT for each test mammogram, in which the pixel values represent the probability that a pixel belongs to a spiculated lesion. A median filter is then used to eliminate isolated positive responses, followed by a smoothing filter to reach a consensus within neighboring pixels. Final results are obtained by thresholding the filtered probability image.

5 Experimental Results and Discussions

We tested the proposed multiresolution detection algorithm on the MIAS database [23] provided by the Mammographic Image Analysis Society (MIAS) in the UK. Images in the MIAS database have 50 micron resolution. There is a total of 19 mammograms containing spiculated lesions. The smallest lesion extends 3.6 mm in radius, while the biggest one is nearly 10 times larger and extends to 35.0 mm. These 19 together with another 19 normal mammograms, also taken from the MIAS database, were randomly split into two sets with

approximately an equal number of lesions and normal mammograms in each set. Each set was then used separately as a training set to generate two BCTs, according to a split-half training paradigm similar to Kegelmeyer *et al* [8]. A BCT trained by one set was used to classify mammograms in the other set, and vice versa. Therefore, no mammogram was used both for training and testing.

The wavelet transform described in Section 2 was used to generate the multiresolution representations of each mammogram. Features described in Section 3 were extracted from images of spatial resolutions $(N \times N)$, $(N/2 \times N/2)$, $(N/4 \times N/4)$, and $(N/8 \times N/8)$. Circular regions were used for all feature extractions so that features behave equally in all directions. In this experiment, the neighborhood radius for extracting features $\sigma_{\theta'}$ and σ_{hist} was 60 at all resolutions, which corresponds to a radius of 3 mm, 6 mm, 12 mm, and 24 mm in the original resolution $(N \times N)$, and coarser resolutions $(N/2 \times N/2)$, $(N/4 \times N/4)$, and $(N/8 \times N/8)$, respectively. The neighborhood radius for extracting features \bar{f} and σ_f was 20 at all resolutions, which corresponds to a radius of 1 mm, 2 mm, 4 mm, and 8 mm in the original resolution $(N \times N)$, and coarser resolutions $(N/2 \times N/2)$, $(N/4 \times N/4)$, and $(N/8 \times N/8)$, respectively. Extracting features at coarser resolutions for relatively large regions significantly reduces the amount of computation.

Accordingly, three-level classifications were performed for each mammogram, starting from resolution $(N/4 \times N/4)$ and going to $(N \times N)$. At each resolution, the corresponding five dimensional feature vector, including the four features extracted at that resolution and σ_{hist} extracted at the next coarser resolution, was classified through the BCT trained by the other half test set. The output is a probability image, in which the pixel values represent the probability of belonging to spiculated lesions. We then used a median filter of radius 5 to eliminate isolated positive responses, followed by a smoothing filter to reach a consensus within neighboring pixels. Because we are looking for larger lesions at coarser resolutions

and smaller ones at finer resolutions, the radii used for smoothing filters were chosen to be 18 mm, 6 mm, and 2 mm at resolutions $(N/4 \times N/4)$, $(N/2 \times N/2)$, and $(N \times N)$, respectively. The detection result was obtained by thresholding the smoothed probability image. Using a small threshold is more likely to detect true lesions, but also to generate more false positive responses. Using a large threshold gives fewer false positive responses, but may miss more true lesions. Hence variation of thresholds gives different diagnostic accuracy which can be quantified using FROC (“Free response Receiver Operating Characteristic”) analysis [24], where the true positive fraction (TPF) is plotted as a function of the average number of false positives (FP) per image. FROC analysis [24] is well suited for the assessment of computer aided diagnosis of mammograms because it is applicable to situations that involve any number of reported locations and any number of actual lesions in each image.

If there is a positive detection at a coarser resolution, no feature extraction and detection are needed at the corresponding pixel locations at all finer resolutions. This approach effectively reduces the number of pixels to be classified. Smaller probabilities that give negative responses are propagated to finer resolutions by weighted sum. Figure 13 shows a spiculated lesion of size 6.6 mm detected at the finest original resolution; Figure 14 shows a spiculated lesion of size 12.4 mm detected at the second coarsest resolution; and Figure 15 shows a spiculated lesion of size 35.0 mm detected at the coarsest resolution.

The coarse to fine detection scheme speeds up computation in two ways. First, positive responses at lower resolutions effectively reduces computation at higher resolutions. Secondly, most of the computation is saved by allowing feature extraction within small windows. For example, we only need a 40×40 pixel window for feature extraction in 2 mm, 4 mm, and 8 mm neighborhoods in a 50μ mammogram, because we extract features at different resolutions. If we extracted features in the original image in three spatial scales so as to examine different sizes of possible spiculated lesions, similar to Karssemeijer and Brake [11], the 2

mm, 4 mm, and 8 mm neighborhoods would correspond to 40×40 , 80×80 , and 160×160 pixel window sizes, respectively, which would cost at least 600% more computation.

To evaluate the computer diagnosis results, we adopted the criteria in [8]: a computer finding is considered as a true positive detection if its area is overlapped by at least 50% of a true lesion as indicated by an expert radiologist; a computer finding that does not so overlap a true lesion is considered as false positive; and a true lesion that is not overlapped by any computer finding is considered as false negative. By these criteria, the diagnostic accuracy performance of our algorithm on the MIAS database quantified using FROC curve is shown in Figure 16. We achieved 84.2% true positive detection at less than 1 false positive per image and 100% true positive detection at 2.2 false positive per image. The spiculated lesion that was missed first when the detection threshold was increased, or the sensitivity decreased, is shown in Figure 17. This lesion lacks visible spicules at all directions. Karssemeijer and Brake [11] reported comparable results on the MIAS database — around 90% true positive detection at 1 false positive per image and 100% true positive detection at more than 4 false positive per image. The presented multiresolution algorithm for the detection of spiculated lesions in digital mammograms is capable of detecting spiculated lesions of very different sizes at low false positive rates.

6 Acknowledgments

The authors would like to thank Dr. Valerie P. Jackson, the John A. Campbell Professor of Radiology at Indiana University, and Professor Zygmunt Pizlo of the Department of Psychological Sciences at Purdue University, for their helpful suggestions.

REFERENCES

- [1] D. B. Kopans, *Breast Imaging*. J. B. Lippincott Company, 1989.

- [2] L. W. Bassett, V. P. Jackson, R. Jahan, Y. S. Fu, and R. H. Gold, *Diagnosis of Diseases of the Breast*. W. B. Saunders Company, 1997.
- [3] G. Cardenosa, "Mammography: An overview," *Proceedings of the 3rd International Workshop on Digital Mammography*, June 9–12 1996, Chicago, Illinois, pp. 3–10.
- [4] K. Woods and K. Bowyer, "A general view of detection algorithms," *Proceedings of the 3rd International Workshop on Digital Mammography*, June 9–12 1996, Chicago, U.S.A., pp. 385–390.
- [5] M. L. Giger, F.-F. Yin, K. Doi, C. E. Metz, R. A. Schmidt, and C. J. Vyborny, "Investigation of methods for the computerized detection and analysis of mammographic masses," *Proceedings of SPIE*, February 6 - 8 1990, Bellingham, Washington, pp. 183–184.
- [6] N. Karssemeijer, "A stochastic model for automated detection of calcifications in digital mammograms," *12th International Conference on Information Processing in Medical Imaging*, 1991, Wye, UK, pp. 227–238.
- [7] C. J. Vyborny and M. L. Giger, "Computer vision and artificial intelligence in mammography," *American Journal of Roentgenology*, vol. 162, no. 3, pp. 699–708, March 1994.
- [8] W. P. Kegelmeyer, Jr., J. M. Pruneda, P. D. Bourland, A. Hillis, M. W. Riggs, and M. L. Nipper, "Computer-aided mammographic screening for spiculated lesions," *Radiology*, vol. 191, no. 2, pp. 331–336, May 1994.
- [9] M. L. Comer, S. Liu, and E. J. Delp, "Statistical segmentation of mammograms," *Proceedings of the 3rd International Workshop on Digital Mammography*, June 9–12 1996, Chicago, Illinois, pp. 475–478.
- [10] W. P. Kegelmeyer, Jr., "Computer detection of stellate lesions in mammograms," *Proceedings of the SPIE Conference on Biomedical Image Processing and Three-Dimensional Microscopy*, February 10-13 1992, San Jose, California, pp. 446–454.
- [11] N. Karssemeijer and G. M. Brake, "Detection of stellate distortions in mammograms," *IEEE Transactions on Medical Imaging*, vol. 15, no. 5, pp. 611–619, October 1996.
- [12] S. G. Mallat, "A theory for multiresolution signal decomposition: The wavelet representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 7, pp. 674–693, July 1989.
- [13] M. Unser and A. Aldroubi, "A review of wavelets in biomedical applications," *Proceedings of the IEEE*, vol. 84, no. 4, pp. 626–638, April 1996.
- [14] S. Liu and E. J. Delp, "Multiresolution detection of stellate lesions in mammograms," *Proceedings of the IEEE International Conference on Image Processing*, October 26–29 1997, Santa Barbara, California, pp. 109–112.

- [15] R. N. Strickland and H. I. Hahn, "Wavelet transforms for detecting microcalcifications in mammograms," *IEEE Transactions on Medical Imaging*, vol. 15, no. 2, pp. 218–229, April 1996.
- [16] A. F. Laine, S. Schuler, J. Fan, and W. Huda, "Mammographic feature enhancement by multiscale analysis," *IEEE Transactions on Medical Imaging*, vol. 13, no. 4, pp. 725–752, December 1994.
- [17] C. H. Chen and G. G. Lee, "Image segmentation using multiresolution wavelet analysis and expectation-maximization (EM) algorithm for digital mammography," *International Journal of Imaging Systems and Technology*, vol. 8, no. 5, pp. 491–504, 1997.
- [18] J. Kovačević and M. Vetterli, "Nonseparable multidimensional perfect reconstruction filter banks and wavelet bases for \mathbb{R}^n ," *IEEE Transactions on Information Theory*, vol. 38, no. 2, pp. 533–555, March 1992.
- [19] L. Tabar and P. B. Dean, *Teaching Atlas of Mammography*. Georg Thieme Verlag, 2nd revised ed., 1985.
- [20] S. Liu, *The Analysis of Digital Mammograms: Spiculated Tumor Detection and Normal Mammogram Characterization*. PhD thesis, Purdue University, West Lafayette, Indiana, USA, May 1999.
- [21] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone, *Classification and Regression Trees*. Wadsworth, Inc., 1984.
- [22] S. B. Gelfand, C. S. Ravishankar, and E. J. Delp, "An iterative growing and pruning algorithm for classification tree design," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 2, pp. 163–174, February 1991.
- [23] J. Suckling, J. Parker, D. R. Dance, S. Astley, J. Hutt, C. R. M. Doggis, I. Ricketts, E. Stamatakis, N. Cerneaz, S. L. Kok, P. Taylor, D. Betal, and J. Savage, "The mammographic image analysis society digital mammogram database," *Proceedings of the 2nd International Workshop on Digital Mammography*, July 10–12 1994, York, England, pp. 375–378.
- [24] C. E. Metz, "Evaluation of digital mammography by ROC analysis," *Proceedings of the 3rd International Workshop on Digital Mammography*, June 9–12 1996, Chicago, Illinois, pp. 61–68.

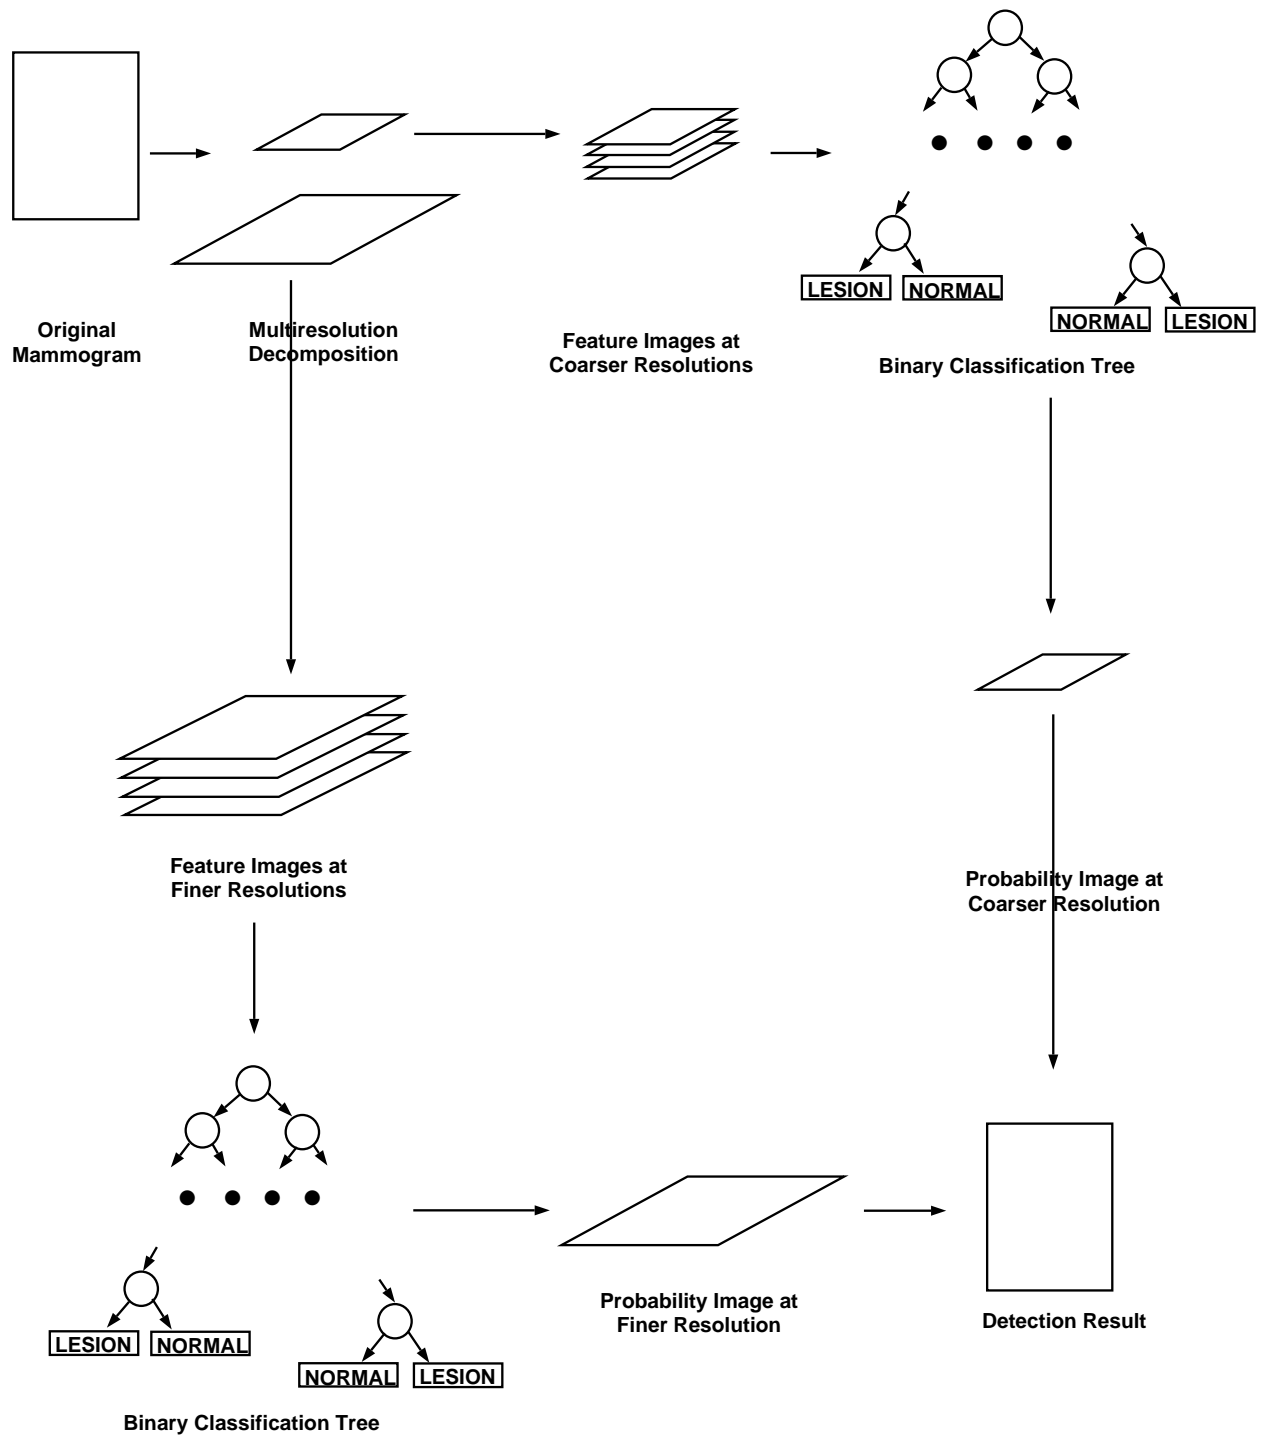


Figure 1: Diagram of system for multiresolution spiculated lesion detection, illustrated here for only two resolutions — fine (left) and coarse (right).






Figure 2: Two dimensional quincunx sublattice, used to create successive generations of lower resolution images for analysis.



(a)



(b)

(c)

(d)

Figure 3: Multiresolution representation of a spiculated lesion (a) spiculated lesion as in the original mammogram of resolution $N \times N$. (b) in the coarser resolution $N/2 \times N/2$. (c) in the coarser resolution $N/4 \times N/4$. (d) in the coarsest resolution $N/8 \times N/8$




Figure 4: Folding edge orientation makes the feature $\sigma_{\theta'}$ insensitive to the nominal value of θ , but sensitive to the actual edge orientation variances.

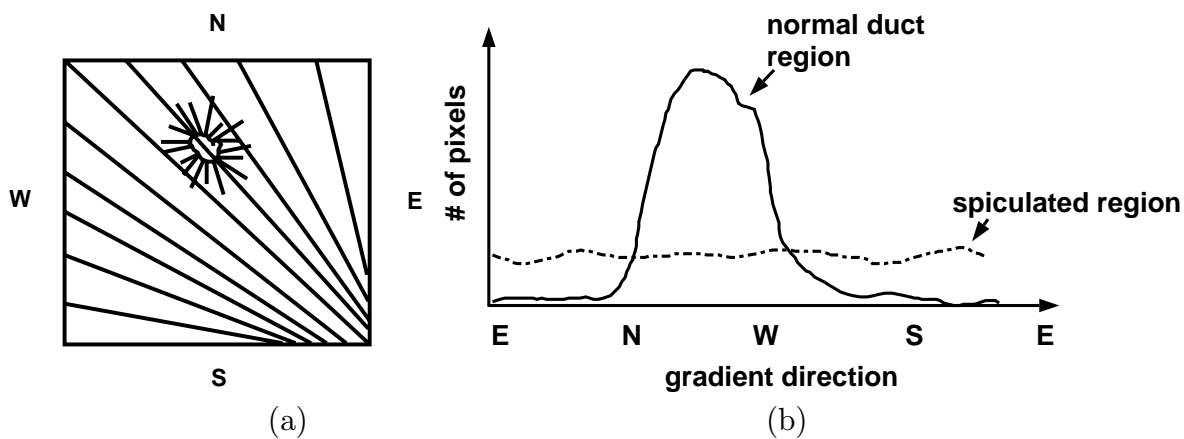




Figure 5: (a) Directions of spicules of a spiculated lesion differ from the directions of normal linear markings in a mammogram; (b) standard deviation of gradient orientation histogram differentiates the area near a spiculated lesion from normal.




(a)

(b)




(c)


Figure 6: (a) A spiculated lesion region. (b) A normal region. (c) Gradient orientation histograms of (a) and (b)



(a)




(b)




(c)


Figure 7: Multiresolution representation of a test pattern consisting of an idealized spiculated lesion and normal structures embedded in uncorrelated Gaussian distributed noise. The lesion has a radius of 128 pixels in the original image: (a) original $N \times N$ image. (b) coarser resolution $N/2 \times N/2$. (c) coarsest resolution $N/4 \times N/4$.



(a)



(b)




(c)




(d)


Figure 8: Feature $\sigma_{\theta'}$ obtained in circular neighborhoods of radius 30 pixels at all resolutions for the test pattern shown in Figure 7. Here large values (bright pixels) indicate spiculated regions as explained in Figure 5. (a) original $N \times N$ image. (b) coarser resolution $N/2 \times N/2$. (c) coarsest resolution $N/4 \times N/4$. (d) a circular neighborhood of size 30 pixels in radius.



(a)



(b)




(c)




(d)

Figure 9: Feature σ_{hist} obtained in circular neighborhoods of radius 30 pixels at all resolutions for the test pattern shown in Figure 7. Here small values (dark pixels) indicate spiculated regions as explained in Figure 5. (a) original $N \times N$ image. (b) coarser resolution $N/2 \times N/2$. (c) coarsest resolution $N/4 \times N/4$. (d) a circular neighborhood of size 30 pixels in radius.



(a)



(b)




(c)




(d)

Figure 10: Feature \bar{f} obtained in circular neighborhoods of radius 20 pixels at all resolutions for the test pattern shown in Figure 7: (a) original $N \times N$ image. (b) coarser resolution $N/2 \times N/2$. (c) coarsest resolution $N/4 \times N/4$. (d) a circular neighborhood of size 20 pixels in radius.



(a)



(b)



(c)



(d)

Figure 11: Feature σ_f obtained in circular neighborhoods of radius 20 pixels at all resolutions for the test pattern shown in Figure 7. Note that feature values get smaller (darker) in the background area at coarser resolutions because of the smoothing effect of the lowpass filtering step in obtaining coarser resolution images. (a) original $N \times N$ image. (b) coarser resolution $N/2 \times N/2$. (c) coarsest resolution $N/4 \times N/4$. (d) a circular neighborhood of size 20 pixels in radius.

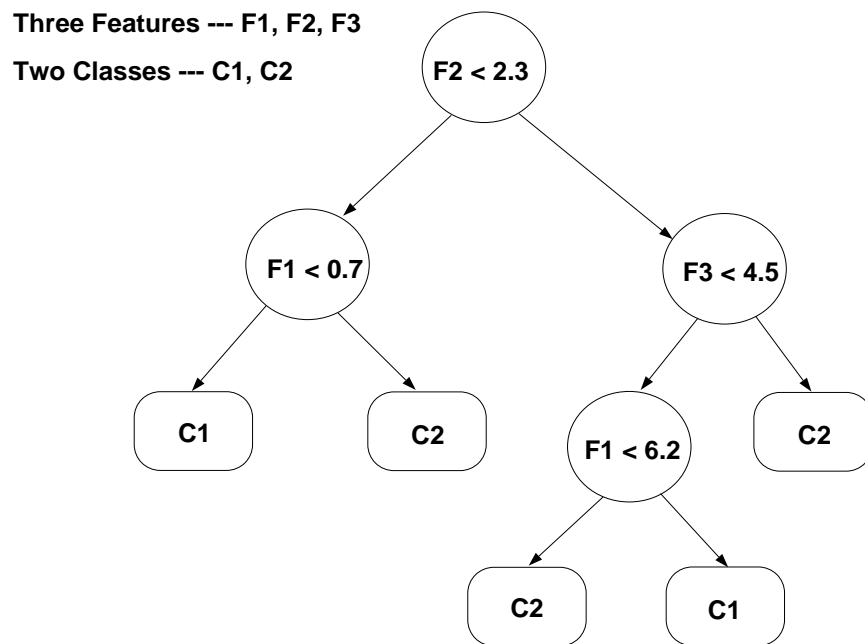


Figure 12: A simple binary classification tree with 3 features and 2 classes. The circular nodes are binary decision nodes and the square nodes are terminal nodes with class labels assigned




Figure 13: A spiculated lesion of size 6.6 mm detected at the finest resolution (a) mammogram with automatic detection. (b) mammogram with tumor as determined by expert radiologists.





Figure 14: A spiculated lesion of size 12.4 mm detected at the second coarsest resolution (a) mammogram with automatic detection. (b) mammogram with tumor as determined by expert radiologists.



(a)



(b)

Figure 15: A spiculated lesion of size 35.0 mm detected at the coarsest resolution (a) mammogram with automatic detection. (b) mammogram with tumor as determined by expert radiologists.




Figure 16: FROC curve for detection of spiculated lesions in the MIAS database




Figure 17: The first missed spiculated lesion when the detection sensitivity is decreased. It lacks visible spicules at all directions.