Pergamon

Pattern Recognition, Vol. 30, No. 9, pp. 1547–1554, 1997
Pattern Recognition Society
Published by Elsevier Science Ltd
Printed in Great Britain

PII: S0031-3203(96)00173-2

AN ANALYSIS OF LOCAL FEATURE EXTRACTION IN DIGITAL MAMMOGRAPHY

DAVID J. MARCHETTE,[†,*] RICHARD A. LOREY[†] and CAREY E. PRIEBE[‡]

[†]Naval Surface Warfare Center, B10, Dahlgren, VA 22448, U.S.A.
[‡]Johns Hopkins University, Department of Mathematical Sciences, Baltimore, MD 21218, U.S.A.

(*Received* 16 *June* 1995; *in revised form* 24 *May* 1996; *received for publication* 24 *October* 1996)

Abstract—A fundamental problem of automating the detection and recognition of abnormalities in digital mammograms utilizing computational statistics is one of extracting the appropriate features for use in a classification system. Several feature sets have been proposed although none have been shown to be sufficient for the problem. Many of these features tend to be local in nature, which means their calculation requires a connected region of the image over which an average or other statistic is extracted. The implicit assumption is that the region is homogeneous, but this is rarely the case if a fixed window is used for the calculation. We consider a method of using boundaries to segment the window into more homogeneous regions for use in the feature extraction calculation. This approach is applied to the problem of discriminating between tumor and healthy tissue in digital mammography. A set of 21 images, each containing a biopsied mass, is described. The results of the boundary-gated feature extraction methodology on this image set shows a difference in distribution between tissue interior to the mass and tissue far away from the mass. Less difference is discernible when boundaries are not used in the feature extraction. Published by Elsevier Science Ltd.

Image processing Feature extraction Local features Boundary gating

1. INTRODUCTION

The use of computational statistics for general image analysis and for application to digital mammography in particular has been described elsewhere in some detail.[1–4] These papers address the fundamental problem of the extraction of appropriate features for use in a classification system. Specifically, a set of features based on the fractal nature of an image has been used with some limited success based on a very small number of mammographic cases.[5–7] Other local features have also been used[8,9] and again have had only marginal success. Our goal in this work is to describe a method of feature extraction which improves the local features by attempting to ensure homogeneity of the region from which the feature is extracted.

A local feature, for the purpose of this work, is a value calculated from a window centered on a given pixel. Thus, each pixel can be assigned a feature value, even though the feature requires a region of pixels for its calculation.

We will give three examples of local features and briefly discuss two related methods for improving the feature extraction. The first feature is statistical in nature, the coefficient of variation. Although this is a simplistic feature we show that it does have some utility as a discriminator between healthy tissue and tumorous tissue.

We then consider fractal dimension and investigate two features that result from the fractal calculation. Our work shows that the fractal dimension and a related feature obtained from digitized gray scale mammograms are appropriate features for extraction and might be used as one or two of several features for a classification system for detecting abnormalities.

The problem with computing local features, which are defined in a region of the image, is that care must be taken to insure that the region is homogeneous. If the region used in the fractal calculation contains more than a single texture then the fractal dimension calculation yields a feature which is neither the fractal dimension for either texture nor the average of the fractal dimensions for the two textures.[10] Hence, this is a less useful feature for discrimination. This is particularly troublesome near the boundary of a tumor, a region which may be critical for the final determination of malignancy. It is the purpose of this work to show that this is one of the reasons that the fractal dimension has received mixed reviews in the literature. To ameliorate this situation we explore the notion of boundary gating which leads to a purer calculation of local features.

In order to evaluate the utility of the features for classification, we construct the probability density functions for the features associated with each class. These can be used for classification by the standard likelihood ratio test and so can give a good measure of the utility of each individual feature. Experience has shown that these densities are not well modeled by any of the more common families of densities, and so a nonparametric approach is taken. This is also the reason for using density estimation techniques rather than linear or quadratic classifiers.

* Author to whom correspondence should be addressed.

Although we are considering each feature separately, we acknowledge that no single feature is going to perform well enough on this problem to be clinically useful. In fact, the three features considered here, even if taken together, are not sufficient to adequately detect tumors in these images. Our purpose is not to solve the detection problem but rather to indicate methods of feature extraction which produce better features and improve the overall detection. This will be critical in improving classification systems based on a large number of features. Our results show that with the improved extraction process incorporating boundary gating, even these simple features do have some utility in detecting tumors.

2. EXPERIMENT

2.1. *Description of the data*

Of a set of 50 mammograms (provided by the H. Lee Moffit Cancer Center and Research Institute and the Department of Radiology of the University of South Florida) with a biopsied proven malignancy, we chose a set of 21 for this study. Since one of our goals was to determine if the features could be useful in a detection system for tumorous masses, we did not consider other kinds of abnormalities. Hence, we eliminated any mammogram that contained any microcalcifications and/or architectural distortion outside of the mass so that we could be reasonably sure that the local feature calculations were based solely on tumorous tissue characteristics or solely on healthy tissue characteristics.

For the study, each mammogram was supplied with a radiologist-drawn boundary of the tumorous mass. This boundary was used by us to determine the location of the mass and to determine the interior pixels. The boundary was not used in the boundary-gated calculations described below. Figure 1 shows a typical mammogram with the radiologist-drawn boundary overlaid.

In estimating the probability density function (pdf) for the healthy tissue, we avoid using any pixels that are close to the radiologist's boundary. We do this to avoid any changes in texture that may be associated with the edge of the tumor. To estimate the pdf of the tumor we use the radiologist's boundary to determine which pixels are interior to the tumor.

The method of estimation of the pdf used in this study is the filtered kernel estimator.[11] This is a nonparametric method which uses local variance information to determine a small set of smoothing parameters. First, a mixture of normals is fit to the data to estimate the variances of the modes and tails of the distribution. This mixture of normals is used as a pilot estimator for a modified kernel estimator which allows different smoothing parameters in the different regions of the support of the density.

2.2. *Boundary gating*

To calculate a local feature within a given window it is necessary to determine which pixels within that window should be used in the calculation. Figure 2 gives an example of what is meant here. This figure shows a window in which two distinct textures are evident. In this case the image is of bark rather than a mammogram, since in this case it better illustrates the issue. The center pixel clearly should not use pixels from the right texture in calculating its local feature since these are from a different texture. If we used a boundary map to separate the two regions we could determine which pixels are of the same type as the center pixel and extract a feature which better represents the texture surrounding the center pixel. This is the idea behind boundary gating. If the boundaries are always closed it is a simple problem to determine which pixels to use. However, there are few reliable methods of constructing closed boundaries

Fig. 1. Mammogram (histogram equalized for display purposes) with radiologist-drawn boundary overlaid.

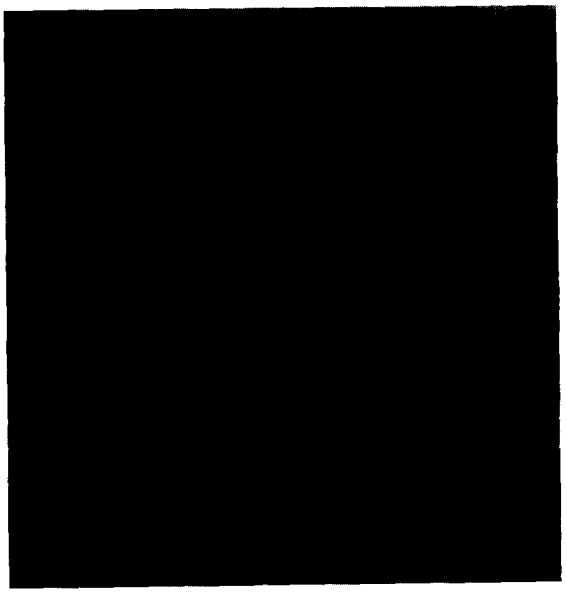

Fig. 2. Window (box) with center indicated by solid circle.

which are computationally tractable. We consider (below) two methods of boundary gating which will work reasonably well with any boundary map.

2.3. *Nonorthogonal wavelet edges*

The desire to include boundary information in the calculation of local features presents a problem when dealing with images such as mammograms, since boundaries are not obvious. This is especially vexing when attempting to generate features for a classification system to differentiate breast tissue types (healthy and tumorous). However, it is possible to generate continuous valued boundaries using wavelets.[12–14] While these boundaries do not always define explicitly the tumorous tissue boundary, it has been shown that using such boundaries results in a pdf which is similar to a pdf generated using a radiologist's determined tumor boundary.[10] For the purposes of this work we are interested in binary boundaries and so we threshold the gray scale boundaries. An example of the boundaries generated using wavelets is shown in Fig. 3.

The wavelet transform is applied to the image and the maxima of the wavelet coefficients are extracted at each scale. These maxima provide an edge map for each scale. For this work we choose to extract edges at four scales and use the second scale edges which are thresholded to produce a binary edge map. This seems to do well on these images and as will be seen below is sufficient to produce improved features.

2.4. *Coefficient of variation*

The simplest and most familiar local features are the gray scale mean and variance of a window on the image. The problem with the mean is that it needs to be normalized across images because images that are darker overall will have lower mean values. Similarly, in general, small differences in intensity are significant to the variance if the overall intensity is low, while these same differences are not significant if the overall intensity is high. To alleviate these problems, one sometimes uses the coefficient of variation. The coefficient of variation (CoV) is defined to be the standard deviation divided by the mean:

$$\mathrm{CoV} = \frac{\sigma}{\mu}. \tag{1}$$

This variable tries to treat the variance in high- and low-intensity regions equally and thus should be fairly robust to changes in overall image intensity.

Fig. 3. Wavelet boundaries.

2.5. *Simple boundary gating*

The simplest form of boundary gating uses what we call the light source algorithm. Imagine a light source at the center pixel and think of the boundaries as opaque walls. The pixels used in the calculation are those which are not shadowed by the boundaries. This is a very crude method which will omit pixels which are merely hidden from the center by a corner or singleton edge pixel. However, it will not select any pixels which require crossing an edge and it is extremely simple to implement. This is the algorithm which will be used in the coefficient of variation example.

2.6. *Fractal features*

The first fractal feature with which we are concerned is the fractal dimension of a gray scale image. To calculate the fractal dimension at a pixel we consider the gray scale to represent a third dimension in the image. We then use the covering method of Peli[15] to estimate the surface area in the region about the pixel. The idea is to bound the surface above and below by erosion and dilation operators at a range of scales and use these to estimate the surface area for the different scales. A linear regression is then performed relating the log of the surface area with the log of the scale. The slope of the regression line is then the difference between the normal Euclidean dimension (d) and the fractal dimension (D). Details of this can be found in references (15–18). This is shown for a general measured property (M) in Richardson's Power Law:[19]

$$M(\varepsilon) = K\varepsilon^{d-D}, \tag{2}$$

where $M(\varepsilon)$ is the measured property of a fractal at a scale ε and K is a proportionality constant.

Using the technique described by Solka *et al.*,[4] we extract the local fractal dimension for two classes of tissue, healthy and tumorous. We describe below our more complicated boundary-gating technique for taking account of boundaries between textures. After calculating the fractal dimension associated with each pixel of each class (healthy and tumorous), we estimate the class pdf. Plots of the class pdfs indicate whether a fractal dimension based texture is a likely candidate for inclusion as an appropriate feature in a classification system.

The second fractal feature which we use is the y-intercept of the regression line. The y-intercept is related to the image contrast. This feature has been used in

previous studies as one of a set of fractal features for image analysis and classification.[1]

2.7. Boundary gating II

For the application of boundary gating described here, a threshold was used on the wavelet generated boundaries. That is, any boundary gray scale value above the threshold was set to 255. Any value below the threshold was set to zero. Thus, any variation in intensity of the boundary was removed. Even though the boundaries in a pixel's window were effectively solid boundaries they often were not continuous. In this technique, the feature calculations were weighted according to the location and distance of each pixel from the central pixel. For those pixels in a straight line from the central pixel, the weight is one. Boundaries that are continuous in the pixel window are not crossed. For discontinuous boundaries (i.e. for those pixels that can be reached by going around the end of a boundary), the weight is a function of the distance from the central pixel. This technique was adopted here to simplify our calculations. However, boundary gating with continuous valued boundaries is possible and is reported elsewhere.[10,11]

3. PERFORMANCE MEASURES

Boundary gating can produce a visible difference in the densities, as indicated in Fig. 4. While this can be convincing in some cases, in others the difference is not so pronounced. Thus we need objective criteria for determining if an improvement has been obtained. We consider two such criteria in this section.

3.1. Kullback–Leibler and divergence

Although the pdfs obtained can be plotted for visual comparison, Kullback–Leibler and divergence integrals are performed to yield quantitative performance measures.[20] The Kullback–Leibler information is defined as

$$\mathrm{KL}(f_\mathrm{T},f_\mathrm{H}) = \int f_\mathrm{T} \log\left(\frac{f_\mathrm{T}}{f_\mathrm{H}}\right), \tag{3}$$

where f_T and f_H are pdfs (associated with tumorous and healthy tissue, respectively) and the divergence, or Jeffreys–Kullback–Leibler Information,[20] is defined to be

$$\mathrm{div}(f_\mathrm{T},f_\mathrm{H}) = \mathrm{KL}(f_\mathrm{T},f_\mathrm{H}) + \mathrm{KL}(f_\mathrm{H},f_\mathrm{T}). \tag{4}$$

The div statistic gives a measure of overall discriminatory power for distinguishing tumor from healthy tissue.

3.2. Probability of detection

The divergence does not tell us what the trade-off is between detecting the tumor and getting a false alarm. In order to assess this, we need to consider the method of classification taken. Since our classifier is based on the pdfs, we classify a pixel according to which pdf is larger, or more generally, we classify a pixel as tumor if the ratio of tumor to healthy tissue pdf evaluated at the pixel is larger than some threshold,

$$\frac{f_\mathrm{T}(x)}{f_\mathrm{H}(x)} > T. \tag{5}$$

Varying the value of T allows us to set the level of false alarms. Using the pdfs for the different classes we compute the threshold at which a given probability of false alarm (PFA) is obtained. This threshold is then used to compute the probability of detection (PD) at this fixed PFA. This allows us to report the probability of correct classification at four fixed values for the probability of false alarm. These give a quantitative measure of the improvement due to boundary gating.

3.3. Hypothesis testing

Two tests are performed for each statistic to assess whether the boundary gating has improved the performance. Thus, for each statistic we compute the difference between the statistic computed with and without boundary gating. The hypotheses are

H_0: S(boundary gating)=S(no boundary)
H_1: S(boundary gating) > S(no boundary)

for a given statistic S. We consider both a parametric and nonparametric test. The parametric test is the familiar paired t-test[21] which assumes, for each statistic, normality of the differences, S(boundary gating) − S(no boundary).

With only 21 data points it is unclear whether the normality assumption is warranted. Therefore, a nonparametric test, the Wilcoxon test[22] which only assumes symmetry of the differences, is also reported. The Wilcoxon test sorts the absolute value of the differences by rank and then sums the ranks of those differences which were positive. If this value is large we can reject the null hypothesis in favor of the alternative.

As a rule, we feel that the Wilcoxon test is probably the safest of the two tests. However, we report both and allow the readers to draw conclusions for themselves.

4. EXPERIMENT OUTCOME

4.1. Results for CoV

We first consider the coefficient of variation, using the simple boundary gating described above. In all these examples we use the same boundary map as extracted by the wavelet edge detector described above. In Table 1 we show the results of considering the div statistic where we test to see if the divergence has improved (become larger) with the incorporation of boundary gating. Normally, a p value of 0.05 is considered sufficient to reject the null hypothesis. Therefore, Table 1 shows we can safely

Table 1. Coefficient of variation

Test	div p value
Wilcoxon	0.0008
Paired t-test	0.0039

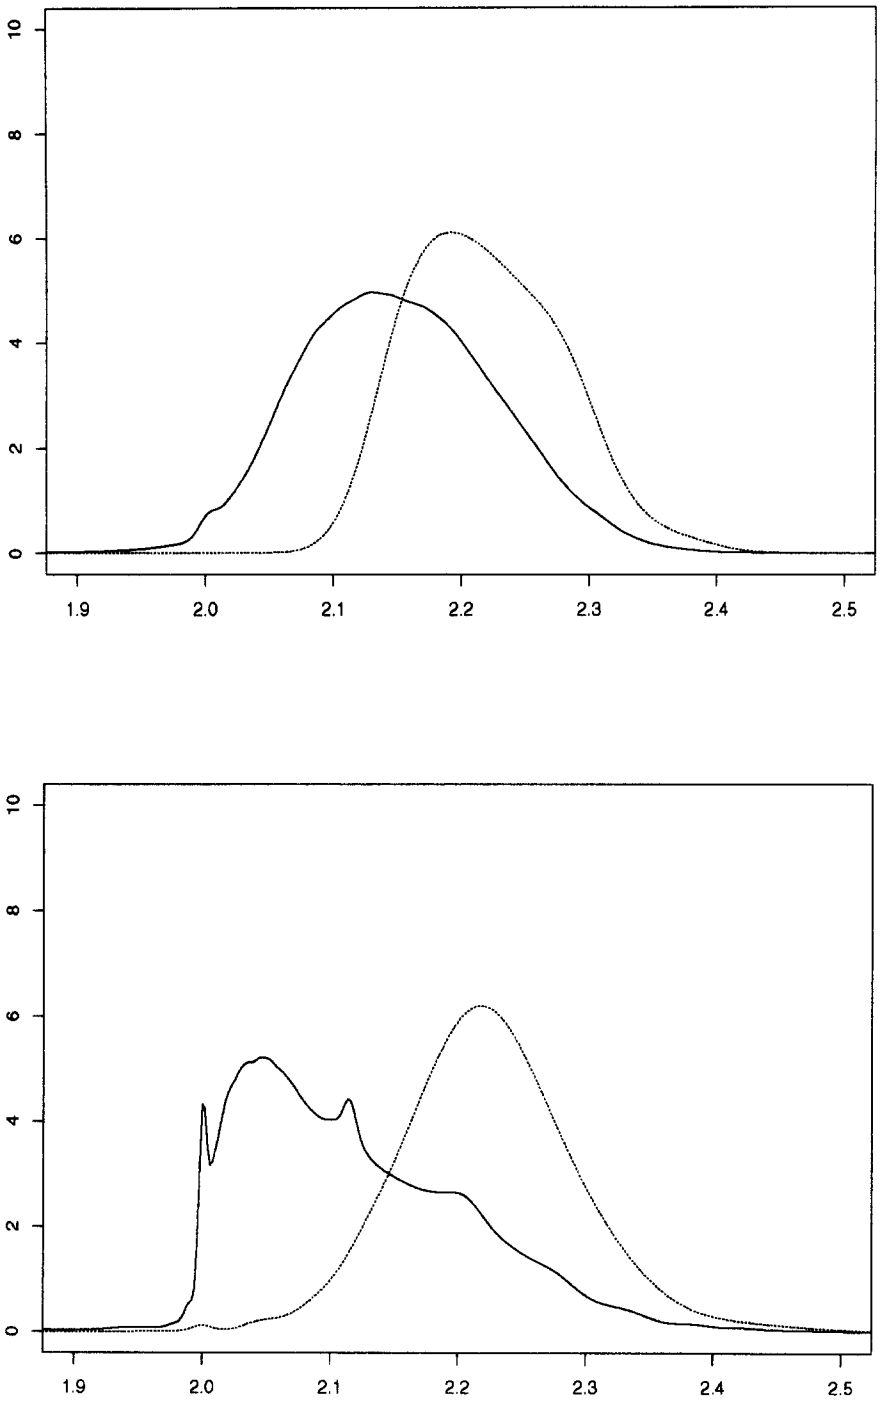

Fig. 4. Probability density functions for healthy tissue (solid curve) versus tumor (dotted curve) without boundary gating (top) and with boundary gating (bottom).

reject H_0 (no improvement). In this case, it is clear from both tests that a significant improvement has been obtained through the use of even this simple boundary gating. Table 2 shows the results of the tests for the probability of detection (PD) at various probabilities of false alarms (PFA). We also include in Table 3 the average PD for each false alarm rate for the two approaches for the purposes of comparison. Once again we have strong evidence for the utility of boundary gating.

The PD values in Table 3 underscore our earlier statement that we do not intend to argue for the use of a single feature for the classification task. A PFA of 1% is completely unacceptable for this problem due to the large number of pixels in an image. Recall that we are classifying at the pixel level. This means that in each image we can expect to have several thousand false alarms. To a degree this problem can be alleviated by combining the pixels in a region to give an overall classification for the

Table 2. Coefficient of variation: Probability of detection

Test	PFA=20%	10%	5%	1%
Wilcoxon	0.0005	0.0009	0.0238	0.174
Paired *t*-test	0.0012	0.0014	0.0236	0.0755

Table 3. Average PD for coefficient of variation

PFA=20 No bnd	PFA=20 Bnd	PFA=10 No bnd	PFA=10 Bnd	PFA=5 No bnd	PFA=5 Bnd	PFA=1 No bnd	PFA=1 Bnd
46.2	51.8	29.5	34.5	10.9	22.1	8.3	9.6

Table 4. Fractal dimension

Test	div p value
Wilcoxon	0.0735
Paired *t*-test	0.0622

region rather than a pixel-by-pixel classification. This would be done in a clinical system. However, the false alarm rate is still far too high. Clearly we need to combine several features together to improve the classification performance. As we stated earlier, it is our purpose here to give a method of obtaining better individual features and the results in Tables 1–3 show that we have succeeded for this feature.

4.2. *Results for fractal dimension and y-intercept*

Proceeding with the fractal dimension and y-intercept as we did with the coefficient of variation. we obtain Tables 4–6 for the fractal dimension and Tables 7–9 for the y-intercept. Although the results are not as dramatic, we still see some improvement in the divergence. A purist might insist that since we fail to reject at the 0.05 level in Table 4, we cannot state that we have a significant improvement. We do not agree with this assessment and would rather state that the improvement does not show up as well in the divergence statistic. However, Table 5 (column one) does show a significant improvement.

Again we see that the probability of detection is far too low for these false alarm rates. Once again we stress that the classifier would not use a single feature in a clinical system.

It should be noted that the improvement drops off as the false alarm rate decreases. This is due in part to the small probability of detection at these levels for these individual features.

The y-intercept actually shows both better performance and better improvement than the fractal dimension. Once again we see an improvement in the feature by using boundary gating to get a more homogeneous region in which to calculate the features.

Although the differences are not nearly as dramatic as with the CoV, it is still evident that the boundary gating case offers an improvement in these features. One point that should be considered is that the boundaries are intensity based while the feature boundaries may not be. While it is reasonable that the boundary between the tumor and healthy tissue would be detectable both in intensity differences and in feature differences, the intensity boundaries within different types of tissues do not necessarily correspond to boundaries of homogeneous regions in feature space. It is possible in principle to improve on these features by using the extracted features to produce feature boundaries to be

Table 7. y-Intercept

Test	div p value
Wilcoxon	0.064
Paired *t*-test	0.032

Table 5. Fractional dimension: probability of detection

Test	PFA=20%	10%	5%	1%
Wilcoxon	0.0303	0.0839	1.1695	0.4792
Paired *t*-test	0.0269	0.0741	0.1939	0.4865

Table 6. Average PD for fractal dimension

PFA=20 No bnd	PFA=20 Bnd	PFA=10 No bnd	PFA=10 Bnd	PFA=5 No bnd	PFA=5 Bnd	PFA=1 No bnd	PFA=1 Bnd
34.9	41.1	21.7	25.2	13.4	14.9	4.7	4.7

Table 8. y-Intercept: probability of detection

Test	PFA=20%	10%	5%	1%
Wilcoxon	0.0199	0.0609	0.2327	0.7566
Paired t-test	0.008	0.0516	0.145	0.3446

Table 9. Average PD for y-intercept

PFA=20 No bnd	PFA=20 Bnd	PFA=10 No bnd	PFA=10 Bnd	PFA=5 No bnd	PFA=5 Bnd	PFA=1 No bnd	PFA=1 Bnd
39.6	44.0	23.9	26.6	14.4	16.2	4.9	5.4

used in place of the intensity boundaries. This is an area of current research. An important point to reiterate is that this study is not concerned with detecting tumors using a single feature. We most emphatically do not suggest using a single feature such as fractal dimension or coefficient of variation to detect tumors. What we are concerned with is the reliable calculation of features which, when used in conjunction with other features, will allow the detection and classification of tumors. Thus, here we are more concerned with an improvement in the features rather than in attaining any fixed performance levels.

5. CONCLUSIONS AND DISCUSSION

It is clear that local features should not be computed without taking into consideration the boundaries between regions of different feature values. The two methods for boundary gating described here show improvement in the feature extraction from the perspective of better classification performance.

Fractal dimension, y-intercept and coefficient of variation are all appropriate features for consideration with other features in a classification system for tumor detection in mammograms. Individually these features do show some ability to discriminate and this ability is enhanced through the use of boundary gating.

In any practical system, multiple features would be utilized both for detection and classification. Additionally, a set of features which is good for detection may or may not be appropriate for classification (for example, between benign and malignant tumors). Finally, it is most probable that a working computer-assisted diagnosis system would use regions for detection and classification rather than dealing with images on a pixel-by-pixel basis. In this manner the system would classify a region based on the features extracted from the region, and thus boundary gating would be critical in the calculation of the features for the regions.

The boundaries used in this study were imperfect for two reasons. First, since they were intensity based, they may not have been perfectly matched with the homogeneous feature regions. Second, since they were not guaranteed to be closed, they allowed some contamination from other regions. Further work is necessary to produce improved boundaries.

There are several possible methods for improving these boundaries. One has already been mentioned. By using the features to recompute boundaries, then using these boundaries to repeat the boundary gated feature extraction, improved features could be obtained. This procedure could probably be improved through noise reduction and Markov random fields methods[23] and is an area of current research. Also, for features like the fractal calculations which perform calculations at several scales, it may be possible to utilize the multiscale nature of the wavelet boundaries. This too is an area of current interest. Finally, there are many different methods for extracting edges from images and some of these are guaranteed to produced closed regions. These need to be investigated, particularly in light of some of the ideas sketched above.

In spite of the caveats in the previous paragraph, the boundaries used in this study were sufficient to produce a significant improvement in the features. It is our belief, as evidenced by this study, that any system which uses local features such as these will be improved using boundary gating.

It should be noted that in some cases intensity-based boundaries are not desirable for use with boundary gating. For example, tumors showing spiculated edges may well be detectable through the texture of these edges. Thus, it is the edge region itself which is used to construct the feature. This points out the necessity of choosing boundaries which are appropriate to the classification task. Given appropriate boundaries, however, it is clear that boundary gating will improve the features extracted.

REFERENCES

1. C. E. Priebe, R. A. Lorey, D. J. Marchette, J. L. Solka and G. W. Rogers, Nonparametric spatiotemporal change point analysis for early detection in mammography, *Digital Mammography*, A. G. Gale, S. M. Astley, D. R. Dancy and AlY. Cairns, eds, pp. 111–120. Elsevier, Amsterdam (1994).
2. J. L. Solka, W. L. Poston, C. E. Priebe, R. A. Rogers, R. A. Lorey, D. J. Marchette, K. Woods and K. Bowyer, The detection of microcalcifications in mammographic images using high dimensional features, *Proc. 1994 IEEE Seventh Symp. on Computer-Based Medical Systems*, pp. 139–145 (1994).

3. K. S. Woods, C. C. Doss, K. W. Bowyer, J. L. Solka, C. E. Priebe and W. P. Kegelmeyer Jr., Comparative evaluation of pattern recognition techniques for detection of microcalcifications in mammography, *Int. J. Pattern Rece. and AI* **7**, 1417–1436 (1993).
4. J. L. Solka, C. E. Priebe and G. W. Rogers, An initial assessment of discriminant surface complexity for power law features, *Simulation* 311–318 (May 1992).
5. C. E. Priebe, J. L. Solka, R. A. Lorey, G. W. Rogers, W. L. Poston, M. Kallergi, W. Qian, L. P. Clarke and R. A. Clark, The application of fractal analysis to mammographic tissue classification, *Cancer Lett.* **77**, 183–189 (1994).
6. X. Liu, G. Zhang and M. Fox, Fractal description and classification of breast tumors, *Proc. Annual Conf. on Eng. in Med. and Biol.*, pp. 112–113 (1991).
7. P. K. Gurski, Fractal-based texture segmentation of digital X-ray mammograms, Masters Thesis, Simon Fraser University (1991).
8. K. S. Woods and K. W. Bowyer, Computer Detection of Stellate Lesions, *Digital Mammography*, A. G. Gale, S. M. Astley D. R. Dancy and AlY. Cairns, eds, pp. 221–229. Elsevier, Amsterdam (1994).
9. C. E. Priebe, J. L. Solka and G. W. Rogers, Discriminant analysis in aerial images using fractal based features, *SPIE, Vol. 1962: Adaptive and Learning Systems II*, 196–208 (1993).
10. C. E. Priebe, E. G. Julin, G. W. Rogers, D. M. Healy, J. Lu, J. L. Solka and D. J. Marchette, Incorporating segmentation boundaries into the calculation of fractal dimension features, *Proc. 26th Symp. on the Interface*, pp. 52–56 (June 1994).
11. D. J. Marchette, C. E. Priebe, G. W. Rogers and J. L. Solka, The filtered kernel estimator, Technical Report, No 104, Center for Computational Statistics, GMU (October 1994).
12. J. Lu, Signal recovery and noise reduction with wavelets, Ph.D. Dissertation, Dartmouth College (1993).
13. J. Lu, D. M. Healy and J. B. Weaver, Contrast enhancement of medical images using multiscale edge representation, *Opt. Eng.* **33**, 2151–2161 (1994).
14. J. Lu, J. B. Weaver, D. M. Healy and Y. Xu, Noise reduction with multiscale edge representation and perceptual criteria, *Proc. IEEE-SP Int. Symp. on Time-Frequency and Time-Scale Analysis*, Virginia, BC, October (1992).
15. T. Peli, Multiscale fractal theory and object characterization, *J. Opt. Soc. Am. A* **7**, 1101–1112 (1990).
16. E. G. Julin, G. W. Rogers, C. E. Priebe and J. L. Solka, Calculation of power law features in the presence of segmentation utilizing a dijkstra potential based algorithm, *SCS Simulation Multiconference* (1994).
17. G. W. Rogers, C. E. Priebe and E. G. Julin, Calculation and comparison of fractal dimension distributions, *Proc. 1994, Summer Computer Simulation Conf.*, pp. 223–229 (1994).
18. G. W. Rogers, C. E. Priebe, H. Hayes and J. L. Solka, A parallel distributed processing algorithm for power law features which requires only nearest neighbor communication, *Proc. Fifth Workshop on Neural Networks, SPIE*, Vol. 2204, pp. 269–275 (1993).
19. B. Mandelbrot, *The Fractal Geometry of Nature*. Freeman, New York (1977).
20. G. McLachlan, *Discriminant Analysis and Statistical Pattern Recognition*. Wiley, New York (1992).
21. R. L. Winkler and W. L. Hays, *Statistics: Probability, Inference, and Decision*. Holt, Rinehart and Winston, New York (1975).
22. J. D. Gibbons and S. Chakraborti, *Nonparametric Statistical Inference*. Marcel Dekker, New York (1992).
23. H. I. Hayes, C. E. Priebe, G. W. Rogers, D. J. Marchette, J. L. Solka and R. A. Lorey, Improved texture discrimination and image segmentation with boundary incorporation, *SPIE Int. Symp. on Aerospace/Defense Sensing and Control and Dual-Use Photonics*, Orlando, FL, April (1995).

About the Author—DAVID J. MARCHETTE earned a B.A. degree in Mathematics from the University of California at San Diego (UCSD) in 1980 and an M.A. in Mathematics from UCSD in 1982. He earned a Ph.D. in Computational Sciences and Informatics at George Mason University under the direction of Professor Edward J. Wegman in 1996. He has been working in pattern recognition and nonparametric density estimation since 1985, first for the Naval Ocean Systems Center in San Diego, CA, and currently for the Naval Surface Warfare Center in Dahlgren, VA.

About the Author—RICHARD A. LOREY earned a B.S. degree in Physics from the University of Pittsburgh in 1962 and an M.S. in Physics and a Ph.D. in Physics from the Georgia Institute of Technology in 1967 and 1969, respectively. He has been working in statistical pattern recognition for three years at the Naval Surface Warfare Center in Dahlgren, VA.

About the Author—CAREY E. PRIEBE earned the B.S. degree in Mathematics from Purdue University in 1984, the M.S. degree in Computer Science from San Diego State University in 1988, and a Ph.D. in Information Technology (Computational Statistics) from George Mason University, under the direction of Professor Edward J. Wegman, in 1993. Since 1985, he has been working in nonparametric estimation and statistical pattern recognition, first for the Naval Ocean Systems Center, San Diego, CA, and the Naval Surface Warfare Center, Dahlgren, VA, and currently as Assistant Professor in the Department of Mathematical Sciences, The Johns Hopkins University, Baltimore, MD.