A NEW APPROACH TO AUTOMATED LABELING OF INTERNAL FEATURES OF HARDWOOD LOGS USING CT IMAGES

Daniel L. Schmoldt
USDA Forest Service
Brooks Forest Products Center
Virginia Tech
Blacksburg VA 24061-0503

Pei Li
Bradley Dept. of Electrical Engineering
Virginia Tech
Blacksburg VA 24061-0111

A. Lynn Abbott
Bradley Dept. of Electrical Engineering
Virginia Tech
Blacksburg VA 24061-0111

INTRODUCTION

In a typical sawmill, logs enter the mill and go through a de-barking process. Following this operation they go to the headrig where a sawyer moves the log repeatedly past a saw to remove boards one at a time. As more of the log interior is exposed with each board removed, the sawyer may re-orient the log periodically to cut from the best side. Sawn boards go through subsequent operations of edging and trimming, where defects near the edges and/or ends of the boards are removed to increase each board's grade, and therefore its value. The cant (the cubical center section of the log) remaining from initial breakdown enters a resawing operation where additional boards are cut. These are also edged and trimmed.

Knowledge of internal log defects, obtained by scanning, is a critical component of efficiency improvements for future mills [1]. Nevertheless, before computed tomography (CT) scanning or any other type of internal log scanning can be applied in industrial operations, there are several hurdles that must be overcome. First, there needs to be some way to automatically interpret scan information so that it can provide the saw operator with the information needed to make proper sawing decisions. A sequence of x-ray tomographs cannot be readily synthesized into a three-dimensional (3D) mental model by human operators [2]. For the purposes of sawing the log cylinder into high-value boards, this means accurately locating, sizing, and labeling internal defects. Second, this defect recognition procedure must operate at real time speeds, so that scanning, image reconstruction, and image interpretation and display can be integrated into mill processing. Third, a 3D display of a log and its defects for the sawyer is only the first step toward real efficiency. Eventually, the sawyer must be guided by computer-analyzed suggestions for the best log breakdown sequence, or have the sawing completely controlled by computer processing [3].

The work described here addresses the first and second of these processing needs. The next section discusses previous work in these areas. This is followed by a detailed

Review of Progress in Quantitative Nondestructive Evaluation, Vol. 15
Edited by D.O. Thompson and D.E. Chimenti, Plenum Press, New York, 1996

description of the neural-net based classification technique that we have developed. Following a description of our experimental methods, performance results are given, including a qualitative comparison with previous approaches. The final section contains conclusions that we have drawn from this work, and some directions for further research.

PREVIOUS WORK

Because most defects of interest are internal, a nondestructive sensing technique is needed which can provide a 3D view of a log's interior. Several different sensing methods have been tried, including nuclear magnetic resonance [4], ultrasound [5], and x-ray. Due to its efficiency, resolution, and widespread application in medicine, x-ray computed tomography has received extensive testing for roundwood applications [6-11]. As noted above, however, CT images require computer analysis before they can be useful in an industrial setting.

Previous work on automatically labeling internal log defects has established the feasibility of utilizing CT images. These researchers have employed a variety of methods to segment different regions of a CT image and then to interpret, or label, those segmented regions. Often, image segmentation methods are based on threshold values derived from image histograms [8, 9, 12]. Texture-based techniques have been applied to defect labeling only [11, 13], not segmentation. Knowledge-based classification [14, 15], shape examination [8, 11], and morphological operations [8] have been used to label defects, also. Hagman and Grundberg [6] used normalized pixel values in a scaled, 8×16 window to label knot types on veneer slices using either a partial least squares classifier or an artificial neural network (ANN). While this approach is interesting, the methods employed were contrived in the sense that objects to be labeled were pre-selected and centered in the analysis window.

While previous efforts have demonstrated feasibility, they have some serious limitations. First, reports of defect labeling accuracy are often either anecdotal, based on success in a training set, or based on a single test set. No statistically valid estimates of labeling accuracy can be found in the literature. Second, there has been no effort to assess or to achieve real-time operability of the developed algorithms. Third, texture information is critical for human differentiation of regions in CT images (i.e. image segmentation), and automated recognition algorithms should exploit this fact for computer-based processing.

This paper presents an alternative to the above approaches that has been developed with these limitations in mind. In contrast to the previous global approaches that separate the tasks of segmentation and region labeling, our approach operates using local pixel neighborhoods primarily, and combines segmentation and labeling into a single classification step. A feed-forward artificial neural network has been trained to accept CT values from a small 3D neighborhood about the target pixel, and then classifies each voxel as knot, split, bark, decay or clear wood. In order to accommodate different types of hardwoods, a histogram-based preprocessing step normalizes the CT density values prior to ANN classification. Morphological postprocessing is used to refine the shapes of detected image regions. These steps are described in the next section.

METHODS

The CT image interpretation system that has been developed here consists of three parts: (1) a preprocessing module, (2) a neural-net based classifier, and (3) a post-processing module. The preprocessing step separates wood from background and internal voids, and normalizes density values. The classifier labels each non-background pixel of a CT slice using histogram-normalized values from a 3×3×3 window about the classified pixel. Morphological operations are performed during post-processing to remove spurious misclassifications.

Preprocessing - Background Thresholding

The first objective of preprocessing is to identify background regions, so that these regions can be ignored by the classifier. Our initial approach was to extract histograms for

individual CT slices and apply Otsu's thresholding method [16]. This method assumes bimodal histograms, and minimizes within-group variance. In our application, it automatically determines a correct threshold for many CT log images (Figure 1), because the histograms are typically bimodal. The two peaks can be found at very low gray-level values (background) and at relatively high CT values, corresponding to clear wood and high-density areas, such as knots and bark. Figure 2 illustrates this with a histogram of densities for the CT slice shown in Figure 1. In Figure 2, the rightmost histogram peak represents clear wood and bark. Knots are denser than clear wood, and tend to cluster at the right side of this peak when present. A large peak representing background is partially shown at the left.

Unfortunately, one of the defect types–decay–has density values which are roughly the average of background (air) and clear wood density values. This appears as a small peak in Figure 2, near the midpoint of the two larger peaks. If Otsu's method is applied directly to this histogram, the threshold indicated by t_1 is detected. Unfortunately, this causes decay regions to be treated as background. We address this problem by weighting the histogram values, using the function

$$w(t) = 1 - e^{-(\frac{t-t_1}{b})^2} \qquad (1)$$

where t_1 is the threshold determined by applying Otsu's method initially, and b = 2000. This value for b was chosen experimentally. The effect of weighting the histogram is essentially to remove the decay peak and reduce the size of the clear wood peak. When Otsu's method is applied to the resulting histogram, the threshold t_2 is found, which successfully distinguishes decay from background. This method has been tested using a large number of CT samples. The weighting function modifies histogram values only for the purpose of determining a threshold value for background pixels. The original pixel CT values are not modified in this step.

Preprocessing - Density Normalization

The second objective of preprocessing is to normalize CT values, so that the classification step can work with different types of wood. Normalization is especially important because neighborhood pixel values are used as features by the classifier. If pixel values are not normalized there will be no consistent relationships among similar regions across CT images, and the ANN classifier will be unable to learn any patterns.

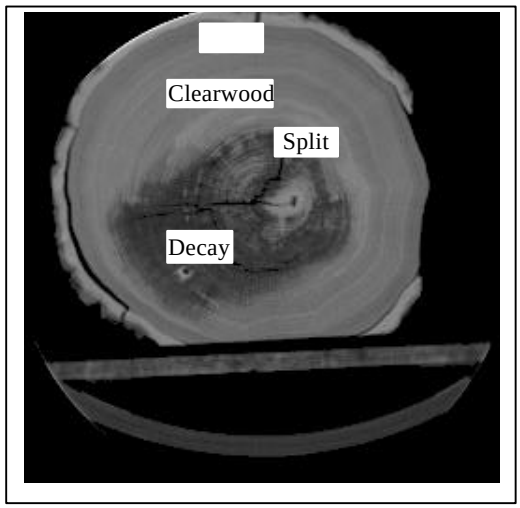

Figure 1. Different densities are depicted by different gray-level values in this computer-generated x-ray tomograph of a red oak log. Regions of clear wood, decay, bark, and splits are visible. The slice shown here contains 256 × 256 elements, each corresponding to a volume of 2.5 × 2.5 × 2.5 mm^3.

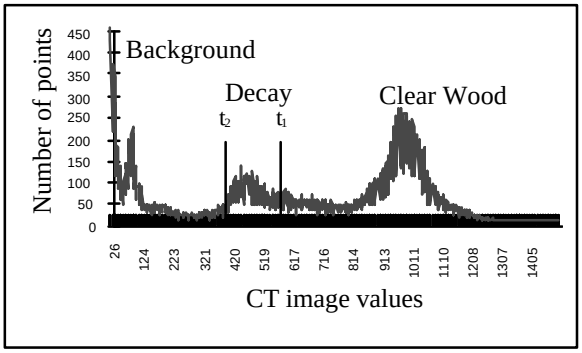

Figure 2. Histogram of a log section. Background pixels produce a very large peak, part of which is omitted from the figure to improve clarity. The t_1 threshold is obtained using Otsu's method directly; t_2 is obtained after introducing a weighting function to the histogram.

All hardwood CT histograms that we have examined have the characteristics of the histogram in Figure 2. That is, there is a large peak of background pixel values at the far left, a large peak of clear wood, bark, and knot pixel values at the far right, and decay pixel values (if present) located at approximately the midpoint of the clear wood values.

To ensure consistency of defect region values across images, we want to be able to do several things with any histogram of CT density values. First, we want to shift the rightmost peak–containing clear wood, bark, and knot values–so that these regions always have the same values and so that the shape of this peak does not change. Second, we want the lower CT values, representing background, to remain about the same following the transformation, so that zero values stay near zero. Third, we want the CT values between the leftmost and rightmost peaks for each original histogram to have the same relative position in a transformed histogram. This type of transformation will give the important regions of any CT image the same density values, and allow us to apply our pixel-value dependent classifier to those normalized values.

The method used here applies a transformation to each CT value in the image. The transformation includes two components: (1) a variable translation component and (2) normalization by an arbitrary parameter. The transformation function is given in Eq. 2:

$$x_t = \frac{x_o + f(x_o ; x_{cw})(x_a - x_{cw})}{x_a} \tag{2}$$

where

x_t	transformed CT value
x_o	original CT value
x_{cw}	original CT value of clear wood peak
x_a	arbitrary translation anchor value, greater than the CT value of the clear wood peak
f	translation multiplier.

The translation anchor x_a is an arbitrary parameter selected to be greater than the CT value of the clear wood peak. The rightmost histogram peak (including clear wood, knot, and bark values) will be shifted to the right by the amount x_a-x_{cw}, so that the clear wood peak is now at x_a. The resulting values are normalized by x_a so that the clear wood peak of a normalized histogram is always located at 1. In order for the translation of the rightmost peak to be consistent for all histograms it is necessary for the translation anchor value to be the same for all histograms. Otherwise, the shape of the rightmost peak will change with respect to the range of transformed density values.

The translation multiplier f is an asymptotic function of the original CT value x_o and is parameterized by the clear wood peak value x_{cw}. It adjusts the amount of the maximum translation x_a-x_{cw} that is added to the original value x_o to arrive at x_t after normalization by x_a. The actual equation for f is as shown below, Eq. 3.

$$f(x_o; x_{cw}) = 1/(1 + e^{(\frac{x_{cw}}{2} - x_o)\alpha}). \quad (3)$$

The range of f is $0 \le f \le 1$, where (1) the slope of f is very steep about the inflection point $x_{cw}/2$; (2) the value of f quickly approaches 0 at values of x_0 less than $x_{cw}/2$; and (3) the value of f quickly approaches 1 at values of x_o greater than $x_{cw}/2$. At $x_o = x_{cw}/2$, f is exactly 1/2. The scale factor α adjusts the steepness of the curve about the inflection point, i.e. how quickly f rises from 0 to 1 as x_o increases. Larger values of α increase the steepness. Initially we have chosen $10/x_{cw}$ as a reasonable value for α.

If we treat all CT values x_o as a proportion β of the clear wood peak value x_{cw}, i.e. $x_o = \beta x_{cw}$ for some β, then Eq. 3 can be rewritten as in Eq. 4, assuming $\alpha = 10/x_{cw}$.

$$f(\beta) = 1/(1 + e^{5-10\beta}). \quad (4)$$

From equations (2-4), we can observe that the following transformations will hold regardless of the original histogram:

$$x_t \cong 1 \text{ for } x_o = x_{cw} \text{ or } \beta = 1$$
$$x_t \cong 0 \text{ for } x_o = 0 \text{ or } \beta = 0$$
$$x_t = 0.5 \text{ for } x_o = x_{cw}/2 \text{ or } \beta = 0.5.$$

A Neighborhood-Based Neural-Net Classifier

A multilayer feed-forward neural network is used to perform the primary classification step. There were two initial goals in this research: (1) to determine if the tasks of segmentation and region labeling could be combined into a single step and (2) to determine whether an ANN classifier could perform well using only simple features obtained from local neighborhoods. Aside from the initial background thresholding, both segmentation and defect labeling are performed simultaneously by the classifier. We have found that such a classifier works quite well, although performance is improved if information concerning distance from the center of the log slice is also included. This distance measure provides contextual information that aids in classification, because some entities (such as splits) tend to lie near log centers and others (such as bark) lie near the outside edge of the log.

Each histogram-normalized value in a 3×3×3 neighborhood about the target pixel serves as an input to the ANN. One additional input is the &radius™ of the element under consideration, which is the distance of the target pixel from the centroid of the foreground region of the CT slice. There are 5 output nodes of the ANN, one for each of the classes to be detected: knot, split, bark, decay or clear wood. The class associated with the output node that has the largest value for a given input is selected as the classification.

The network was trained using the conventional back-propagation method [17]. Because network topology has a large impact on classification accuracy and on convergence time during training, several topologies were compared. Networks using one, two, and three hidden layers were generated, with the total number of weights for each network topology kept constant [18, 19].

At this date, the image interpretation system has been trained using only two hardwood species, northern red oak (*Quercus rubra*, L.) and water oak (*Quercus nigra*, L.). Although these two species are from the same family of oaks, they are from different geographic regions and growing conditions. Training/testing samples were selected from multiple CT

slices. The entire training/testing set consists of 1973 samples. Ten-fold cross-validation was used to estimate the true accuracy rate of the ANN classifier [20].

Postprocessing

Because local neighborhoods are the primary source of classification features that are used by the ANN, spurious misclassifications tend to occur at isolated points. A post-processing procedure is used to remove small regions, thereby improving overall classification accuracy. This method is effective since the defects of interest typically have relatively large sizes in an image. We chose to use the gray-scale operations of erosion followed by dilation for this purpose. A 3×3 structuring element is used for both operations. An added benefit is that labeled region borders are smoothed somewhat during this process.

RESULTS

A sample histogram is presented in Figure 3 to illustrate the effect of our density transformation procedure. Histogram appearance is invariant under this transformation, but values of critical regions are automatically adjusted to be consistent across different CT images.

Four different ANN topologies were trained/tested using ten-fold cross-validation. The results are shown in Table 1. The ANN with two hidden layers exhibited the best performance with an accuracy of just over 90% for pixel classification. The next best classifier, with a single hidden layer of 12 nodes, exhibited practically the same classification accuracy. Because the latter network requires much less processing time, it was chosen as the optimal classifier among those evaluated. It is interesting to note that classification performance decreased slightly as the number of hidden layers increased.

The chosen classifier has been applied to two CT images for illustration (Figure 4). As anticipated, the ANN produces some isolated pixel misclassifications, as shown in the middle column of the figure. The classification regions are improved with post-processing, however, as shown at the right. In the second example of Figure 6, for example, the ANN classified partial regions of several growth rings as split defects; these were removed by subsequent postprocessing. In the first example in that figure, incorrect labels near the outside border of the CT slices are removed by postprocessing steps.

The image interpretation system is currently implemented on a Macintosh[1] Quadra 650 containing an MC68040/33MHz processor. Analysis of a single 256×256 CT slice requires about 25 seconds. This is considerably faster than the previous approach (Zhu 1993) which requires 9 minutes of processing time on a VAX 11/785. Because the algorithms are implemented in C, however, they can be transported easily to any other computer hardware.

In comparison to previous hardwood log inspection systems, our system has a simple implementation, but high classification speed and accuracy. Other systems are reported to be able to successfully identify or locate some internal defects, but few statistical results are available. Most previous work is limited to 2D image analysis, which does not make full use of the 3D nature of CT images. Finally, most research has dealt with a single type of wood, whereas our approach successfully deals with two different wood species.

Table 1. Network topologies and classification performance

Network topology	Number of weights	Number of training iterations	Classification accuracy
28-12-5	396	6699	0.898275
28-10-8-5	400	8299	0.902442
28-7-16-5	388	10499	0.869596
28-8-8-8-5	392	60499	0.852903

[1] Tradenames are used for informational purposes only. No endorsement by the U.S. Department of Agriculture is implied.

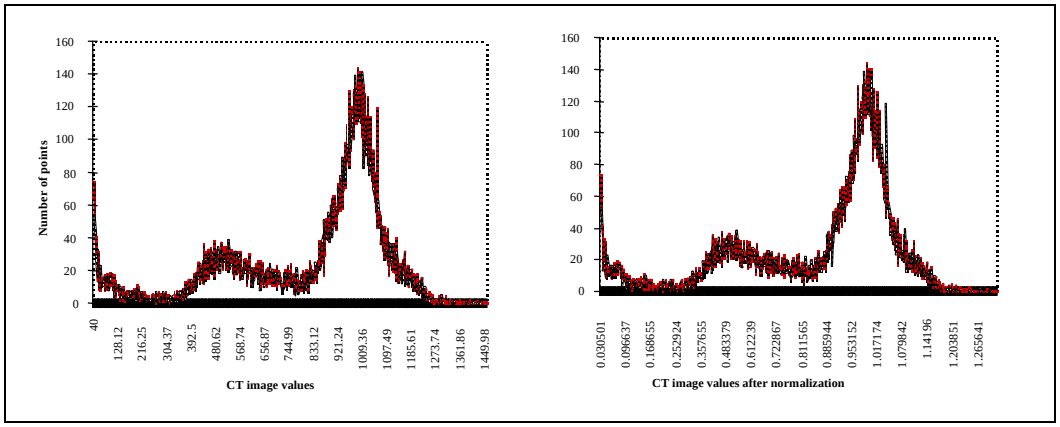

Figure 3. These CT image histograms illustrate the effect of transforming density values. The original CT image histogram appears on the left and the transformed histogram appears on the right. Visually, the histograms do not change, while values for critical regions become approximately the same.

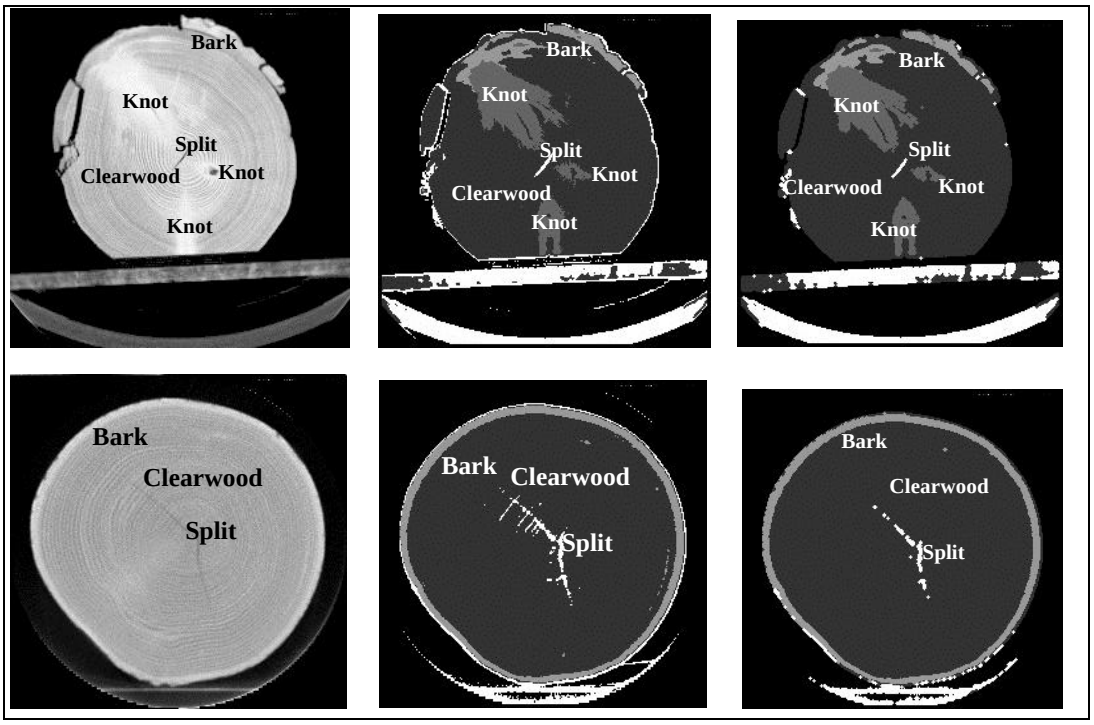

Figure 4. Two log CT images demonstrate defect recognition results. Original CT images appear at the left, middle images are ANN classified images, and the rightmost images depict the classification results following postprocessing.

CONCLUSIONS

In general, the ANN classifier, operating primarily with local, pixel values, is able to classify regions of CT images with high accuracy. The resulting classification performance is 90% accuracy at the pixel level. Postprocessing improves this value considerably, but we do not have an exact numerical estimate for this improvement. Most regions are detected and correctly labeled; however, in some cases the classifier fails to correctly size defects. It is possible that by the addition of further postprocessing, e.g., high-level, rule-based analysis

of defect regions, we may be able to size defects more accurately and to remove any remaining misclassified regions.

As noted above, the entire classification operation requires only about 25 seconds on the current hardware. By using newer RISC-based hardware, this defect recognition time can be reduced drastically, by a factor of 8-10. This places defect recognition speed on a par with scanning and image reconstruction times. Because each of these 3 operations takes 2-3 seconds, they can be performed in parallel on successive slices. Therefore, this defect recognition technique can easily be implemented in real time as logs are scanned and images reconstructed.

Because of the success of the trained ANN classifier on oak samples, we feel confident that we can develop species-dependent classifiers that are very accurate. It is not clear, however, whether we will be able to create a classifier that is entirely independent of species. Should a generalized classifier prove to be infeasible, species-dependent classifiers can still be useful in actual mill operations because typically a single species is sawn over an extended period. Additional samples of CT images for other species need to be collected. This will enable us to verify the efficacy of our density normalization technique and the ability of our classifier (or a newly trained classifier) to correctly label and size internal features of logs.

REFERENCES

1. L. G. Occeña, Journal of Forest Engineering 3, (1991).
2. D. L. Schmoldt, D. P. Zhu, and R. W. Conners, in *Review of Progress in Quantitative Nondestructive Evaluation*, Vol. 12, eds. D. O. Thompson and D. E. Chimenti (Plenum Press, New York, 1993)
3. L. G. Occeña, W. Chen, and D. L. Schmoldt, in *Proceedings of the 4th Industrial Engineering Research Conference*, ed. B. Schmeiser (1995)
4. S. J. Chang, P. C. Wang, and J. R. Olson, in *2nd International Conference on Scanning Technology in Sawmilling*, ed. R. Szymani (Forest Industries/World Wood, San Francisco CA, 1987)
5. W. Han and R. Birkeland, Industrial Metrology 2, (1992).
6. P. O. G. Hagman and S. A. Grundberg, Hols als Roh- und Werkstoff 53, (1995).
7. C. W. McMillin, Wood Science 14, (1982).
8. S. Som, P. Wells, and J. Davis, in *Second International Conference on Automation, Robotics and Computer Vision*, (1992)
9. F. W. Taylor, J. F. G. Wagner, C. W. McMillin, I. L. Morgan, and F. F. Hopkins, Forest Products Journal 34, (1984).
10. D. Zhu, R. W. Conners, F. M. Lamb, D. L. Schmoldt, and P. A. Araman, in *4th Internation Conference of Scanning Technology in Sawmilling*, ed. R. Szymani (Forest Industries/World Wood, San Francisco CA, 1991)
11. B. V. Funt and E. C. Bryant, Forest Products Journal 37, (1987).
12. D. Zhu, R. W. Conners, and P. Araman, in *IEEE International Conference on Systems Engineering*, (1991)
13. D. Zhu, A. A. Beex, and R. W. Conners, in *Stochastic and Neural Methods in Signal Processing, Image Processing, and Computer Vision*, (SPIE-The International Society for Optical Engineering, Bellingham WA, 1991)
14. D. Zhu, *A Feasibility Study on Using CT Image Analysis for Hardwood Log Inspection* (Virginia Tech University, 1993),
15. D. Zhu, R. W. Conners, D. L. Schmoldt, and P. A. Araman, in *Proceedings of the 1991 IEEE International Conference on Systems, Man, and Cybernetics*, (1991)
16. N. Otsu, IEEE Transactions on Systems, Man, and Cybernetics SMC-9, (1979).
17. J. McClelland and D. Rumelhart, *Parallel Distributed Processing*. (MIT Press, Cambridge MA, 1986).
18. R. Nekovei and Y. Sun, IEEE Transactions on Neural Networks 6, (1995).
19. M. Özkan, B. M. Dawant, and R. J. Maciunas, IEEE Transactions on Medical Imaging 12, (1993).
20. S. M. Weiss and C. A. Kulikowski, *Computer Systems That Learn*. (Morgan Kaufmann Publishers, Inc., San Mateo, 1991).