Labeling Defects in CT Images of Hardwood Logs With Species-Dependent and Species-Independent Classifiers

Pei Li*, Jing He*, A. Lynn Abbott*, and Daniel L. Schmoldt**

*The Bradley Department of Electrical and Computer Engineering
Virginia Polytechnic Institute and State University
Blacksburg, Virginia 24061-0111 USA

** USDA Forest Service
Brooks Forest Products Center
Blacksburg, Virginia 24061-0503 USA

Abstract:

This paper analyses computed tomography (CT) images of hardwood logs, with the goal of locating internal defects. The ability to detect and identify defects automatically is a critical component of efficiency improvements for future sawmills and veneer mills. This paper describes an approach in which 1) histogram equalization is used during preprocessing to normalize pixel values; 2) a feed-forward neural network assigns tentative labels to individual image pixels; and 3) a morphological post-processing step removes noise and refines image regions. The normalization step facilitates the classification of wood features across different logs and different species. The neural network assigns tentative labels using normalized pixel values from small three-dimensional (3D) neighborhoods. We demonstrates the utility of this approach when the the network is trained using a single species of wood. This paper also considers the effect of training the network with samples from more than one species. Because small neighborhoods are used in either case, the classifier can be made to operate at real-time rates. Tests of the method using ten-fold cross-validation and CT images from three different logs resulted in a classification accuracy of approximately 95%.

1 Introduction

Several steps are required for processing hardwood logs. Logs are transported to a sawmill, and an assessment of each log's quality is performed. Logs with the highest quality are shipped to

veneer mills, where they are sawn into flitches and then sliced to produce veneer. The remaining logs are sawed into lumber. Quality is inversely related to the presence of defects such as knots, splits, voids, and decay in the wood. When a log is to be sawed, a cutting strategy must be selected that preserves large areas of clear wood on board faces. There is a strong incentive to perform a correct assessment, for both veneer and saw logs, since the economic return can improve considerably (possibly up to a factor of 10) when a correct decision is made.

In a typical sawmill, logs enter the mill and go through a de-barking process. Following this operation they go to a 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. A long-term goal of this research is to automate these sawing processes, and this requires the ability to detect the presence of defects automatically.

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 (Chang, et al. 1987), ultrasound (Han and Birkeland 1992), and x-rays. Because of its efficiency, resolution, and widespread use in medical applications, x-ray computed tomography has received extensive testing for roundwood applications (e. g., McMillin 1982, Zhu, et al. 1991). An x-ray CT scanner produces image slices that capture many details of a log's internal structure. A typical slice in our data set contains 256 x 256 elements, each corresponding to a volume of 2.5 mm x 2.5 mm x 2.5 mm. CT numbers are directly related to density, and CT images can therefore vary considerably for different species and different moisture content. For example, a log that is freshly cut will produce different CT values than one that has had time to dry. In previous research, CT image analysis has focused on a single CT slice, although in a few cases neighboring slices have been used for 3D filtering during preprocessing steps.

This paper presents an alternative to previous methods. A feed-forward artificial neural network (ANN) has been employed to classify each pixel in an image slice. The ANN accepts CT values from a small 3D neighborhood around each pixel, and then classifies the center pixel of the neighborhood as knot, split, bark, decay, or clear wood. In order to accommodate different types of hardwoods, a histogram-based preprocessing step normalizes pixel values in each CT image. Following initial classification by the ANN, a postprocessing step is performed to refine the shapes of detected image regions. The major benefits of this classification approach are high computational speed and relatively high classification accuracy. The system

has been extensively tested for a single species of wood, and ten-fold cross-validation indicates a classification accuracy of 95% by the ANN before postprocessing. The potential for parallelization is high, since local neighborhoods are used, and since the classifier can be applied to all pixels in parallel. This work is a continuation of that reported in (Schmoldt 1996) and (Li 1996).

The next section of this paper describes the preprocessing steps that are used by the system. Section 3 describes the topology, features, and training method for the neural network classifier. Section 4 presents results for different combinations of wood species. Section 5 summarizes the paper.

2 Preprocessing

The first objective of preprocessing is to identify background regions, so that these pixels can be ignored by the classifier. Our initial approach was to extract histograms for individual CT slices and apply an adaptive thresholding method (Otsu 1979). This method assumes bimodal histograms, and automatically selects a threshold for a histogram h to minimize within-group variance. In our application, it automatically determines a correct threshold for many CT log images. At very low density values, a large peak is present which represents the background. Another peak is present at relatively high CT values, corresponding to clear wood and high-density areas such as knots and bark. Knots are denser than clear wood, and tend to cluster at the right side of this peak when present. Decay values lie near the midpoint of the two major peaks.

Unfortunately, decay causes a histogram peak that violates this bimodal assumption. Using our original thresholding method, we found that decay was often treated as background and was therefore not detected. We then developed the following mapping w that addresses this problem,

$$w(t) = 1 - \exp\left[-\left(\frac{t - t_1}{b}\right)^2\right], \tag{1}$$

where t is a given CT density value, t_1 is the threshold determined by applying Otsu's method initially, and b is a constant that is chosen empirically. If the modified histogram $h'(t) = w(t)h(t)$ is now considered, the effect of the weighting function is to remove the decay peak and reduce the size of the clear wood peak relative to the background peak. If Otsu's method is now used on the modified histogram h', the chosen threshold is shifted to the left. This places the decay peak to the right of the chosen threshold, so that decay pixels will be retained when the background is removed. Note that this weighting is used only for the purpose of choosing a threshold value. The original pixel CT values are not modified in this step.

The second objective of preprocessing is to normalize CT values so that a single classifier can work with different types of wood. Normalization is especially important because the resulting density (pixel) values are used directly as features by the classifier, as described below. If pixel values were not normalized, then an ANN classifier would not be able to distinguish internal features of logs with even modest differences in moisture content or intrinsic density characteristics.

To ensure consistency of defect region values across images, we developed the transformation

$$x_{norm} = \frac{1}{x_a}\left[x_0 + \frac{x_a - x_{cw}}{1 + \exp a\left(\frac{x_{cw}}{2} - x_0\right)}\right] \tag{2}$$

which maps original CT values x_0 to normalized values x_{norm}, giving roughly the same density values to identical features of CT log images. For example, knot defects on one log will have the same normalized CT values as knot defects on another log. This allows us to use a classifier that has been trained using those normalized values. In this equation, the translation anchor x_a is arbitrarily selected to be greater than the CT value of the clear wood peak x_{cw} for any scanned log. The quantity a is a constant, and has been set to $10 / x_{cw}$. Intuitively, small and large values of x_0 pass through (almost) linear mappings, whereas values of x_0 near $x_{cw} / 2$ are expanded into a larger range of values. Perhaps most importantly, the clear wood peak is mapped approximately to the normalized value 1.0 for all CT scans.

3 A Neighborhood-Based Neural-Net Classifier

Using normalized CT values, we have successfully used a multilayer feed-forward neural network to perform the primary classification step. A major objective of this work has been to determine whether an ANN classifier could perform well using only normalized CT values obtained from small, local neighborhoods. We have found that such a classifier works reasonably well, although performance is improved if information is also included concerning the distance of the target pixel from the center of the log slice. 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.

We have tested this approach using several scanned logs, and using both 3 × 3 and 3 × 3 × 3 neighborhood windows. Each histogram-normalized value in the neighborhood serves as an input to the ANN. One additional input is the radial distance of the element under consideration,

which is the distance of this 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. In classification mode, the ANN assigns a label to the pixel based on the output node that has the largest value for the given input values.

Another major goal of this work has been to evaluate the performance of this method for individual species, and to determine the extent to which multiple species can be accommodated by a single ANN. Our first training/testing set consisted of 1973 samples for two species of oak, 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. The training/testing samples were selected from multiple CT slices.

Results of this "oak network" are summarized in Table 1. (This will be compared with other classifiers in the next section.) The network was trained using the conventional back-propagation method. 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 (Nekovei and Sun 1995).

Ten-fold cross-validation was used to estimate the true accuracy rate of the ANN classifier. In ten-fold cross-validation, the set of all samples is divided into 10 partitions. At each stage of the ten-step process, one of the partitions is reserved for testing, the classifier is trained on the remaining 9 partitions, and after training is complete the classifier is tested on the reserved partition. This process is repeated 10 times; final classification accuracy for the classifier is the average of the 10 test partitions. Cross-validation provides an objective and statistically valid estimate of the true classification rate (Weiss and Kulikowski 1991).

As indicated in Table 1, the ANN with two hidden layers exhibited the best performance with an accuracy of just under 95%. 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.

We compared this 3D classifier with a similar ANN that used two-dimensional (2D) neighborhoods only. Using only 9 pixels from a 2D neighborhood, rather than 27 from the corresponding 3D neighborhood, classification accuracy dropped from 94.7% to 93.7%. Consequently, we be-

Table 1: Several neural network topologies were compared with respect to classification accuracy and speed of training

Network topology	Number of weights	Number of training iterations	Classification accuracy
28-12-5	396	6699	0.948
28-10-8-5	400	8299	0.949
28-7-16-5	388	10499	0.940
28-8-8-8-5	392	60499	0.854

Table 2: Three classifiers of 28-12-5 neural network topologies were compared with respect to classification accuracy and the number of training samples

Three classifiers	Number of training samples	Classification accuracy
Oak classifier	1973	0.948
Yellow poplar classifier	1018	0.891
Combined classifier	1983	0.963

gan to consider 2D neighborhood classifiers as having approximately the same accuracy as their 3D counterparts.

All of the neural networks considered here were trained using the delta rule with a momentum term. The effect of learning parameters on the speed of training convergence was studied by experimenting with various learning coefficients and momentum terms. The final choice of the learning parameters is a small learning coefficient (0.1) and a medium momentum term (0.6). Experiments using different initial weights to train the networks show that the choice of initial weights has a negligible effect on the training process and on the performance of the classifier.

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. Classification accuracy is greatly improved by this step (Schmoldt et al. 1996), but because re-classified pixels cannot be quantitatively

tracked we only have visual and qualitative indication of the improvement

4 Results and Conclusions

To date, we have generated 3 distinct image classifier systems, each represented by a neural network having the 28-12-5 topology: 1) an "oak classifier" that was trained/tested using 1973 samples of two hardwood species, northern red oak *(Quercus rubra,* L.) and water oak *(Quercus nigra,* L.), as described in the previous section; 2) a "yellow poplar classifier" that was trained and tested using 1018 samples of yellow poplar *(Liriodendron tulipifera,* L.); and 3) a "combined classifier" that was tested and trained using 1983 samples chosen randomly and at equal proportions from the oak and yellow poplar CT slices. All three classifiers accept as input features the histogram-normalized CT density values from 3 × 3 × 3 neighborhoods, along with radial distance, as described in Section 3. The oak classifier was trained to label each non-background pixel as one of five classes: knot, split, bark, decay, and oak clear wood. The yellow poplar classifier was trained to recognize knots, splits, bark, yellow-poplar sapwood, and yellow-poplar heartwood. In oak, heartwood and sapwood are not very different, but for yellow poplar, heartwood and sapwood CT values are very different. Therefore, we needed to distinguish those 2 yellow poplar features for classification purposes. Subsequent application of classified yellow poplar images will combine those 2 classes into the class "clear wood." The yellow poplar classifier was not trained on samples of decay because our data set did not exhibit that defect type in yellow poplar. Also, we only had one small sample of the feature "split" in our yellow poplar data set. The combined classifier was trained to recognize knots, splits, bark, decay, oak clear wood/yellow-poplar heartwood, and yellow-poplar sapwood. Clear wood in oak and heartwood in yellow poplar are very similar with respect to CT values and can be treated as the *general* clear wood class.

Figures 1 through 4 compare results obtained using these classifiers. The example CT images were chosen because they exhibit all defects of interest. Figure 1 compares results from the oak classifier, and from the combined classifier, for a red oak slice. There is no decay on this image. The major defect regions on this image are identified correctly. It is clear from visual inspection that both classifiers work very well for red oak images.

Figure 2 compares the same two classifiers for another red oak slice. This image has decay and an included split. The split is large enough to be classified as void during background thresholding. Again, both the species-specific classifier and the species-independent classifier seem to work equally well.

Figure 3 compare outputs from the yellow poplar network and from the combined classifier for a yellow poplar input image. There is no sapwood in this image, so both the yellow poplar classifier and the combined classifier give similar results.

Figure 4 compares the same two classifiers for an other yellow poplar slice. In this slice, sapwood and heartwood are both present. The yellow poplar classifier identifies most features correctly. There is some difficulty identifying the entire bark region, though. On the other hand, the combined classifier can still identify yellow poplar heartwood as clear wood. Sapwood, however, is classified as a knot region and only small amounts of the bark region are correctly labeled. Also, there is a large region at the log center that has been labeled as split. This is partially due to the paucity of yellow poplar images containing split features. In contrast to the results of the combined classifier in Figure 4, Figure 5 illustrates results for the combined classifier on another yellow poplar image from another log. Some sapwood pixels are still misclassified, but a large percentage are also correctly labeled. We expect that by retraining the combined classifier with some samples from the CT slice of Figure 4, it will be able to discriminate sapwood regions more accurately. Similarly, the current misclassification of central heartwood regions as splits could be partial mitigated by incorporating more yellow poplar split examples in our training data.

5 Summary

This paper has described a system that is capable of locating and identifying defects in CT images of hardwood logs. This represents a major new component that is needed for further automation of hardwood mill operations. In comparison to previous hardwood log inspection systems, our system has a simple implementation, but relatively 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. Our approach gives statistically valid estimates of classification accuracy. 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 three different wood species.

Already, one x-ray scanner has been installed at a sawmill in Canada (Aune 1995), although defects are currently detected by manual examination of CT slices. As similar systems are developed and deployed, automated defect detection software will become increasingly important.

References

[1] Aune, J. 1995. "The development of a log scanner for sawmills in Canada," in O. Lindgren (Ed.), *Proceedings: 2nd International Seminar on Scanning Technology and Image Processing on Wood,* Luleå University of Technology, Sweden, pp. 51-55.

[2] Chang, S. J., P. C. Wang, and J. R. Olson. 1987. "Nuclear magnetic resonance imaging of hardwood logs," in R. Szymani (Ed.), *2nd International Conference on Scanning Technology in Sawmilling,* October 1-2, 1987, Oakland/Berkeley Hills CA, Forest Industries/World Wood, San Francisco CA.

[3] Han, W., and R. Birkeland. 1992. "Ultrasonic scanning of logs," *Industrial Metrology* 2(3/4): 253-282.

[4] Li, P., A. L. Abbott, and D. L. Schmoldt. 1996. "Automated Analysis of CT Images for the Inspection of Hardwood Logs", *Proceedings: International Conference on Neural Networks,* June 3-6, 1996, Washington, D. C., pp. 1744-1749.

[5] McMillin, C. W. 1982. "Applications of automatic image analysis to wood science," *Wood Science* 14(3): 97-105.

[6] Nekovei, R., and Y. Sun. 1995. "Back-propagation network and its configuration for blood vessel detection in angiograms," *IEEE Transactions on Neural Networks,* 6(1): 64-72.

[7] Otsu, N. 1979. "A threshold selection method from gray-level histograms," *IEEE Transactions on Systems, Man, and Cybernetics,* SMC-9: 62-66.

[8] Schmoldt, D. L., P. Li, and A. L. Abbott. 1996. "Machine vision using an artificial neural network and 3D pixel neighborhoods," to appear in *Computers and Electronics in Agriculture.*

[9] Weiss, S. M., and C. A. Kulikowski. 1991. *"Computer Systems That Learn."* San Mateo: Morgan Kaufmann Publishers, Inc.

[10] Zhu, D., R. W. Conners, F. M. Lamb, D. L. Schmoldt, and P. A. Araman. 1991. "A computer vision system for locating and identifying internal log defects using CT imagery," in R. Szymani (Ed.), *Proceedings: 4th Internation Conference of Scanning Technology in Sawmilling,* Oct. 28-31, San Francisco CA, Forest Industries/World Wood, San Francisco CA, 13 p.

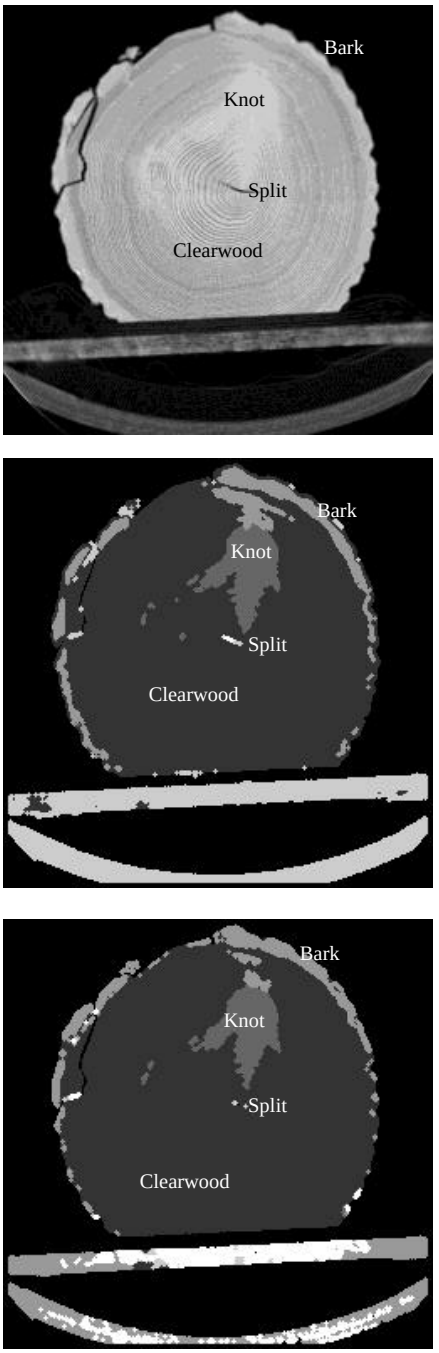

Figure 1: Comparison of classifiers for a red oak image. (a) Original CT image. (b) Output of oak classifier. (c) Output of combined classifier.

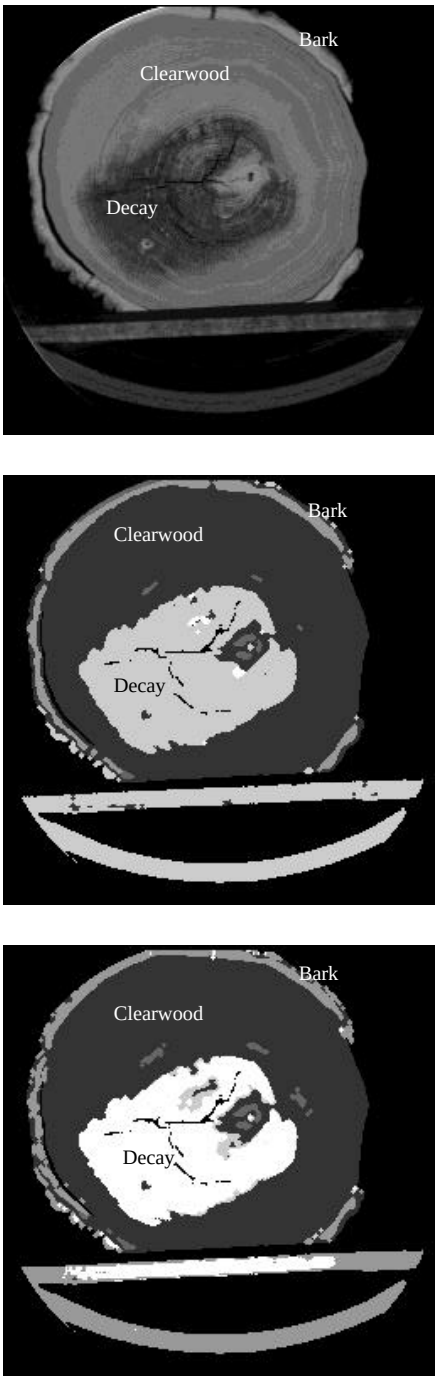

Figure 2: Comparison of classifiers for another red oak image, (a) Original CT image. (b) Output of oak classifier. (c) Output of combined classifier.

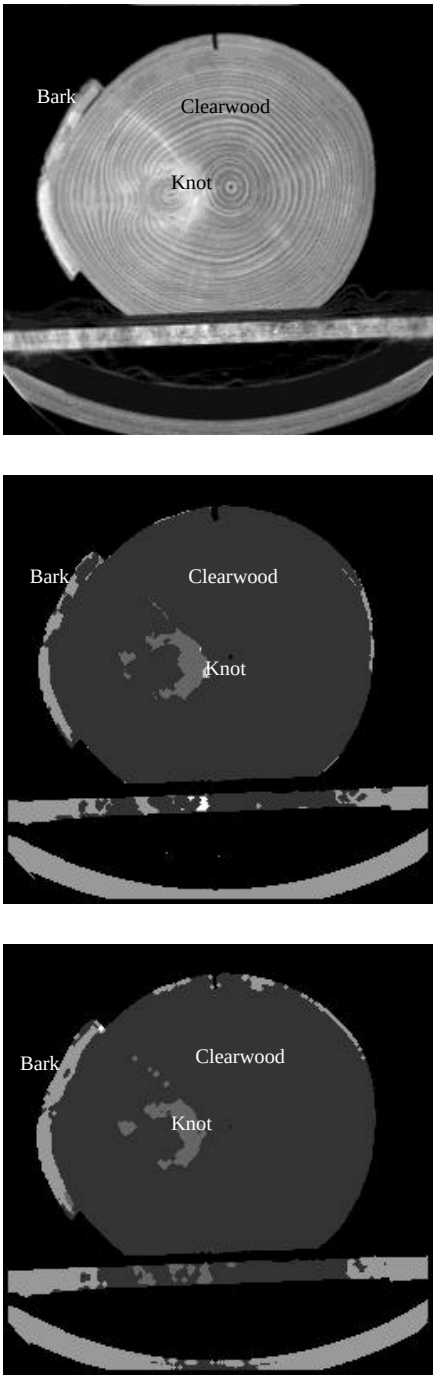

Figure 3: Comparison of classifiers for yellow poplar image. (a) Original CT image. (b) Output of yellow poplar classifier. (c) Output of combined classifier.

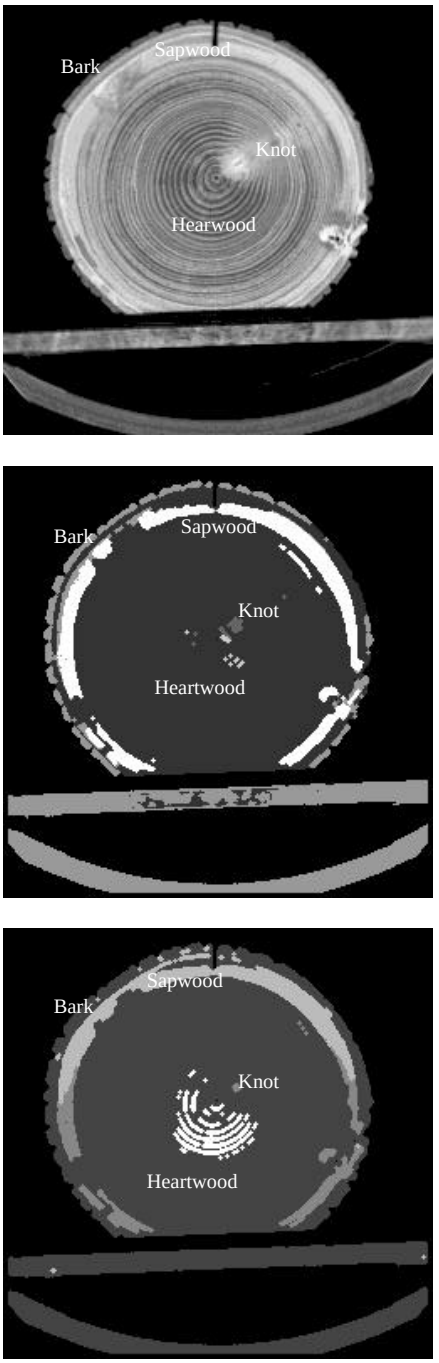

Figure 4: Comparison of classifiers for another yellow poplar image. (a) Original CT image. (b) Output of yellow poplar classifier. (c) Output of combined classifier.

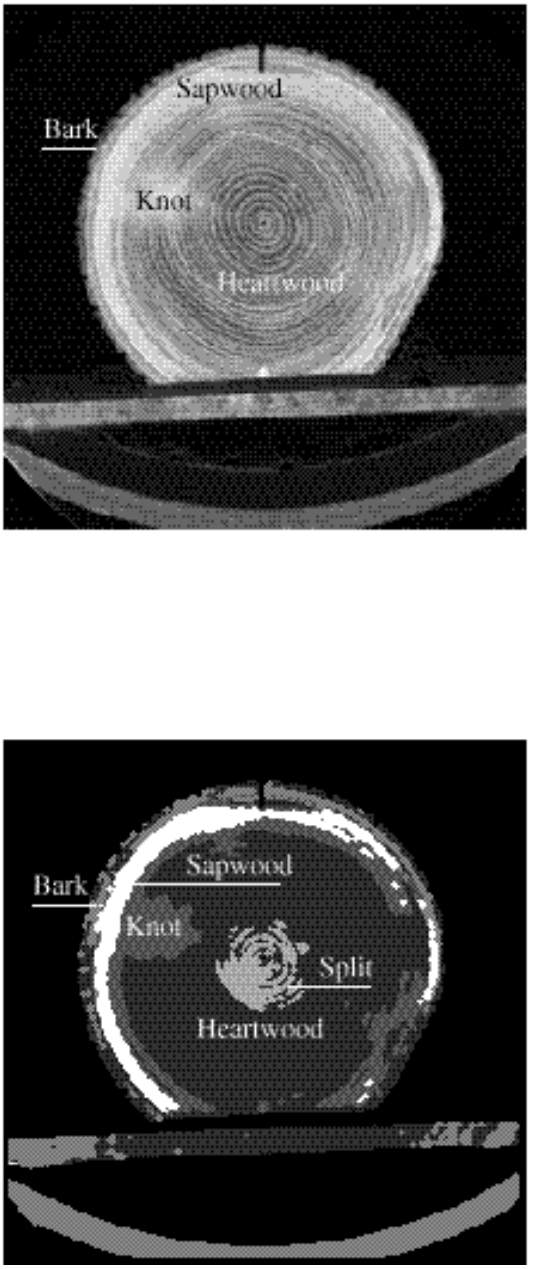

Figure 5: Example of the combined classifier for another yellow poplar image. (a) Original CT image. (b) Output of combined classifier.

Schriftenreihe

A. Pinz; W. Pölzleitner (eds.)

Machine Perception Applications

Proceedings of the IAPR TC-8 Workshop
on Machine Perception Applications
Technical University Graz, Austria
2-3 September, 1996

Österreichische Computer Gesellschaft
R. Oldenbourg Wien München