An Interactive Machine-Learning Approach for Defect Detection in Computed Tomography (CT) Images of Hardwood Logs

Erol Sarigul[a], A. Lynn Abbott[a], Daniel L. Schmoldt[b], and Philip A. Araman[c]

[a]Bradley Department of Electrical and Computer Engineering
Virginia Tech
Blacksburg, VA 24061

[b]USDA Cooperative State Research, Education, and Extension Service (CSREES)
800 9th Street SW, Waterfront Centre, Room 3122
Washington, DC 20024

[c]USDA Forest Service, Southern Research Station
Brooks Forest Products Center
Virginia Tech
Blacksburg, VA 24061

Abstract

This paper describes recent progress in the analysis of computed tomography (CT) images of hardwood logs. The long-term goal of the work is to develop a system that is capable of autonomous (or semiautonomous) detection of internal defects, so that log breakdown decisions can be optimized based on defect locations. The problem is difficult because wood exhibits large variations in texture along with irregular defect placement, particularly for hardwood species. In an earlier project, we developed a classification system that utilizes artificial neural networks (ANN) for this purpose. The system uses small neighborhoods in a CT image to make a preliminary classification decision for every pixel, using labels such as "knot," "split," and "bark." This approach has yielded high accuracy statistically. Subjectively, however, the results can often be improved through further processing steps. For that purpose, we have developed a prototype system called IntelliPost, which can refine a segmented image. During its "learn mode," IntelliPost observes image-editing operations performed by a human operator, and develops its own rules based on those actions. Later, the user can place the system into "run mode" and provide new segmented images. The system automatically refines these new images by using the rules that it has developed. This approach allows IntelliPost to be tailored for different application domains (e.g., species and grading criteria) and for different user preferences. In tests involving CT datasets of red oak and sugar maple logs, the use of IntelliPost resulted in pixel-wise accuracy improvements ranging from 1.61% to 19.47%.

Introduction

Growing demand for lumber products and limited forest resources are forcing the hardwood industry to seek more productive conversion of logs to lumber. Conventional log sawing practices waste a considerable amount of valuable wood, largely because defects that significantly lower board quality are at unknown locations inside the logs.

The conventional breakdown method of a log relies on visual examination. The sawyer first considers the exterior of a log and chooses an initial breakdown strategy. This is modified as sawing incrementally reveals the log's interior. This method has several drawbacks. Among the most noticeable is that exterior bark distortion provides only limited information concerning internal features. In addition, repeated manual operations are subject to fatigue and subjective variation.

Computed tomography (CT) imaging is one possible method to obtain information concerning the internal structure of logs. This nondestructive technique provides image "slices" representing cross-sectional density distribution of a scanned object. Due to the fact that a large amount of data is collected during typical CT scanning, it is preferred to have an automated technique that can quickly analyze the images, locate and identify defects, and suggest breakdown strategies [4, 5, 10, 13].

A collaborative research effort involving Virginia Tech and the Southern Research Station of the USDA Forest Service has resulted in a method that utilizes artificial neural networks (ANNs) to classify CT image pixels individually. This system uses small kernel windows of CT density values as input feature vectors [10,16,17]. The trained neural network assigns a label such as "knot", "split", or "clear wood" to each nonbackground pixel in the image. Several experiments have demonstrated good accuracy in labeling of log defects in CT slices. Although the classification performance of the ANN-based classifier has been quantitatively high, the experiments have shown a need for a postprocessing module to refine classified images further. The primary reason is that the ANN-based classifier depends heavily on local information, and this can result in misclassifications, particularly in small, isolated locations. The initial approach for postprocessing was a fixed, nonadaptive method that was based on mathematical morphology to refine the classified image regions. While that approach has shown some success, it was limited in its ability to handle a wide variety of postprocessing needs.

Several difficulties are faced in improving the postprocessing module, however. Although many fixed postprocessing steps can be easily implemented, different situations may require different types and degrees of postprocessing. For example, different species of wood, particular defect types, the intended use of a log, and personal preferences suggest the need for a more flexible approach. For these reasons, we have developed an adaptive system called IntelliPost that is capable of "learning" new sets of postprocessing rules. The system operates by observing the postprocessing operations as a human user interactively refines segmented CT images. As the user manipulates the images, IntelliPost stores information related to those manual operations, and develops internal rules that can be used later for automatic postprocessing of other images. After one or more training sessions, the system accepts new images, and uses its rule set to apply postprocessing operations in a manner that is modeled after those learned from the human user.

IntelliPost does not simply memorize a particular sequence of postprocessing steps during a training session, but instead generalizes from the image data and from the actions of the human user so that new CT images can be refined appropriately. Because it learns from a human "teacher", this approach represents a form of supervised machine learning. However, the level of supervision is relatively mild by traditional machine-learning standards, because the teacher does not need to be knowledgeable concerning the system's internal feature space or its rule representation and selection methods. Because of its ability to accept new training inputs over time, the system is said to perform "incremental" learning. This contrasts with many machine-learning systems, which require all training data to be made available prior to training. Such systems perform "batch" learning.

The next section of this paper provides a brief overview of the complete defect detection system. This is followed by a description of the IntelliPost postprocessing system, and then a section that presents experimental results that have been obtained with IntelliPost. The final section presents concluding remarks.

Architectural Overview

The work described in this paper builds on a previously developed method that uses artificial neural networks for image segmentation [10,14,15,16,17]. In this approach, an ANN classifies pixels individually using small neighborhoods of CT density values as input and assigns a label (e.g., "knot" or "decay") to each pixel in the image. The overall classification system consists of three modules: (1) a preprocessing module, (2) an artificial neural network module that performs tentative image segmentation, and (3) a postprocessing module. The preprocessing module separates wood from background (air), and normalizes CT density values. The ANN module labels each nonbackground pixel of a CT slice using histogram-normalized values from small windows, typically of size $3 \times 3 \times 3$ or 5×5, centered on each pixel location to be classified. In the postprocessing module, learning-based postprocessing operations are applied to remove spurious regions and refine region shapes.

Because the ANN primarily uses local information, incorrect misclassifications can occur, as described above. The example shown in Figure 1 illustrates the nature of the problem. In the figure, a sugar maple slice is processed, and the output from an ANN is shown without postprocessing. In this case, the ANN classifier is misled by density information near the center of the log. Small density changes in the center cause the ANN to label regions as heartwood instead of sapwood. Some annual rings and low density regions are incorrectly labeled as split and decay, respectively.

It is interesting to note that many of these incorrect labels have negligible effect on statistical measures of accuracy, which depend on pixel counts alone. Qualitatively, however, the removal of small regions and the refinement of larger regions can be desirable. Most of the needed refinement can be accomplished with relatively simple postprocessing steps. The difficulty lies in the development of rules that determine when to apply these simple steps. For example, the relatively large region that is classified as heartwood in Figure 1 should be removed, along with several small spurious regions. A filter that indiscriminately removes all regions smaller than some threshold will fail to remove large heartwood regions such as in this case. Because of the difficulty in developing an exhaustive set of rules that will work well for all possible situations, the emphasis of this study has been to let the machine develop its own postprocessing rules, based on observations of a human user.

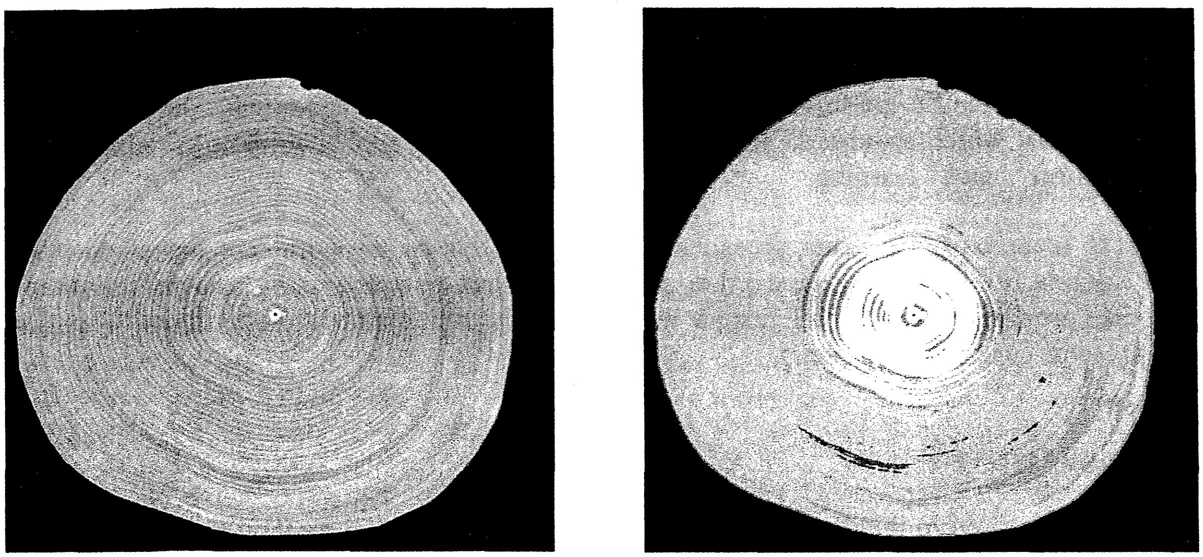

(a) (b)

Figure 1. Example of ANN-based defect detection, without any postprocessing steps. For this sugar maple slice, density information alone is not sufficient. (a) CT slice to be analyzed. (b) Output of ANN classifier.

IntelliPost: an Intelligent Postprocessing System

This study claims that the segmentation performance of an ANN-based system can be improved by applying learning-based postprocessing. The postprocessing system has been designed so that it takes human experience and preferences into account, and then stores this information for future use. Initially, the system requires manual postprocessing by a human user. This step is necessary to obtain domain knowledge and store that in a knowledge base. The system uses the knowledge base to postprocess other images in a similar way, but without user intervention.

Many machine-learning systems, including the one described here, operate in two different modes: a "learn mode" and a "run mode". During the *learn* mode, our system provides a graphical user interface that allows a human user to edit segmented images. The user selects postprocessing operations from a menu, designates portions of the image to be processed, and observes the selected refinement operation. The interactive refinement operations can continue until the user is satisfied with the resulting segmentation for any number of training images. The architecture of the *learn* mode is shown in Figure 2. This mode of the system has two main components: a postprocessing operations library that provides region-based image editing operations to the user, and a domain knowledge extraction and storage module that saves relevant information for later use.

When a user selects the *run* mode, IntelliPost automatically generates a set of rules based on its stored knowledge database. A user, possibly different than the training user, can load a new image, and the system will automatically apply its rules to the regions in the image. Based on geometric features of those regions, the system selects operations and applies them. Ideally, the system will generate a postprocessed image that is the same as the original human operator (trainer) would produce. The architecture of *run* mode is shown in Figure 3. In this mode, the postprocessing operation module is utilized by the inference engine instead of by a human operator.

We have chosen a demonstrational approach to obtain and formulate domain knowledge. This simplifies the process of creating the knowledge database, which is often a bottleneck in the design process of knowledge-based systems. Traditionally, a knowledge engineer works closely with a domain expert to develop rules and guidelines for a particular application. The rules are encoded, and an interference engine (often using MYCIN-style reasoning) is used to analyze new inputs to the system. Since this traditional process is inherently difficult and time-consuming, efforts have taken place within the artificial intelligence community to automate the knowledge acquisition process [1,7,8,9,11,12]. For our IntelliPost system, we have adopted a variation of supervised learning strategies in which the system can learn by observing a domain expert.

Many different inductive inference techniques could be adopted for our purpose, including support vector machines or explanation-based learning. For the current implementation of IntelliPost, we selected an approach based on the induction of decision trees [6]. A decision tree is a graph-theoretic tree in which each interior node represents a decision point, conceptually incorporating IF-THEN-ELSE statements, and each leaf node represents a final class label that should be assigned. Decision-tree induction is typically a supervised learning method, in which rules are generated for classifying objects that are represented as vectors from a feature space. We have developed a method in which computed region properties (along with postprocessing-related properties) constitute the feature space for decision tree induction. As a simple example, consider region size and radial distance from the center of a log slice as two features that can be computed for a given region in an image. It is possible to map any particular region onto a point in this feature space, and to assign a label indicating a desired action, such as "remove" for that region.

Using this feature-driven approach, it is possible to use information-theoretic methods to construct a decision tree that can select an action for any point in the feature space. Among the most common systems for decision-tree induction are ID3, C4.5, and CART [7,19,20]. Relatively recent systems known as MSM-T and OC-SEP [2,3,18] are based on linear and nonlinear optimization methods. The former group creates a classification tree by repeatedly subdividing the feature space using linear univariate thresholds, and the latter group creates more general nonlinear decision rules. The result is a set of separating hyperplanes that may not be parallel to the feature-space axes, with resulting subsets forming a partition of the feature space. In classical decision tree induction algorithms, entropy measures are often used to select which feature variable is to be considered at each node of the decision tree. But in the case of linear programming based decision tree induction, separating hyperplanes are used to separate the training examples at each node of the decision tree.

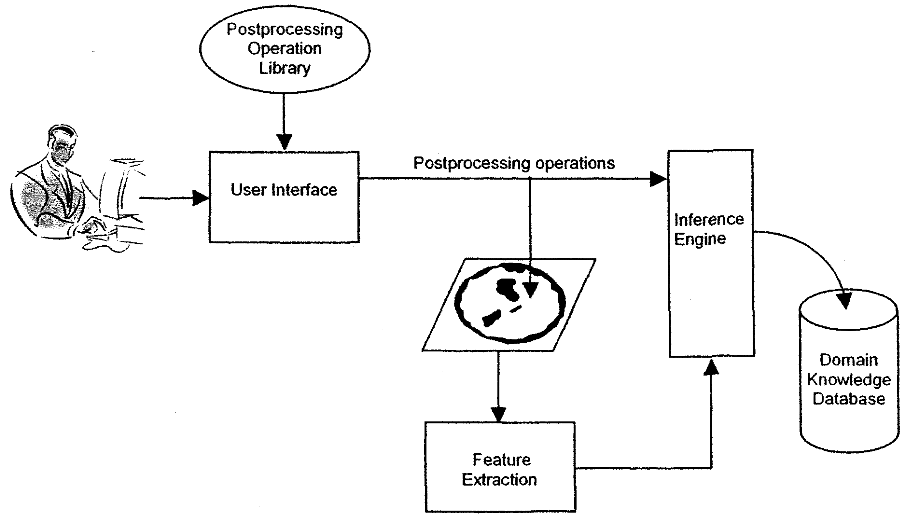

Figure 2. The architecture of IntelliPost's *learn mode*. A human operator edits a segmented CT image interactively, as the system observes and extracts information to be used later.

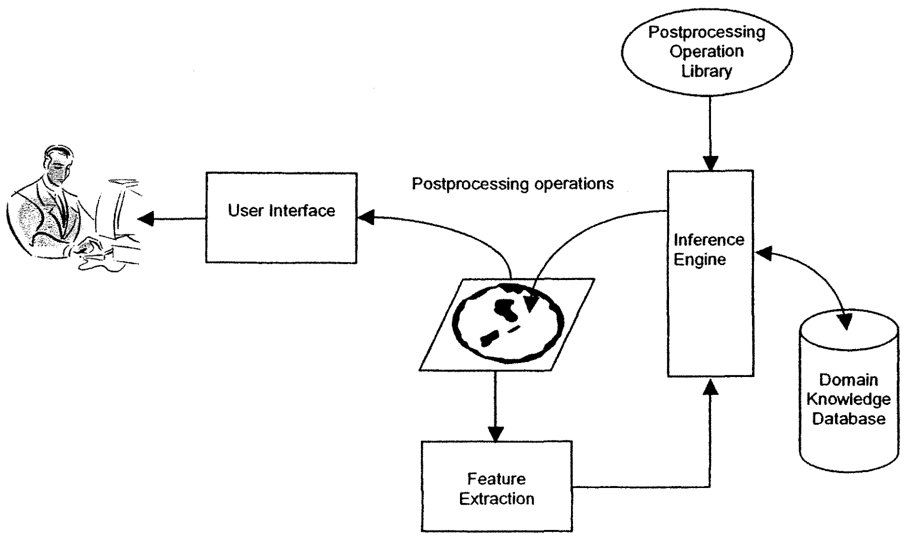

Figure 3. The architecture of *run mode*. The user provides segmented images, and the system automatically modifies the image in a manner similar to the human's earlier editing steps.

Results and Discussion

The IntelliPost postprocessing system represents a novel approach for learning postprocessing rules. Its graphical user interface, shown in Figure 4, lets the user select either *learn* mode or *run* mode for the current session. During *learn* mode, the user can select regions and modify them individually. IntelliPost supports several operations, including region removal, region thinning, and boundary smoothing. For illustration purposes, region removal is the only action considered in the example results shown here. As the user designates particular regions, the system removes them. Simultaneously, it computes and stores in the knowledge database information related to those regions and to the underlying images themselves.

When a user later invokes IntelliPost's *run* mode, the system automatically generates a decision tree using information stored in the knowledge base. It also considers each region in the given image automatically, and consults the decision tree to determine which of those regions should be removed and which should be retained.

The segmentation performance of IntelliPost has been tested using several different hardwood log datasets. This paper reports results for 2 red oak datasets that were scanned at the Virginia-Maryland Regional College of Veterinary Medicine, and for 2 sugar maple datasets that were provided by Forintek Canada Corp. An overview of the experimental setup is illustrated in Figure 5. An input CT image is processed by an initial segmentation algorithm (the artificial neural network). The output is a presegmented slice that needs refinement, and is refined by IntelliPost.

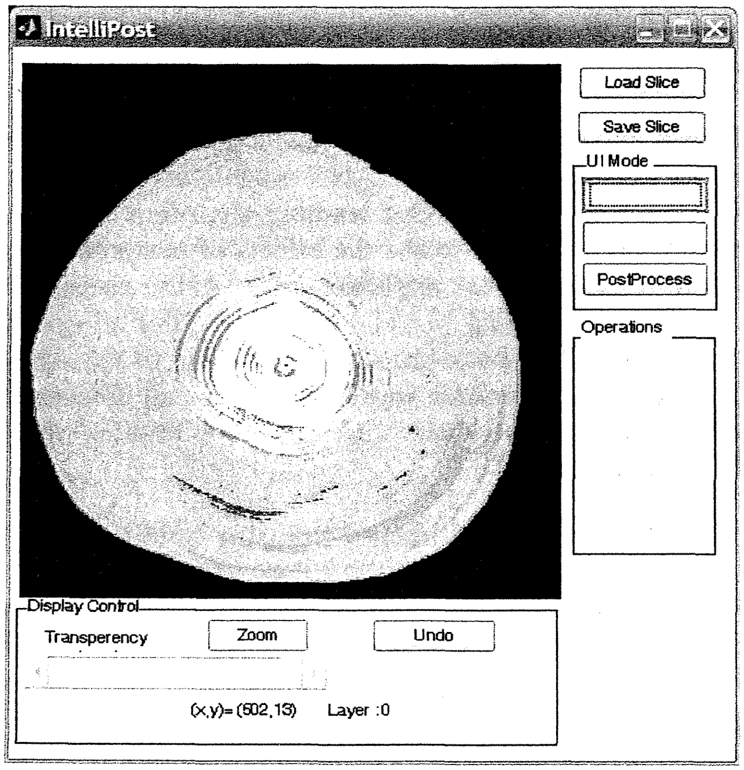

Figure 4. Example screen capture of IntelliPost's user interface. The original image and the resulting segmented image can be displayed simultaneously. The user can adjust the transparency level so that original image can be overlapped with its segmented image for comparison.

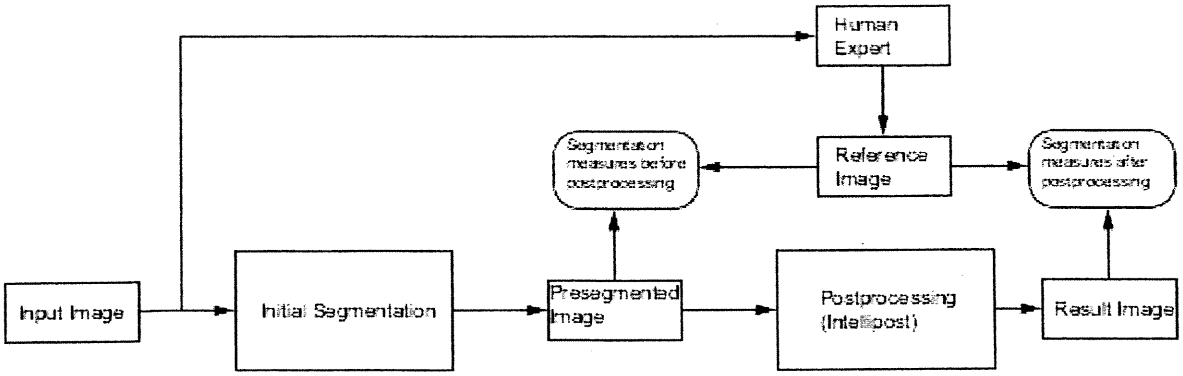

Figure 5. Comparing segmentation improvement between a presegmented image and the result image. Reference images were generated by human experts for assessing segmentation performance.

To assess the results, two USDA Forest Service researchers from the Brooks Forest Product Center at Virginia Tech provided "ground truth" information by manually delineating the boundaries of defects on given CT images. From these boundaries, segmented images were generated for use as references for comparison. Segmentation performance was measured, using these reference images, before and after refinement by Intellipost.

Each dataset was divided into two equal-sized subsets: a training partition and a test partition. Each dataset's training set was used to train the ANN module and IntelliPost independently of the others. For the ANN, the back-propagation learning algorithm was used. The ANN was slightly overtrained in each case, in order to see the effects of postprocessing better. We used MATLAB's neural network toolbox to implement the ANN classification module, and IntelliPost is currently implemented using MATLAB as well. For all datasets, the size of the ANN sample window was 5×5. For back-propagation training of the ANN, the learning rate was set to 0.2 and the momentum parameter was set to 0.8 for all datasets. Those parameters were not changed throughout the experiment. The goal for training was to reach a mean-squared-error threshold of 0.02.

After ANN training, several segmented images from the training sets were provided to a human expert for use in training IntelliPost. A USDA Forest Service researcher from the Brooks Forest Product Center at Virginia Tech used IntelliPost to postprocess those images. After these IntelliPost training sessions, IntelliPost was used to automatically refine ANN-segmented images from the test sets. These images were not used during training. Example results are shown in Figures 6 and 7. Figure 6 shows two image slices from the veterinary medicine datasets, and Figure 7 shows two image slices from the Forintek datasets. For all of these cases, removal of small regions improved the quality of the results considerably.

Quantitative results, obtained using ground-truth references, are presented in Table 1. In one case, IntelliPost correctly removed a relatively large region, and this contributed to a large improvement in classification accuracy (19.47%). For the other cases, small improvements in accuracy of about 2 or 3% resulted from the removal of small regions only. These improvements, seemingly very small, corresponded to substantial qualitative improvement.

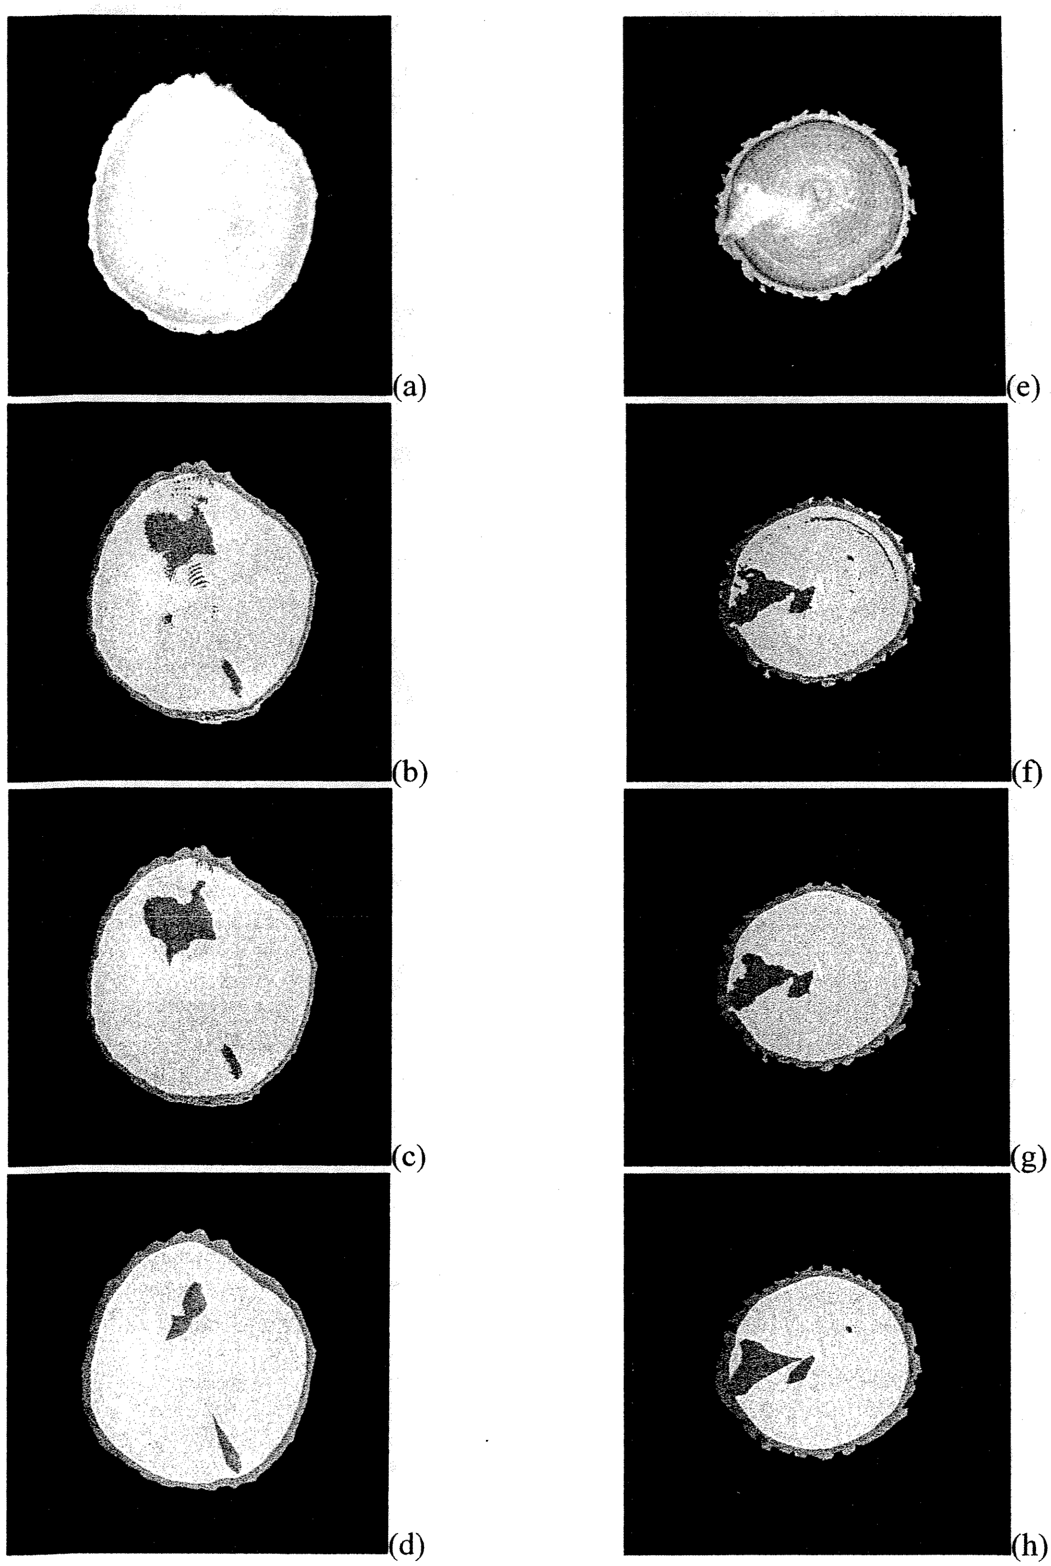

Figure 6. Classification results for two red oak log slices, taken from two different logs in the veterinary medicine dataset. (a, e) Original CT images. (b, f) Initial classifications performed by the ANN. (c, g) Automatic postprocessing results, in which many spurious regions have been removed. (d, h) "Ground truth" images that were used to assess segmentation performance.

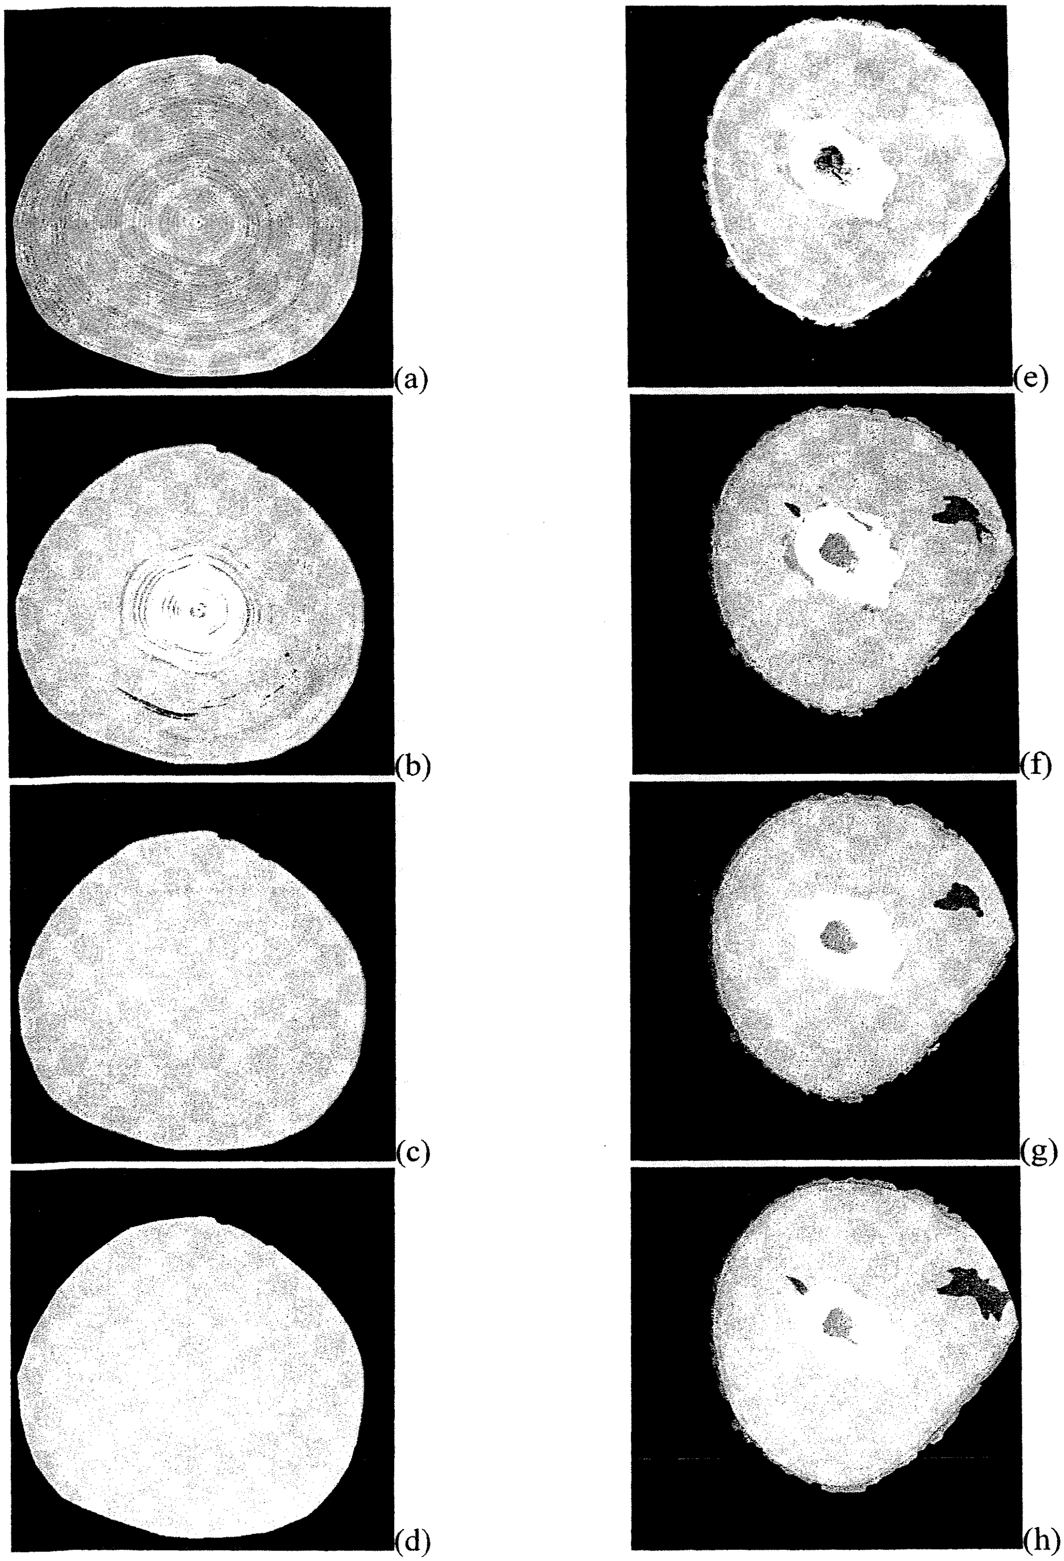

Figure 7. Classification results for sugar maple slices form Forintek dataset. (a, e) The original CT images. (b, f) Initial classifications performed by the ANN. (c, g) Automatic postprocessing results, in which many spurious regions have been removed. (d, h) Reference "ground truth" images that were used to assess segmentation performance.

Table 1. Overall correct segmentation rates for selected images in veterinary medicine and Forintek datasets.

Dataset	Slice Number	Before Postprocessing	After Postprocessing	Segmentation improvement
Forintek	Bille3-1	0.9378	0.9534	1.66%
Forintek	567b	0.8205	0.9803	19.47%
Vetmed	2049	0.9044	0.9190	1.61%
Vetmed	5357	0.9345	0.9660	3.37%

Conclusion

IntelliPost is a new tool that has been developed to assist with the detection of internal defects in hardwood logs. The system draws heavily from two different areas of research: machine learning and image analysis. One focus of this study was to explore the possibility of using demonstrational learning methods for image analysis applications. The resulting prototype is capable of observing the actions of a human operator who interactively edits a set of training images. The system then applies an automatic inferencing method to develop its own postprocessing rules based on those actions. After this learning process, the system is capable of automatically applying similar refinement steps to other images.

IntelliPost provides two modes of operation: *learn* mode and *run* mode. During the *learn* mode, the user refines a segmented image interactively, using operations in the postprocessing operations library. The system extracts high-level information from regions that are being edited by the user, and stores that information into the domain knowledge base for later use. In *run* mode, the system automatically constructs postprocessing rules by using an OC-SEP decision tree induction algorithm. After the decision tree is constructed, the system automatically refines image regions based on their geometric characteristics, as well as other high-level information that can be extracted from the image automatically. A quantitative assessment has demonstrated the viability of this approach.

Acknowledgements

This work was partially supported by USDA National Research Initiative competitive grant #98-35504-6581. We would like to thank Forintek Canada Corp. for providing CT data sets of sugar maple and black spruce. We also thank Matt Winn, of the USDA Forest Service, for providing domain expertise in this project.

References

1. Barstow, D., *A Knowledge-Based Program Construction*, North-Holland, New York, 1979.
2. Bennett, K., "Decision tree construction via linear programming," in *Proceedings of the 4th Midwest Artificial Intelligence and Cognitive Science Society Conference*, Utica, IL, pp. 97–101, 1992.

3. Bennett, K. P., and Mangasarian, O. L., "Multicategory separation via linear programming," Tech. Report No. 1127, Computer Science Department, University of Wisconsin, Madison, WI, 1992.
4. Bhandarkar, S.M., Faust, T., and Tang, M., "A system for detection of internal log defects by computer analysis of axial CT images", *Proceedings of IEEE Workshop on Applications of Computer Vision*, Sarasota, FL, pp. 258-263, Dec. 1996.
5. Burgess, A.E. "Potential application of imaging techniques to wood products," *Proceedings of 1st International Conference on Scanning Technology in Sawmilling*, San Francisco, CA, pp. VII 1-13, 1985.
6. Breiman, L., Friedman, J. H., Olshen, R. A., and Stone, P. J., *Classification and Regression Trees*. Belmont, CA: Wadsworth, 1984.
7. Fisher, D., "Knowledge acquisition via conceptual clustering", Technical Report No. 8722, Department of Information and Computer Science, University of California-Irvine, 1987.
8. Halbert, D., "Programming by example", Ph.D. dissertation, Department of Electrical Engineering and Computer Science, University of California-Berkeley, 1984.
9. Kurlander, D., "Graphical editing by example", Ph.D. dissertation, Department of Computer Science, Columbia University, 1993.
10. Li, P., Abbott, A. L., and Schmoldt, D. L., "Automated analysis of CT images for the inspection of hardwood logs", *Proceedings of the IEEE International Conference on Neural Networks*, Washington, DC, pp. 1744-1749, 1996.
11. Lieberman, H., "Capturing design expertise interactively by example," *Proceedings of the East-West International Conference on Human-Computer Interaction*, St. Petersburg, USSR, pp. 360-366, August 1992.
12. Maulsby, D., Witten, I., Kittlitz, K. and Franceschin, V., "Inferring graphical procedures: the complete metamouse", *Human-Computer Interaction*, vol. 7, no. 1, pp. 47-89, 1992.
13. Occeña, L. G., "Computer integrated manufacturing issues related to the hardwood log sawmill," *Journal of Forest Engineering*, vol. 3, no. 1, pp. 39-45, 1991.
14. Sarigul, E., Abbott, A. L., and Schmoldt, D. L., "Nondestructive rule based defect detection and identification system in CT images of hardwood logs," in *Review of Progress in Quantitative Nondestructive Evaluation* (Thompson, D. O. and Chimenti, D. E., eds.), vol. 20, pp. 1936–1943, 2000.
15. Sarigul, E., "Interactive machine learning for refinement and analysis of segmented CT/MRI images", Ph.D. dissertation, Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University, Blacksburg, VA, 2004.
16. Schmoldt, D.L., Li, P., and Abbott, A. L., "Machine vision using artificial neural networks and 3D pixel neighborhoods," *Computers and Electronics in Agriculture*, vol. 16, no. 3, pp. 255-271, 1997.
17. Schmoldt, D.L., He, J., and Abbott, A.L., "Automated labeling of log features in CT imagery of multiple hardwood species," *Wood and Fiber Science*, vol. 32, no. 3, pp. 287-300, 2000.
18. Street, W. N., "Oblique multicategory decision trees using nonlinear programming," *INFORMS Journal on Computing*, vol. 4, no. 5, 2004.
19. Quinlan, J.R., *C4.5: Programs for Machine Learning*, Morgan-Kaufman, San Manteo, CA, 1993.
20. Quinlan, J.R., "Induction of decision trees", *Machine Learning*, vol. 1, pp. 81-106, 1986.

SCANTECH 2005

THE 11TH INTERNATIONAL CONFERENCE ON SCANNING TECHNOLOGY AND PROCESS OPTIMIZATION IN THE WOOD INDUSTRY

JULY 25-26, 2005
TREASURE ISLAND HOTEL
LAS VEGAS, NEVADA

PROCEEDINGS

SPONSORED BY

WOOD MACHINING INSTITUTE

IN COOPERATION WITH

FOREST PRODUCTS SOCIETY

AND

INTERNATIONAL UNION OF FOREST RESEARCH ORGANIZATIONS