To appear in Proceedings of the IEEE Southwest Symposium on Image Analysis and Interpretation,

DESIGN OF A COMPUTER VISION BASED TREE RING DATING SYSTEM

W. Steven Conner and Robert A. Schowengerdt

Electrical & Computer Engineering
University of Arizona
Tucson, Arizona 85721 U.S.A.
conner@ece.arizona.edu
schowengerdt@ece.arizona.edu

Martin Munro and Malcolm K. Hughes

Laboratory of Tree-Ring Research
University of Arizona
Tucson, Arizona 85721 U.S.A
mmunro@ltrr.arizona.edu
mhughes@ltrr.arizona.edu

ABSTRACT

The purpose of this paper is to describe the design and implementation of a computer vision based analysis system for dendrochronology. The issues involved in the detection and analysis of tree rings are not unique to the application, but are likely of interest to anyone developing automated image analysis systems.

I. INTRODUCTION

Computer vision is a diverse research area that has applications in many fields. One scientific field which has the potential to greatly benefit from computer vision is dendrochronology, the study of tree rings. Researchers at the Laboratory of Tree Ring Research of the University of Arizona and similar research sites around the world invest many hours manually examining tree ring samples under a microscope, measuring such information as ring boundaries and widths and correlating data from different samples. This work is tedious, but is currently the main approach for collecting data.

This paper describes the computer vision issues involved in the design of a semi-automated, computerized tree ring analysis and dating system. It is based on work being done cooperatively between the Electrical and Computer Engineering Department and the Laboratory of Tree Ring Research at the University of Arizona. A brief overview of dendrochronology and the computer system being developed to automatically detect tree rings is provided, followed by a description of the computer vision techniques used to solve this problem.

This work was supported by the National Science Foundation, *Grant SBR9601867*.

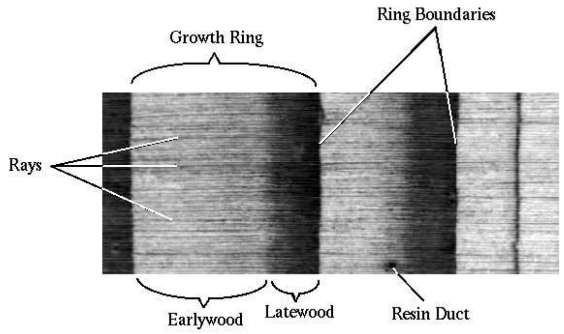

Figure 1: Major Features of Conifer Tree Rings.

II. BACKGROUND

A. Dendrochronology

Dendrochronology is the science of tree ring dating. During each year in the lifetime of many trees, a single tree ring is created. Through the techniques of dendrochronology, it is possible to assign each tree ring in a sample to a specific calendar year. The innermost ring at the center of a tree cross-section corresponds to the first year in the tree's life, while the last ring at the tree's outer edge corresponds to the last year of growth.

The most common features in the image of a cross-section of a tree are summarized in Figure 1. In this image, tree rings run from the top to bottom along primarily vertical lines. These are the edges that must be detected in the image, while ignoring other extraneous features. Resin ducts, which may occur anywhere in a sample and are often circular in nature, can be distinguished from near-linear tree rings without too much difficulty simply by their physical characteristics. However, lines that run from the center of a tree outward toward the bark, referred to as "rays", and aligned cell boundaries are also nearly linear, and therefore less easily distinguished from rings.

Having identified the rings in a tree sample, it is necessary to find the features of individual rings. Ring widths and other characteristics may be used to match ring patterns from one tree to another, allowing cross-dating to be performed.

B. Overview of Tree Ring Dating System

The computerized tree ring dating system consists of a digital camera attached to a microscope used for imaging a cross section of a tree that is situated on a positioning table under the microscope. The camera, microscope, and positioning table are all controlled by a Sun ULTRA 1 workstation.

To begin tree ring analysis, a series of overlapping images are acquired across the tree sample, from the center of the tree (the pith) to the outer edge (the bark). These images are normalized for any variation in camera gain during capture and are mosiaced together into a single sample image. At this point, a region of interest window traverses the image data, detecting tree rings and performing measurements on them. Once ring characteristics are determined, ring patterns from one sample may automatically be matched to patterns from another sample.

At the time of writing, the computer vision problems described in this paper are work in progress. The algorithms and techniques included here represent the best methods found to-date to solve the computer vision problems of this project.

III. IMAGE PREPROCESSING

Before attempting to detect tree rings, two important pre-processing steps must be performed on the image data. First, the individual image frames must be mosaiced into a single image representing the entire sample that is being studied. Next, gain adjustments must be made to ensure contrast is normalized between frames.

During image capture, it is important to ensure a minimum amount of overlap between adjacent frames. This is necessary to allow frames to be matched to one another by their overlapping edges. A straightforward technique to mosaic a series of captured images is to present the images to an analyst in pairs, in the order in which they were captured. The analyst could then indicate, with a mouse or similar pointing device, the locations of common features in each of the two images. Auto correlation techniques for image registration may then be applied to identify the "best" relative shift positions between the two images, as described by Schowengerdt, [6, section 8.3.1].

Due to the electronic properties of frame-grabber cameras, it may be necessary to perform contrast adjustments between frames to avoid artificial lines at frame overlap lines that may be confused as tree ring edges by edge detection algorithms. As long as the variations between frames is linear, it is possible to adjust their relative contrast by computing image statistics for the overlapping boundaries. The appropriate gain adjustment between two frames a and b is: $gain = deviation_{overlap_a} / deviation_{overlap_b}$. The bias adjustment between the same two frames is then: $bias = mean_{overlap_a} - (gain)(mean_{overlap_b})$. Once the relative contrast adjustments have been determined between each pair of captured images, one image may be used as a reference and each remaining image may be adjusted through a simple linear contrast stretch.

IV. RING DETECTION

A. Orientation Tracking

In order to optimize the tree ring detection process, it is useful to know the general orientation of the rings ahead of time. This information can be used in an algorithm to distinguish true tree rings from other features. As described earlier, the rays in a tree cross-section may easily be confused for tree rings. The main distinguishing factor between tree rings and rays is the fact that they are nearly orthogonal to one another. Therefore, if the orientation of the tree rings within a region of interest is known, it becomes straightforward to distinguish the two features.

Tree ring orientations are found by computing the gradient direction at each pixel location within a region of interest via a directional filter such as Sobel or Roberts, and quantizing the results into one of eight sectors. The gradient magnitude can then be used to compute the average magnitude for each direction sector. The direction with the largest average magnitude corresponds to the orientation that has the highest probability of being the ring orientation in the region of interest.

B. Modified Canny Edge Detection

After performing experiments with various gradient filters, zero-crossing filters, and other common edge detection techniques, the Canny edge detector was found to yield the best overall results for detecting tree rings for different types of wood.

The non-maxima suppression feature of the Canny edge detector allows a one-pixel wide edge to be detected. As described by Jain, et. al. [4, section 5.6.1],

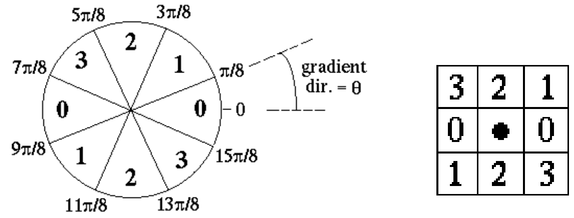

(a) (b)

Figure 2: An example of sector definitions for the Canny edge detector. (a) Gradient Directions and their Corresponding Sectors. (b) Sectors Corresponding to Each 8-Connected Pixel Neighbor.

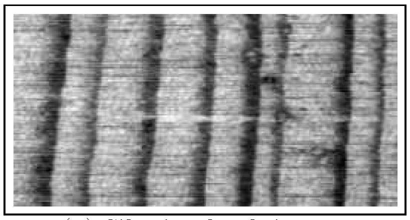

(a) Siberian larch image.

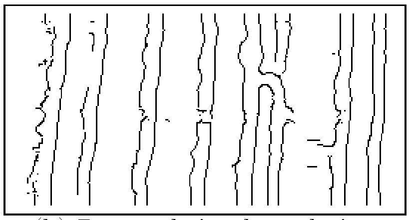

(b) Detected ring boundaries.

Figure 3: Example Canny edge detector output, using an appropriate threshold. Problems include double edges and horizontal intra-ring connections.

a simple version of the Canny edge detector can be implemented by convolving an image with a Gaussian smoothing filter, computing the gradient magnitude and direction at each pixel location using two orthogonal 2x2 first-difference approximations for the x and y partial derivatives, and performing rectangular-to-polar conversion on the results. Non-maxima suppression is then performed by quantizing the gradient direction of each pixel into one of four sectors, shown in Figure 2. The magnitude of each pixel is compared to the two 8-connected neighbors in its sector. If the pixel's magnitude is not greater than that of both of the neighbors of interest, it is set to zero in the output image. Otherwise, its magnitude is used as the value for the corresponding pixel in the output image. This technique results in edge boundaries that are one pixel wide and represent the location of the maximum gradient magnitude in each edge.

The Canny edge detection results for an example tree ring image are shown in Figure 3. Although the Canny edge detector produces very accurate tree ring boundaries, it also finds two edges: one at the correct boundary transition between latewood and earlywood, and the other at the intra-ring transition between earlywood and latewood. Also, several cases exist where one ring is connected to the next by edges that run along prominent rays and other horizontal features.

Two types of *a priori* knowledge are useful at this point. First, if the overall orientation of tree rings within a region of interest is known, as described earlier, cross-ring connections can be eliminated by suppressing edges that are nearly orthogonal to the orientation of tree rings. Second, if the locations of the pith and bark are known, it is possible to suppress any edges that are not a transition from latewood to earlywood. Modifications to the Canny algorithm will be described with the assumption that tree rings in the region of interest are nearly vertical, with the pith to the left of the region of interest and the bark to the right. Similar modifications can be made to detect tree rings with any of the four orientations shown in Figure 2.

The directional modifications to the Canny edge detector algorithm are primarily at the non-maxima-suppression stage. Assuming that within a region of interest the tree rings are approximately vertical, edge detection should be optimized to detect vertical edges. This is accomplished by comparing the gradient magnitude of each pixel with its left and right neighbors, regardless of the gradient direction of the pixel. If the gradient magnitude is not larger than either of its horizontal neighbors, then it will be suppressed as a non-maximal edge point. In order to eliminate double edges at each ring, the greylevel value of the right neighbor is compared to the greylevel value of the left neighbor. Assuming that the pith of the wood is to the left and the bark is to the right, the transition from latewood to earlywood is represented in the image as a transition from dark pixels to light pixels during left-to-right traversal across the image. Thus, if a given pixel has a local maximum gradient magnitude, and its right neighbor has a higher greylevel value than its left neighbor, the pixel is considered to be a tree ring edge and the value of the corresponding output pixel is set to the gradient magnitude at that location. Otherwise, the output pixel value is set to zero.

Two examples of the modified Canny algorithm results, using an appropriate threshold, are shown in

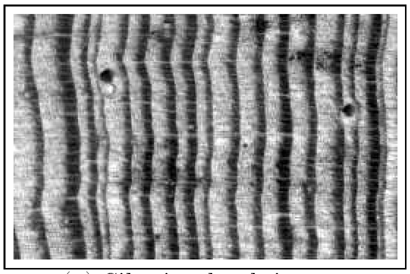

(a) Siberian larch image.

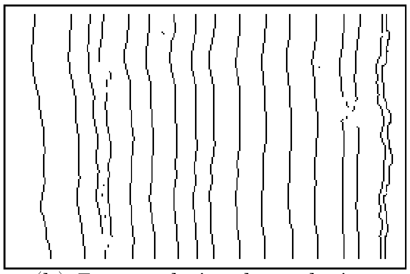

(b) Detected ring boundaries.

Figure 4: Example ring boundaries from modified Canny Algorithm for a Siberian larch sample.

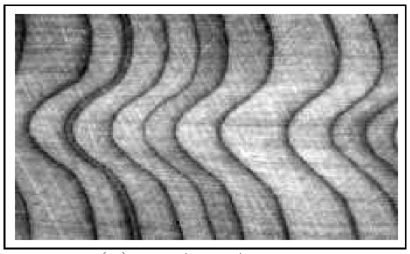

(a) Juniper image.

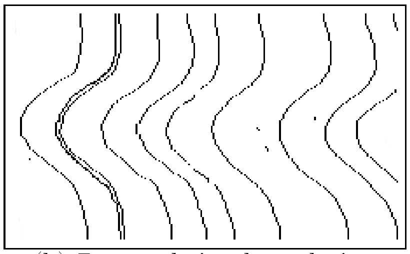

(b) Detected ring boundaries.

Figure 5: Example ring boundaries for a Juniper sample from modified Canny Algorithm . Note that these results were obtained assuming that tree rings are nearly vertical in the image. This edge detection technique is robust enough to maintain good edges, even in the areas where the rings are very curved. Even better results might be obtained by using a smaller region of interest and optimizing edge detection in the directions of Sectors 1 and 3 (Figure 2) in the curved regions.

Figures 4 and 5. Clearly, the problem of double edges has been eliminated, and although a few breaks exist in some of the ring boundaries (especially in the areas where the ring orientations diverge from the assumed vertical orientation), they are largely connected from the top of the image to the bottom.

Since it is not necessary to examine the gradient direction at any stage in the modified Canny algorithm, this computation might be eliminated to increase efficiency. The algorithm could be further optimized by eliminating the computation of the gradient magnitude all together by simply computing the first partial derivative at each pixel in the assumed direction of tree ring orientation using simple directional filters. However, after experimenting with this method it was found to yield less than satisfactory results in many sample images, due to highly broken ring edges. In general, the computation of the gradient magnitude provides a relaxation of assumed ring orientations, allowing more deviation from the assumed ring orientation than if a directional first derivative is used.

C. Edge Linking

Although a modified version of the Canny edge detector produces good tree ring edges across varying types of wood samples, the detected ring boundaries tend to be broken in several locations. In addition, unwanted artifacts resulting from knots, rays, resin ducts, and other troublesome features in the wood are present. Post-processing is necessary to successfully link tree ring fragments and to remove unwanted noise from the edge map.

For the first step of post-processing, each edge fragment in the image is assigned a unique label via an 8-connected components labeling algorithm, [4, section 2.5.2]. Once the edge fragments have been labeled, it is necessary to determine which fragments belong to a given ring and which edge fragments do not belong to any rings at all and should simply be discarded.

One method to eliminate the majority of the unwanted edge fragments is to apply a size filter to the labeled edge map. However, in some cases curved edges will tend to be fragmented into numerous edge fragments and may be removed by the filter.

A robust ring linking technique utilizes information about the physical structure of tree rings to add a certain degree of intelligence to the process. One of the most basic features of tree rings is that the width of a given ring relative to the width of neighboring rings remains nearly constant along the length of the ring. This allows a measure of certainty to be obtained as to whether various edge segments belong to the same tree ring or not.

The method currently used to link tree ring edge segments begins by partitioning regions between known rings. If an edge exists that is fully connected from the top to the bottom, then it is considered to be a complete ring. By the nature of the edge detection technique used to detect vertical tree rings, only one pixel per line will exist on a given edge segment. Therefore, there is a good probability that the average number of edge pixels per line between two known rings, rounded to the nearest integer, corresponds to the number of tree rings that exist between the two known rings. As an illustration, take rings two and five in Figure 4 as the two known rings. If the number of edge pixels between the two rings is counted for each line, some will have none, some will have two, some will have four, etc. However, the average number will be two, which is indeed the number of tree rings between the two known rings.

Once the number of tree rings has been determined between the two complete rings, it is possible to compute the most common locations of the ring edges. In the example described above, the positions of each edge pixel between the two rings can be computed by dividing the edge pixel's distance from the left ring by the distance between the two rings on a given line. The resulting percentage may then be quantized into one of say 100 integer values from zero to 99%. If an array of 100 bins is used to count the number of occurances of each edge pixel location between the two known rings, the two bins with the largest count represent the locations of the two tree rings between the two know rings. Using this information, it is possible to label any edge segments that are approximately located at a given estimated ring location as belonging to that ring. Any edge segments that do not correspond to one of the estimated ring locations may be discarded as unwanted noise segments.

The main goal of this project is to provide dendrochronologists with a tool that improves the efficiency of their work, while still maintaining a high degree of accuracy. In some cases, especially near the edges of a tree ring sample, or on low contrast rings that are fragmented with large gaps, it may become necessary for an analyst to intervene. In this case, the analyst might be allowed to either manually fill in the gaps of the tree ring edge by drawing it with a mouse, or a semi-automatic edge tracking algorithm may be applied to search for a tree ring within a boundary constrained by the user.

V. RING MEASUREMENTS

Once a clean tree ring edge map has been produced, the next step in the semi-automated tree ring analysis system is measuring the rings in the image. The most common tree ring features measured in dendrochronology are ring widths. Assuming that each ring is uniquely labeled and fully connected from one edge of a region of interest to another, ring widths can be obtained by simply computing the average distance between each pair of ring edges. Since the *relative* width of each ring to its neighbors is of far more importance in dendrochronology than *absolute* widths, the units of measurement are not critical. However, since the orientation of tree rings often varies slightly across a region of interest, it is important to pay attention to the direction of width measurement at each point along a ring to ensure that meaningful width averages are obtained.

In general, the gradient direction at a given ring edge location is not a sufficient representation of the normal to the ring curve because it is sensitive to any dips or bumps in a ring edge. In fact, it is possible for the gradient direction to misrepresent the normal to a tree ring by as much as ninety degrees or more! A superior method of determining the normal direction to a tree ring edge at a given location is to fit a second order polynomial to the ring edge data in a neighborhood around the location of interest. The number of edge pixels used to compute an approximating polynomial curve must be large enough to eliminate small twists and curves in the ring, but small enough to maintain an accurate representation of the local orientation of the ring. Based on the magnification used to capture images, ten to twenty pixels seem to be convenient and sufficient to represent local tree ring edge curves for this application.

Once the normal direction to a ring edge has been determined, the distance in this direction can be computed between the current ring and the ring to its right. The easiest way to compute this distance is to first travel along the two sector directions (Figure 2) that are closest to the normal direction at the current ring location, until the next ring is reached. The ring width at a given point is computed by interpolating the distance in the normal direction from the digitally measured distances in these two sector directions.

In addition to ring widths, other features such as the average greyscale profile between two ring edges can also be measured and recorded. Also, the analyst may annotate additional features for unusual rings, such as frost and micro rings. All of this information is useful when cross-dating between tree ring samples.

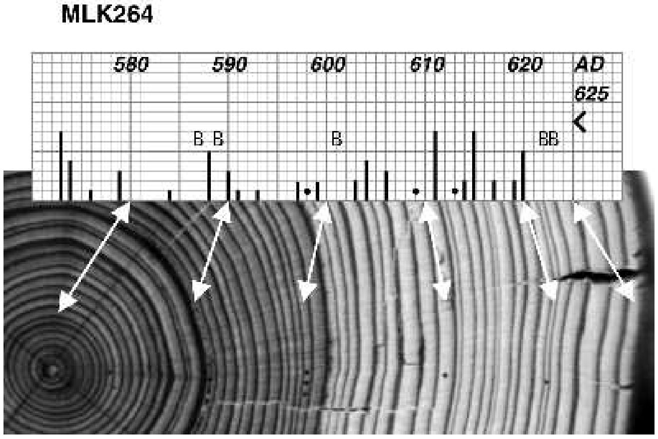

Figure 6: Sample Skeleton Plot.

VI. DATA REDUCTION AND CROSS DATING

The primary purpose of tree ring analysis is to allow cross-dating between samples. This is possible due to the fact that one ring exists for each year in the lifetime of a tree, with the outermost ring representing the last year of growth for the tree. Varying widths in trees are caused primarily by changes in the climate, and thus the same patterns of tree ring widths can often be observed across different trees in a given region. One example of cross-dating involves the matching of tree ring patterns between live trees and trees that were cut many years in the past. This allows dendrochronologists to determine exactly when the tree was cut down, which has many applications in archaeological studies.

The method of cross-dating used in the Douglass method of dendrochronology involves the creation of "skeleton plots" from ring measurements. These plots are simply a way of representing tree rings in a compact, normalized form. An example of a skeleton plot is illustrated in Figure 6. The issues involved in computerized generation of skeleton plots are described by Cropper, [2].

Once a skeleton plot has been generated to represent tree ring widths and other features, it must be compared to skeleton plots from other samples. Skeleton plot patterns that match between samples indicate a high probability that they represent identical calendar years in the lifetimes of the trees from which they were obtained. This type of cross-dating can be performed in a fairly straightforward manner using pattern matching algorithms. One technique to perform computerized cross-dating is described by Munro, [5].

VII. SUMMARY AND CONCLUSIONS

Tree ring analysis has many applications in fields ranging from climatology to archaeology. However, many hours of tedious work by experts in dendrochronology are required to accurately analyze each tree ring sample. This paper describes computer vision techniques that may be applied to perform analysis of tree rings in a semi-automated manner. By adding directional intelligence to the Canny edge detector, it has been shown that accurate representations of tree ring boundaries can be created with a high degree of accuracy. Adding fragment linking techniques based on the physical structure of tree rings allows complete representations of tree ring boundaries to be obtained, while removing remaining noise.

Although the methods described here produce acceptable ring boundary results in many cases, it is still desirable to retain the option of analyst intervention during the process. When emulating complex human vision analysis, as in the case of tree ring detection, computer vision serves the analyst best as a computer-assisted environment.

VIII. REFERENCES

[1] Canny, John, "A Computational Approach to Edge Detection." *Readings in Computer Vision: Issues, Problems, Principals, and Paradigms*, ed. Martin A. Fischler and Oscar Firschein. Los Altos, CA: Morgan Kaufmann Publishers, Inc., 1987.

[2] Cropper, John Philip, "Tree-Ring Skeleton Plotting by Computer." *Tree-Ring Bulletin*, Vol. 39, 1979.

[3] Douglass, A.E. "Tree Rings and Chronology." *University of Arizona Bulletin* 8, no. 4. 1957.

[4] Jain, Ramesh, Rangachar Kasturi, and Brian G. Schunck, *Machine Vision*. New York, NY: McGraw-Hill, Inc., 1995.

[5] Munro, Martin A. R., "An Improved Algorithm for Crossdating Tree-Ring Series." *Tree-Ring Bulletin*, Vol. 44, 1984.

[6] Schowengerdt, Robert A., *Remote Sensing: Models and Methods for Image Processing.* Second Edition. San Diego, CA: Academic Press, 1997.

[7] Stokes, Marvin A. and Terah L. Smiley, *An Introduction to Tree-Ring Dating.* Chicago, IL: University of Chicago Press, 1968.