Histogram Stretching Or Histogram Equalization In Image Processing

William R. Oliver
Armed Forces Institute of Pathology, Washington, DC

Introduction

A few weeks ago, a person posted an interesting question on an internet microscopy mailing list: what is the difference between histogram stretching and histogram equalization when applied to microscopy images? The following is a short intuitive review which compares the two. Since this is a short and general description, I will gloss over some details and make some generalizations which may not be true in all cases and implementations.

The first thing to remember is the basic purpose of contrast enhancement. The idea is simple. In a grayscale (black and white) image you are simply trying to take two levels of gray that are close together, and thus visually similar, and move them apart so you can better see the difference between them. You can think of grayscale values in an image as beads on a string. Contrast enhancement simply moves all or some of those beads farther apart.

In histogram stretching, you are simply stretching the outliers of the pixel values to fit the dynamic range of the representation, and scaling everything in between accordingly. The *relative* position of the pixel intensity values will not change. With histogram equalization, you change the relative positions in order to distribute the values more evenly within the representation. Histogram equalization does not *necessarily* imply a histogram stretch as well, though most implementations go ahead and do it.

Histogram stretching and equalization

Let's look a bit more closely at these two techniques. Consider a histogram that looks like this:

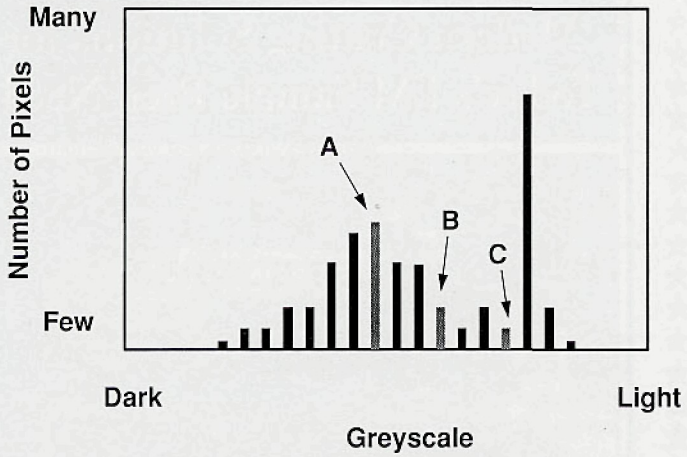

Figure 1: An Unmodified Histogram

Note that the distance between columns A and B is pretty much the same as the distance between columns B and C.

Let's assume that the brightness level denoted by "Dark" represents the minimum value that can be displayed and "Light" represents the brightest level that can be displayed. The distribution of pixels in Figure 1 clearly does not use all the available levels of brightness. The simplest thing to do would be to drag the lowest value of the histogram down to "Dark" and the brightest value of the histogram up to "Light" and scale everything else in between.

Using the beads on a string analogy, all the beads are clumped together in the middle of the string, with lots of string on either side of the clump. An easy way to un-clump them would be to use all the string. That's all a histogram stretch does. The absolute positions of the beads on the string change, but the relative positions do not. So, a histogram *stretch* performed on the histogram in figure 1 would result in something like this:

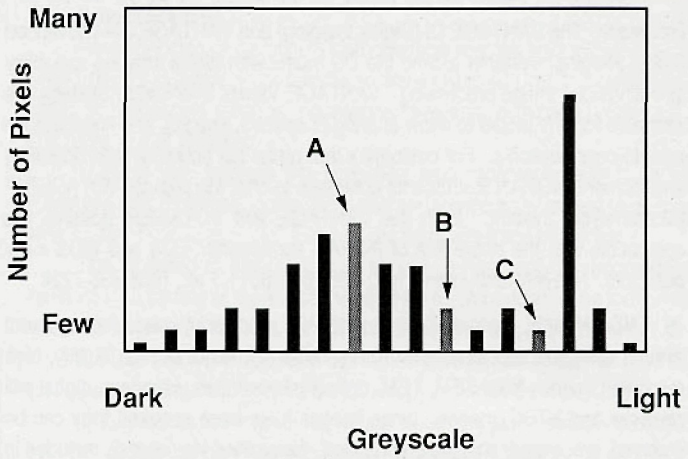

Figure 2: The Results Of Histogram Stretching

Note that the columns denoted by A, B, and C, while farther apart, are still roughly equidistant. As you might expect, there are lots of variations on this theme. You might divide the distribution into regions and scale each one a little differently. You might place barriers that can't move, and so on. Each variation may or may not be helpful, depending on the particular set of images with which you tend to work.

Histogram *equalization* recognizes that by some statistical measures the greatest overall contrast enhancement (in terms, for instance, of information content) is achieved if all the intensity values had the same number of pixels. This ideally results in a flat histogram. This assumes, by the way, that all intensity values are of equal importance – that light stuff is just as important as dark stuff. Using the beads on a string analogy, histogram equalization attempts to distribute the beads evenly along the string. Once this is done, you can't separate any two beads any more without forcing some other beads closer together. Overall, that's the best you can do.

But there's a problem. Looking at the distribution in Figure 2, we want to spread the pixels in the tall column (next to column C) over three or four spaces. However, there is no way to tell which pixels in that column should be a little brighter and which should be a little darker. So, instead, we simply leave those nearby columns empty, and push everything away from that tall column. The result is basically to squeegee the values with low numbers of pixels together and separate the ones with large numbers of pixels. If you take the histogram-stretched distribution above and do a histogram equalization on it, you might get something like Figure 3. This is done by eye, so really doing it would likely be a little different. I will show a real example below.

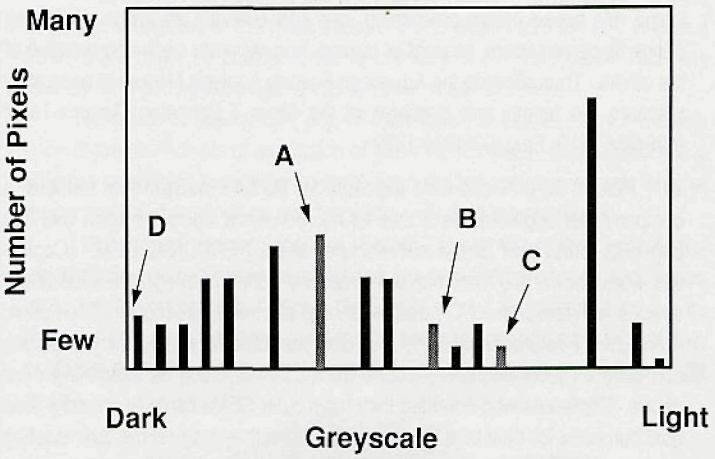

Figure 3: The Results Of Histogram Equalization

Note that columns A and B are now farther apart than columns B and C. The order of the beads on the string are retained. A is still darker than B and B is still darker than C. However, both the absolute and relative positions are lost. Some

implementations even sacrifice small distinctions containing few pixels in order to provide more space for taller columns in the histogram. In Figure 3, for instance, the three darkest small columns of Figure 2 are combined to form column D.

Local or Adaptive Histogram Equalization

Finally, there is a method called "local" or "adaptive" histogram equalization in which equalization is not done for the entire image, but only within small local regions (sometimes called "adaptive regions"). Some algorithms divide the image up into lots of little subimages, and some just look in the region around each pixel. In either case, an equalization is done using only a part of the image. Thus, a dark area in one part of the image will not "use up" all the dark pixels in a bright area. This is useful in the processing of medical images such as Computerized Tomography (CT) images.

In CT images, bone is very light and tissue such as lung is very dark. If you can do contrast enhancement on the bone and the lung independently, then you can present much more information in one image. Otherwise you have to have one image that shows detail in bone with the lungs all black and another image that shows details in the lung with the bones all white.

As you might expect, there are lots of parameters to play with in order to optimize results. Parameters that affect the result include the size and shape of the adaptive regions, whether or not bordering adaptive regions should have any effect, whether or not contrast should be limited to minimize some noise measure, and others. In cases of adaptive histogram equalization, the global ordering of the intensity values is lost, but is retained locally.

"Ringing" of various sorts is a common artifact in this method. In Figure 7, the adaptive regions are about the size of the largest nucleus. Since the large nuclei take up an entire adaptive region, their contrast is optimized. Small nuclei, on the other hand, are contained in adaptive regions which also contain background or cytoplasm, and their distributions are affected by those features. If I had made the adaptive regions smaller the smaller nuclei would have shown more contrast, but I would have introduced other artifacts.

Color Images

Color images are like grayscale images, only more so. Instead of a single axis going from dark to light, you have an axis for each primary color. The most common representation in home computers is to divide color into red, green and blue axes. However, there are many other ways of representing colors, called "color spaces," and the same algorithm applied to different representations can have very different results. Instead of having grayscale values lined up like beads along a string, you have a cluster of color values floating in a three-dimensional

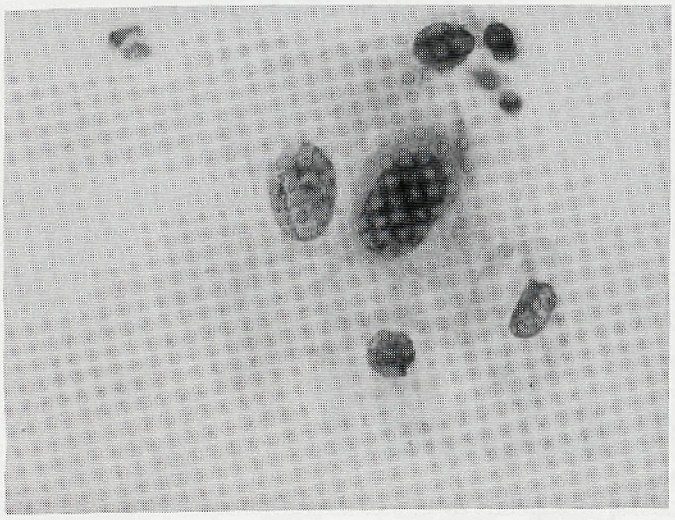

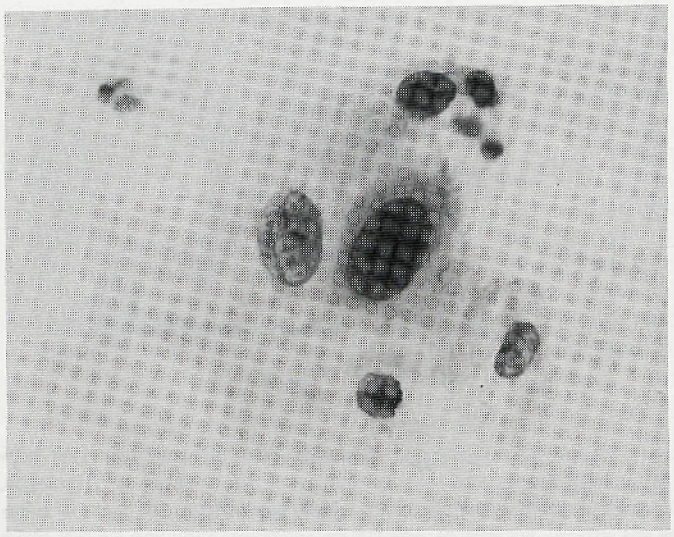

Figure 4: An image as digitized. Note that the lightest areas are still a little dark and the darkest areas are not completely black.

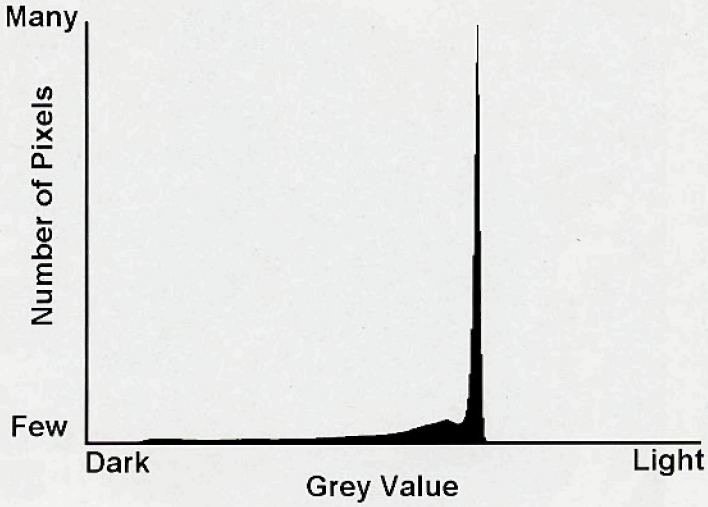

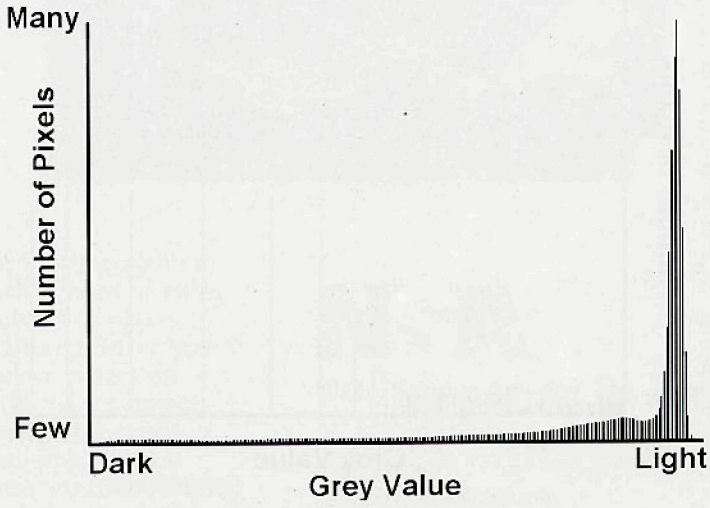

Figure 5: This is a simple histogram stretch. The brightest areas are now white and the darkest areas are now black. Everything in between is scaled. Notice the introduction of spaces between the columns in the histogram, as they are spread apart. Normal directly digitized images usually don't have gaps such as this. This kind of sawtooth pattern on a histogram is a hint that the image has been manipulated.

Continued on following page

Histogram Stretching Or Histogram Equalization In Image Processing - Continued from preceding page

color space. The extension of contrast stretching into three dimensions is straightforward: you blow up the cluster like a balloon.

Contrast equalization in three dimensions is not quite as straightforward. One simple compromise is to do equalization on each axis independently. You break the color image into three grayscale images, one for each primary color or color space dimension, and perform an equalization on each image. When those images are recombined, you can get some useful contrast enhancement that isn't too jarring to one's color sense. Other methods work directly in three dimensional space, forcing the points in color space to diffuse or repulse away from each other.

An Example

Figures 4 through 7 are examples of grayscale contrast enhancement on a 63x Papanicolau preparation of cervical cells. Figure 4 shows the original image and its resultant histogram. Figure 5 shows the result of a contrast stretch. Figure 6 shows a histogram equalization. Figure 7 shows an adaptive histogram equalization. Note that the histograms in these figures all have the same scales. The histogram in Figure 7 really does show many fewer pixels distributed over many more levels of gray.

There are a number of good image processing textbooks available. "The Image Processing Handbook," 2nd ed. by John Russ provides a good general introduction without getting bogged down in the mathematics or implementation trivia. I found the Russ book good for providing a conceptual understanding of what's happening in some of these methods, but I don't think I would find it useful as a guide for writing code. "Digital Image Processing" by Rafael C. Gonzalez and Richard E. Woods provides a bit more mathematical background, but is still quite readable. Kenneth Castleman recently published a new version of his book, "Digital Image Processing," which provides an excellent mathematical foundation and a good general introduction.

Disclaimer: The opinions expressed are those of the author and do not necessarily represent those of the government of the United States, the Department of Defense, or the Armed Forces Institute of Pathology. ■

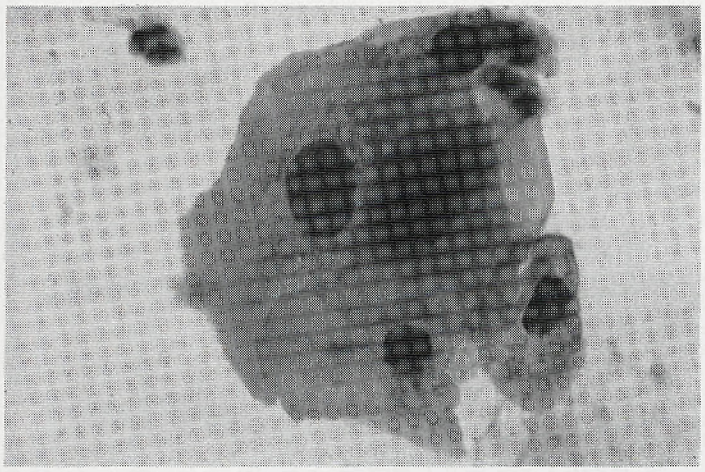

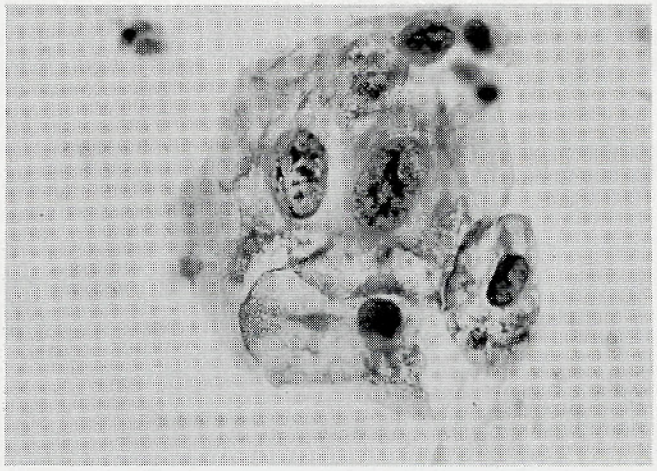

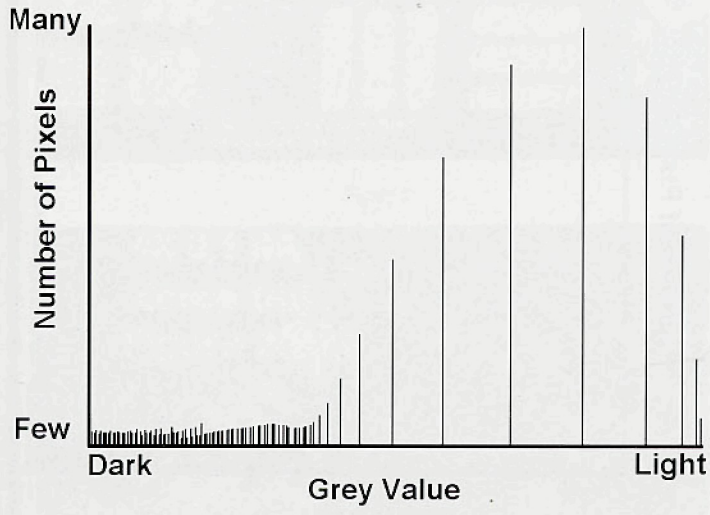

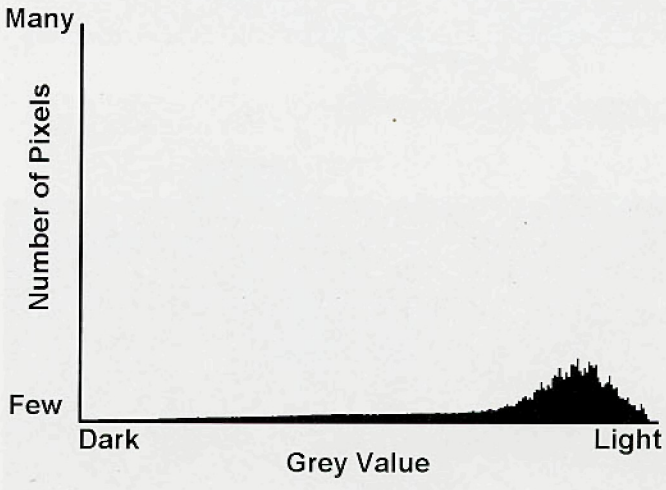

Figure 6: This is the result of a histogram equalization. In this implementation, some dark contrast has been sacrificed to provide more contrast in bright areas, which have more pixels. This method of contrast enhancement on this class of images might be useful if one were interested in looking at cell boundaries, but clearly makes nuclear detail less obvious. This is a result of the assumption that all pixels are equally important. In this image, cell detail is lost to provide information about background noise. In other images, of course, finding details in a background may be exactly what you want.

Figure 7: This is an adaptive histogram equalization. Since each small area of the image has its own histogram, the nuclei can take advantage of more levels of gray. This method, while sometimes providing dramatic results, is highly parameterized and can result in very distracting artifacts. For instance, note that the background away from the cells shows enhanced noise, while the background near the cells do not. This is because the small adaptive regions which include parts of both background and cells will have a very different distribution than adaptive regions which have only background or only cell.

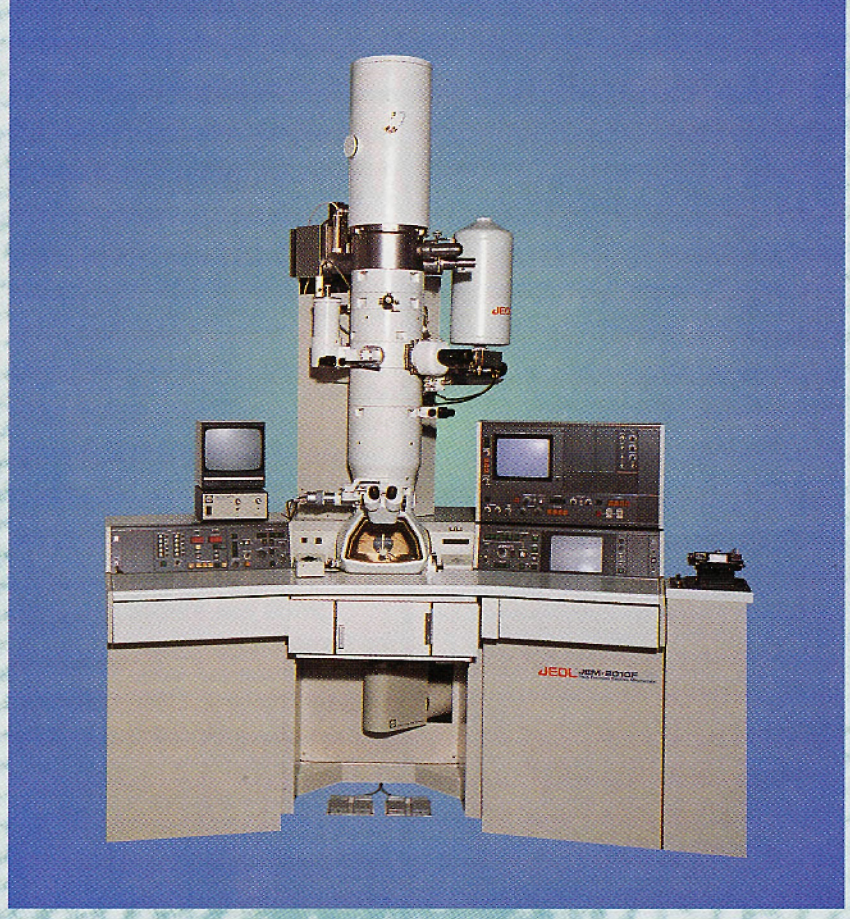

Sights Unseen.

With the New JEM-2010F **Field Emission** Electron Microscope You'll See It ... If It's There.

From JEOL...the newest generation of precision equipment that permits unprecedented resolution in 3-dimensional, subnanometer analysis of microstructures. Featuring user-friendly operation and long-term stability, the JEM-2010F also offers:

- Schottky Emission: High Current High Brightness
- High Probe Current: 0.5nm Probe with 100 pA Current
- High Resolution: Information Limit 1.4Å, Scherzer 1.9Å
- Holography: Option Available
- STEM Resolution: 0.2nm
- Magnification: 8MX

Discover the JEM-2010F and visit sights previously unseen.

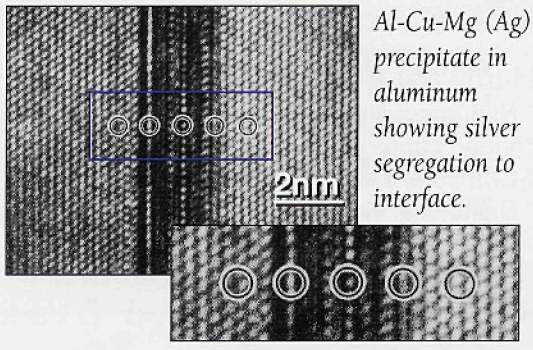

Al-Cu-Mg (Ag) precipitate in aluminum showing silver segregation to interface.

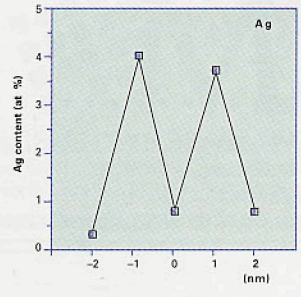

Data courtesy of Dr. James M. Howe, Department of Materials Science & Engineering, University of Virginia, U.S.A.

JEOL
Serving Advanced Technology

JEOL USA, Inc., 11 Dearborn Road, Peabody, MA 01960
Tel: 508-535-5900 Fax: 508-536-2205 e-mail: eod@jeol.com

Circle Reader Inquiry #4

Detection of *Ophiostoma Piceae* in Radiata Pine using Immunofluorescence Labeling and Confocal Laser Scanning Microscopy

Ying Xiao[1], Bernhard Kreber[1] and Colette Breuil[2]

1. *New Zealand Forest Research Institute Limited, Rotorua, New Zealand*
2. *Faculty of Forestry, University of British Columbia, Vancouver, Canada*

The use of fluorescence microscopy to investigate fungal growth in wood often causes interference due to the strong autofluorescence of wood lignin, unless fluorescent probes which specifically react to fungal hyphae, are used. Techniques to enable differentiation of hyphae from wood have been recently reported (Singh *et al.*, 1997; Xiao *et al.*, 1997). The authors demonstrated that while glutaraldehyde can be used to detect fungal native proteins, wheat germ agglutinin (WGA) which reacts with cell wall chitin of hyphae, considerably improved detection of fungi growing in wood.

Confocal laser scanning microscopy (CLSM), a new technique in wood biodeterioration research, has recently been employed because it produces blur free images and allows optical sectioning across the thickness of the specimen (Xiao *et al.*, 1998). CLSM also enables multi-channel, fluorescent imaging which can be monitored on a big screen and controlled by computer.

In the current study, an immunofluorescence technique was developed to detect *Ophiostoma piceae*, a common sapstain fungus in New Zealand, in radiata pine wood using a monoclonal antibody, 1F3(1)), and CLSM. Production and characterisation of the monoclonal antibody used in this investigation, has been described previously (Banerjee *et al.*, 1994).

Wood wafers which were infected with *0. piceae*, were cut into 20 µm thick sections using a microtome prior to incubation in 1% (w/v) casein in phosphate buffered saline (PBS, pH 7.4) for 20 minutes to block non-specific antibody binding. Wood sections were then incubated for 1 hour in monoclonal antibody diluted 1:500 in PBS containing 0.01% Triton X-100 (Breuil *et al.*, 1992). Sections were repeatedly washed in water and then incubated in *0.5* mg/mL Oregon green 514 Goat anti-mouse IgG (H+L) fluorescent dye (excitation 511 nm, emission 530 nm; Molecular Probes, Inc., Eugene, OR) in the dark for 1.5-2 hours. After four washes in water, sections were individually mounted in glycerol on glass slides with a cover slip and the edges were sealed with nail vanish. Microscopic examination was performed using a Leica TCS NT CLSM with an Argon-Krypton laser excitation source (488/568/647 nm). Images were captured in green channel through a band pass filter BP 530/30 nm and in red channel through a long pass filter LP 590 nm using a 63x oil lens with numerical aperture setting of 1.4. Digital image stacks obtained on

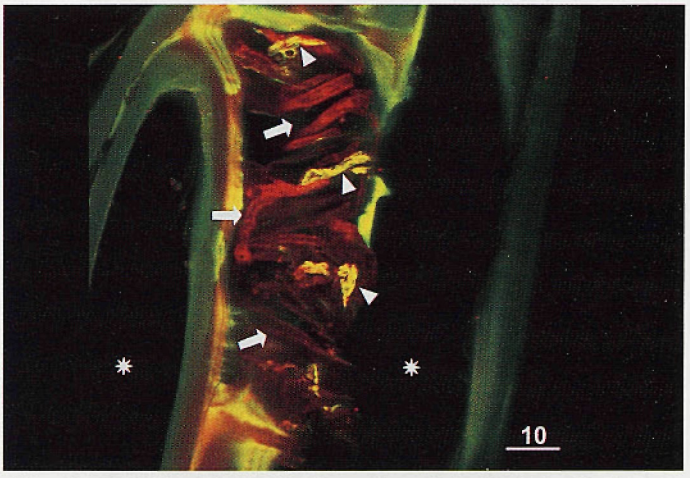

Figure 1: Duel channel projection in maximum intensity mode of an image stack with 2.4 µm thickness. Because the colour of the green and red channel were inverted fungal hyphae appear in red (arrows), and wood cell walls in yellow-green. Arrowheads indicate wood extractives and asterisks wood lumina. Bar = 10 µm.