Appendix A:

Filtering Fundamentals

How digital images differ from photographs

Color photographs display an infinite number of shades, ranging from pure black to pure white with everything in between. These shades are continuously blended together without interruption. Unlike a human, a computer cannot work with continuous tone images . Digitizers and scanners are used to convert the grayscale information in a photograph to numbers. A digital image is created by placing an invisible grid over the photograph and reading the information about the brightness, contrast and color of the image at each grid location. The resulting numbers from each grid location are assigned to a pixel. Information from all the pixels in an image are grouped together to create a pixel matrix. This matrix contains information on:

- Coordinate location—defining the row and column (X, Y) of each pixel in the matrix.
- Brightness information—intensity information at a pixel location.

Image analysis software is designed to enhance and manipulate these pixel matrixes. These enhancements are made by changing the gray values of pixels in the image. The following sections contain a brief background on computer image enhancement terms and techniques.

Humans view grays differently than computers

Before you begin filtering an image, it is important to understand how the human eye sees gray levels. This may help you decide what type of filter to use. It may also explain some of the unexpected results you may get when painting on an image.

Our eyes naturally perform some filtering of an image to help us see details. When looking at a black to white edge, the eye emphasizes the transition to make it appear as sharp and crisp as possible. In effect, the eye creates its own edge enhancement to sharpen the appearance of an image. As a result, it is not always necessary to use edge enhancement filters to sharpen an image. A contrast filter may give you the same results because your eye automatically sharpens the edges.

Our eyes do not view colors and intensities in a linear fashion. Our eyes have different sensitivities to different colors. Human eyes are far more sensitive to darker colors than lighter ones. You may be able to distinguish between two dark colors of a slightly different shade, while two light colors separated by the same difference in shade may appear to be the same color. Thus slightly darkening an image brings out details previously unseen. Scientists can also use the phenomenon to invert bright images to bring out details in the lighter area of the image.

The human eye also has the ability to adjust to the average intensity of a region. As a result, the brightness of an object depends upon the gray levels surrounding the area. For example, if a medium gray object is placed on a black background, it appears to be brighter than the same object on a white background. Because of this, filtering processes will generate different perceptual responses depending on the gray level of the background.

Histogram and Filtering Processes

In technical terms, a histogram displays the frequency distribution of the number of pixels in an image of a specific color. This is represented in a graph with the 256 gray levels displayed (8 bit images containing 256 shades of grays or colors) on the horizontal axis. The number of pixels in the image for each color is displayed on the vertical axis.

A histogram shows information on the number and types of colors in an image. The higher the peak, the more pixels in the image of that specific color. The width of the histogram shows the range of colors in the image. For example, if the histogram has a range of colors between index value 200 and 150, the image appears very washed out with low contrast. But a histogram with colors ranging from an index value of 0 to 255 appears more natural.

The distribution of color in an image also affects its appearance. For an image to appear natural, it is important that the curve be fairly smooth with no sharp or narrow peaks in the histogram.

When colors are not evenly distributed, the image may not have a natural appearance. One such case is for a high contrast image. The histogram of these images has two very large peaks (bimodal histogram): one centered around black and the other around white. These images are often described as being harsh in appearance. Features within the image appear as sharp black to white transitions.

Low contrast images compress the colors into a small range. These types of images are often described as being “washed out.” Since low contrast images often hide subtle details, stretching the histogram significantly improves the appearance of the image.

The location of the peaks in the histogram characterizes the overall tonal quality of the image. For example, if a histogram has one very large peak around an index value of 200, the image appears very bright. If the peak is shifted to an index value of 25, the image now becomes very dark. The sharpness of the peaks gives you a feel for the kind of edge detail in the image. An image with a histogram showing very broad and smooth peaks probably does not have very sharp details. As a result, when you smooth an image, removing edge details, the peaks of the histogram also become much smoother.

Histograms can immediately tell you whether an image is dark, light, low contrast or high contrast. By manipulating an image’s histogram, you can correct many of the problems discussed above.

Effects of Filters

Through filtering operations, images can be changed to improve their appearance or enhanced to bring out hidden image features. All the filtering techniques described below are single pixel point processes. They perform mathematical operations on each pixel within the image. One by one each pixel is chosen, a mathematical or logical operation performed, and the resulting value substituted back into the image. The following list describes which of Color It!'s filters manipulate the histogram:

Invert: Inverts all the gray levels in the image (blacks to whites, whites to blacks, etc.). In technical terms, the filter inverts the histogram along the horizontal axis.

Equalize: For low contrast images, the Equalize filter may give you the best results. The filter is especially suited to bring out details in a low contrast image. In technical terms the filter Equalizes the distribution of gray levels (i.e. broadens the peaks of a histogram).

Contrast: The contrast filter is the most commonly used filter for improving the appearance of an image. This filter gives low contrast, washed-out images new life. The filter linearly stretches or compresses a selected range of index values. A histogram stretch will expand the colors to cover the entire range of colors.

Gamma: The gamma filter is an all-purpose filter which is often used for improving the appearance of image. It should be your second choice of filters when trying to bring new life to a washed out image. A Gamma filter skews the color distribution. The direction the histogram is skewed depends upon the value chosen. Higher values lighten the image. Lower values darken the image.

Highlight/Shadow: The highlight and shadow options help you compensate for your printer's inability to print all colors. All printers have problems printing both the darkest and lightest shades in an image. Printers tend to force these shades either to pure black or pure white. The Highlight and Shadow option compresses the histogram and removes the colors the printer has problems with. Since the new range of colors is closer to the printer's range, the image looks much better when printed. While the screen appearance of the image seems to degrade, the final printout is far more natural looking. In technical terms, the filter compresses the histogram into a smaller range of colors (that is, creates a narrow peak in the histogram).

Note: Though the screen image may appear to degrade after a filter operation, the image will look fine when printed.

Neighborhood Processing

While single-point pixel processing can be used to control the tone of an image, neighborhood processing techniques are most often used to change the sharpness of an image. Filters can be developed to

- Sharpen edges—Enhance the sharp brightness transitions in an image (Laplacian, High Pass filters).
- Smooth images—Only show slow brightness transitions (Low Pass filters).
-

Create line art—Only show the sharp transitions (Edge Detection filters).
How does Neighborhood Processing work?
Neighborhood processors differ from single point pixel processors because they use the information about adjoining pixels. Neighborhood filters perform mathematical calculations on the group of pixels surrounding the current pixel. By looking at this information, filters make decisions about how to change the appearance of the image. The mechanics of a convolution mask are rather simple. The following procedure outlines the calculation of a 3x3 convolution mask on an image. Each pixel in the image is evaluated with its eight nearest neighbors. These nine pixels use the convolution mask to produce a new value which is placed in the output image. In effect, the mask is placed over the nine input values, where each of the pixels is multiplied by a corresponding weighted coefficient in the mask, the values summed and the result placed in the output image.
Mathematically, a weighted average of the neighborhood is calculated for each pixel in the image. The result of these calculations is then placed back into a new filtered image. The average is formed using a group of pixels, called a kernel, which surrounds and includes the center pixel. For a 3x3 kernel, the computer multiplies each of the nine pixels by a weighting factor (convolution coefficient) and then adds the numbers together to generate the weighted average. By changing the weighting factors, certain pixels have a greater or lesser effect on the overall average.
NOTE: As the size of the kernel is increased, computational time increases dramatically.
The following list describes which Color It! filters do neighborhood processing:

Low Pass: This filter smooths, mellows or blurs the appearance of the image. A Low Pass filter removes the sharp white to black transitions in an image and replaces the transition with a blurred edge.

High Pass: This filter highlights the edges or transitions in an image (i.e. sharpens an image but also introduces noise).

Edge Detection and Enhancement Filters
Edge detection filters reduce an image to display only its edge information. This information is used to generate special types of line art or special effects for desktop publishers. The following figure shows how an edge transition is displayed after doing a

Gradient or Laplacian filter.
The following list briefly describes the Edge Detection filters available in Color It!.

Laplacian Edge Enhancement
The Laplacian filter enhances all edges regardless of their direction. The

magnitude and direction of the color transition directly affect the color placed in the output image. Thus, edges appear as alternating dark gray or black and white pixels on a gray background. This filter is useful for determining the noise in the image and is an excellent way to test a scanner. Adding these filtered images back to the original seldom give visually appealing results.

Gradient Filters, Stochastic Filters

These are like the Shadow/Dark Shadow filters discussed below. This class of filters looks for gray level edges in more than one compass direction at a time. After a Gradient or Stochastic filter is performed, the image is turned into a pseudo line art version of the original image. Adding these filtered images back to the original generates some visually interesting results.

Shadow/Dark Shadow Filters

These filters are designed to highlight one of eight compass directions: N, NE, E, SE, S, SW, W or NW. These filters create shadows next to objects in the image (i.e. an East mask will show shadows on the east side of objects). Adding these filtered images back to the original produces a pseudo 3D effect.

Shift and Difference

These filters look for horizontal or vertical edges. In essence, the filter takes an image, shifts it one pixel over and subtracts it from the original. The effect is that of embossing the image. Adding these filtered images back to the original sharpens and enhances the general detail of the image.

Line Segment Enhancement

A Line Segment filter enhances line segments within the image.