
Using Curves and Histograms

Written by Jonathan Sachs
Copyright © 1996-2003 Digital Light & Color

Introduction

Although many of the operations, tools, and terms used in digital image manipulation have direct equivalents in conventional darkroom techniques, histograms and curves remain an unfamiliar and daunting concept to many photographers. This document will try to demystify curves and histograms by explaining what they are and how to interpret them. Some of the information may be a little technical, but mastering it is well worth the effort. Once you understand how to use them you will be able to:

- Lighten images without losing highlight detail.
- Darken images without losing shadow detail.
- Adjust midtones without altering shadows or highlights
- Control brightness and contrast precisely for each part of the tonal range.
- Evaluate the quality of scanners.
- Create a wide range of special effects.


What is a Histogram?

The term histogram has several meanings. Originally a graphical tool developed by statisticians to visualize frequency distributions, it has come to have a very specific meaning when used in the context of digital images.

To keep things simple, we'll start with black and white images; later we'll see how the same concepts work in color. In black and white images, each pixel has a specific brightness level which is stored in the computer as a number between 0 and 255. Zero corresponds to black and 255 to white. Since the eye, even under ideal conditions, cannot barely distinguish 200 different gray levels, the 256 available gray levels in a digital image, if properly used, are more than adequate to represent even the most subtle tonal variations in a black and white image.

How histograms are computed

Imagine a row of 256 buckets, the first one labeled 0, the next one labeled 1, and so on up to 255. The histogram of an image such as the one at the left is computed by examining each of its pixels one at a time and tossing it into the bucket that corresponds to its brightness level. When we're done, we count how many pixels are in each bucket. Picture Window displays histograms as a graph that looks like this:



The gray scale along the bottom of the histogram indicates the brightness level of each bucket, starting with zero at the far left (black) and progressing up to 255 on the far right (white). For each brightness level, there is a vertical white line that represents how many pixels in the image have the corresponding brightness level. To make the display fit the available space, the pixel counts are scaled so that the tallest line runs all the way from the bottom to the top of the graph.

How to Interpret Histograms

The basic idea is simple: where you see a tall line in a histogram display, this means there are a lot of pixels of the corresponding brightness level in the image. Where you see a short line or no line there are very few pixels of that brightness level.

Overall brightness

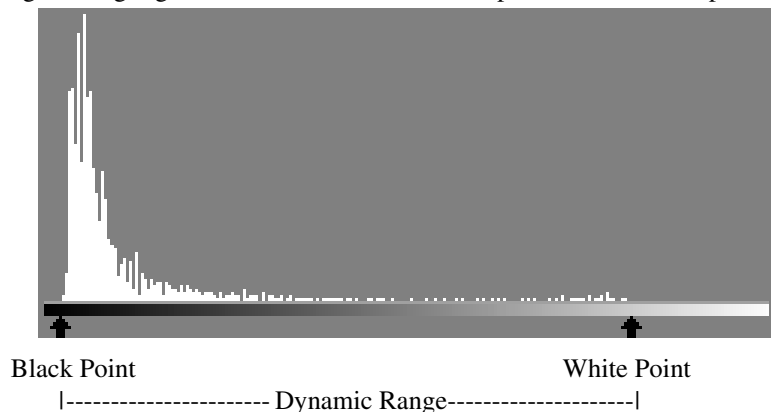
The first thing we can tell about the example image above based on its histogram is that most of it is dark since most of its pixels have relatively low brightness levels as indicated by the lines at the left of the histogram being much taller than those at the right.

Each image has its own unique distribution of light and dark tones—there is no right or wrong histogram. If an image is made up of all one shade of gray, then the histogram of that image consists of a single line corresponding to that gray level. If it is part one gray and part another, then it would have two lines, one for each gray level, and their relative heights would reflect how much area of the image is covered by each gray. If the image is a smooth gradient going all the way from black to white, each of the lines in the histogram would be the same height since there would be equal numbers of pixels of each brightness level.

Bear in mind that the histogram says nothing about where the pixels are located in the image, just how many of them have a given brightness. You get exactly the same histogram if you rearrange all the pixels as long as you don't alter their brightness.

Dynamic Range -- Finding the darkest shadows and lightest highlights

The next thing we can tell about an image from its histogram is its darkest shadow and lightest highlight, sometimes called the black point and the white point.



Remembering that if the histogram is zero there are no pixels of the corresponding brightness, we can see that where the histogram drops to zero on the left (marked Black Point), there are no pixels in the image darker than this. Similarly, where the histogram falls to zero on the right (marked White Point), there are no brighter pixels

in the image. These two points define the dynamic range of the image -- the range between its darkest and lightest parts. We can see from looking at the histogram of this image that its shadows do not go all the way to pure black and that its highlights fall quite a bit short of pure white. The greater the dynamic range of an image, the more contrast it exhibits. If all the lines in the histogram fall in a narrow range, this means that the image has only a small variation of gray levels and will consequently have a very low contrast. Whether this is good or bad depends on the effect you want to achieve.

Posterization -- How smooth is the histogram

Posterization is a special effect where an image is rendered using only a limited set of colors or gray levels. The smoothness of a histogram can tell you how well the image represents subtle tonal variations. The more the histogram consists of isolated spikes separated by empty spaces, the more the image is posterized. While intentional posterization can be a very striking effect, in normal images, excessive posterization, or as it is sometimes called *quantization*, of an image can indicate a loss of subtle tonal variation which in turn can mean lower image quality. Quantization can be caused by starting with a posterized original, poor scanning of the image, or by excessive image manipulation.

Evaluating scanner performance


If you start with a photograph that has a full range of gray levels and your scanner produces a digital image that does not have a fairly smooth histogram, this usually means your scanner is not doing a good job of capturing the brightness range of the image. Less expensive scanners often measure only 256 gray levels per pixel and, once corrections are applied to the image data to adjust its brightness and contrast curves, many of these gray levels can be lost. Consequently, the resulting image may use substantially less than 256 different brightness levels, and you can tell this by examining the histogram of the image and looking for spikes and gaps. More expensive scanners capture 1024 or 4096 different brightness levels per pixel and this oversampling leaves plenty of room to make corrections to the image data and still end up with 256 useful brightness levels. For this reason, the best time to manipulate the brightness curve of an image is during the scanning process. Or, if you are using Picture Window Pro, you may be able to capture and work with a 16-bit black and white or 48-bit color image,

Manipulating images—Less is More

Even if you start with a scanned image that has a nice smooth histogram that covers the entire tonal range, each time you perform operations on the image that change its brightness curve such as trying to bring out shadow, midtone, or highlight detail, you will sacrifice some brightness levels. The more radical the change, the more likely you are to introduce posterization. The only way to avoid this problem is to try to make as many of these changes as possible at the time you scan the image (assuming of course that you have your own scanner and that your scanner captures more than 8 bits of data per color component), or, if you have Picture Window Pro, to capture and work with more than 8 bits per channel.


Two Images and Their Histograms

To help you get a feel for the relationship between an image and its histogram, here are some more examples.




The histogram of this image has two peaks, a broad one in the shadow area and a narrower one in the light gray area. These correspond to the foreground and building which are large and dark and the background sky which is a fairly uniform light gray. The area under a peak is proportional to the percentage of the pixels in the input image that have a particular brightness range; the width of the peak indicates the brightness range it covers. As you can see, while this image has plenty of dark tones, it does not go all the way to white which makes it somewhat dingy. Also, there are essentially no midtones, which is typical of images that silhouette a dark subject

against a light background or vice versa.



This image has a fairly smooth histogram. Note that it does not go all the way down to black or all the way to white, but that is normal for this type of image because the original scene lacked deep shadows or bright highlights.

How to Use the Histogram Display in the Brightness Curve Transformation



Picture Window provides a very powerful transformation accessible from the main menu as *Transformation/Gray/Brightness Curve*. Brightness curves can be used to achieve a wide variety of effects. The unique feature of Picture Window's Brightness Curve transformation is that it lets you view and modify both the curve and the histogram of the image interchangeably.

What is the Brightness Curve transformation?


The Brightness Curve transformation is based on a very simple principle. Each pixel in the result image is computed by applying a curve to the corresponding value in the input image. The curve is just a table of 256 values, each corresponding to a brightness level from 0 to 255. For each input pixel, its brightness is changed to the value found in the corresponding table entry and the result is stored in the output image.

For example, suppose the table simply contains the sequence of numbers from 0 to 255. Then the output image will be an exact copy of the input image because brightness level 0 is set to 0 and so on up to 255. If we draw a graph of this table, we get a curve like the one to the left—a straight line running diagonally from the lower left corner (0,0) to the upper right corner (255,255).

To see what the curve will do to a pixel of a given brightness, locate its gray level along the horizontal axis, run a vertical up from that point to the curve and then run a horizontal from the intersection point over to the vertical axis and read off the corresponding brightness value in the output image.


Switching between Curve and Histogram

The Brightness Curve transformation lets you work with both the brightness curve and with the brightness histogram at the same time. You can either view the curve in the foreground with the input image histogram displayed in the background or you can view both the input and output image histograms with the curve displayed in the background. To switch from one view to the other, use the two button tool bar just below the Option buttons to the right of the curve display.



Curve Display

Histogram Display

When you click the  (histogram view) button, the display switches over to show you two histograms, one above the other. The top histogram is for the input image and remains unchanged for the duration of the transformation. The bottom histogram is displayed upside down and shows you what the histogram of the result image will be after applying the curve to it. This histogram changes as you experiment with different curves. For each control point (see below) on the curve, Picture Window displays a pair of arrow heads connected by a line. The top arrow head defines a brightness level in the input image and the bottom arrow head the brightness level it will be changed to in the output image.

Changing the curve or histogram by adding, removing, or moving control points

There are several ways to alter a curve. The simplest is to click and drag one of the little circles that indicate *control points* on the curve. Initially, there are just two of these, one at either end. To create additional control points shift-click on the curve where you want the new control point to appear. (To shift-click, hold down either Shift key with one hand while clicking once with the left mouse button.) To remove a control point, position the cursor over it and control-click (click on it once with the Ctrl key held down). Neither the first nor the last control point can be removed.

The curve always passes through all of the control points exactly, no matter how many of them you create. The more points you add, the more control you have over the shape of the curve. In practice, however, only a small number of control points are required for most operations. There are many different kinds of curves that can be passed through the control points -- Picture Window lets you select from three different types of curve (smooth, broken line, or stair step) using a three-button tool bar displayed just to the right of the curve or histograms. This feature is described in more detail below.

Control points can be moved freely except you cannot move one control point past another on the horizontal axis. To move a control point, simply position the cursor over it until the cursor changes to a crosshair and then click and drag it to the desired location. In the curve display, for each control point there are a pair of arrow heads displayed one along the bottom edge of the graph and one along the left edge. These arrow heads track the motion of the corresponding control point. You can move a control point in a single direction, either horizontally or vertically, by clicking and dragging the corresponding arrow head.

Adjusting control points from the histogram display is even simpler. All you need to do is click and drag one of the upper or lower arrow heads. To insert a new control point, shift-click anywhere along the base of the upper histogram; to remove a control point, position the cursor over the upper arrow head and control-click.

Controlling dynamic range using histograms




You can either increase or decrease the dynamic range of an image very easily using the double histogram display. Start by positioning the upper arrow heads to the black and white points of the image. These correspond to the darkest and light-

est points in the input image. Then move the bottom arrow heads to where you want the black and white points to be in the final image. In changing the histogram control points, you are simultaneously changing the control points of the curve. For reference, the curve is displaying in the background along with the histograms.

Moving the bottom arrowhead for the black point lightens or darkens the shadow areas. Moving the bottom arrowhead for the white point lightens or darkens the shadow areas. The farther apart you move the two arrowheads, the greater the spread of tonal values in the resulting image and hence the greater the contrast. The closer you move them to each other, the more narrow the range and the lesser the contrast.


Making an image use the full dynamic range

To make the result image use the entire available tonal range from black to white, without losing highlight or shadow detail, position the top arrow heads at the black and white points and the bottom arrow heads at the far left and far right of the brightness scale. Notice that the slope of the corresponding curve gets steeper the more the contrast is increased.



Reducing the dynamic range




To reduce the dynamic range, simply move the bottom arrow heads closer together. Moving the lower left arrow head to the right lightens the shadows; moving the lower right arrow head to the left darkens the highlights. The closer together you move the bottom arrow heads, the lower the contrast. Notice that the slope of the corresponding curve gets smaller the more the contrast is reduced.



Creating extreme contrast





If you move the upper left arrow head right past the black point, you will continue to increase the contrast of the image but you will start to lose shadow detail; if you move the upper right arrow head left past the white point, you will continue to increase the contrast of the image but you will start to lose highlight detail. You can tell this is happening because the bottom histogram will start to show a sharp spike at black or white. If you see an image with this kind of histogram it is a sign that it was either scanned improperly or with a scanner that could not handle the dynamic

range of the image or that somewhere in the manipulation of the image it was lightened or darkened too much. Of course, if you want to create a dramatic high contrast image, this effect may be exactly what you want.



Creating a negative image

If you reverse the positions of the bottom arrow heads, then you get a negative image. The bottom arrow heads correspond to the brightness levels you want in the final image. If you move the lower left arrow head all the way to the right, you are telling Picture Window to make the dark parts of the image light; if you move the lower right arrow head all the way to the left, you are making the light areas dark. Changing dark to light and light to dark gives you a negative. Notice that the corresponding curve slopes downward instead of upward.



Using Curves to control brightness and contrast

While histograms are a powerful tool for adjusting the dynamic range of an image, curves are sometimes better for adjusting its brightness and contrast. You can switch back and forth between the two methods at any time just by clicking one of the two buttons on the tool bar just below the Options button to the right of the curve.

Just like histograms, once you understand them curves are a very powerful tool for manipulating images. Here are a few very important things to keep in mind about curves:

1. Where the curve lies *above* the diagonal line connecting the lower left corner to the upper right corner, it will *lighten* the image; where it lies *below* the main diagonal it will *darken* the image. The closer the curve lies to the main diagonal, the more subtle its effect on the image; the farther away, the more effect it will have.
2. Where the curve rises steeply, it will increase the contrast of the image; where it rises slowly, it will decrease the contrast. Increasing the local contrast in one part of the tonal range always decreases it in another part.
3. Where the curve is flat, the image will be posterized.
4. Where the curve slopes downward, the image will be a negative or solarized.
5. Kinks in the curve usually produce some kind of artifact in the image. Unless you are trying to create special effects deliberately, try to keep your curves smooth.




To fully control brightness and contrast, you will usually need to add control points to the curve. Additional control points give you more control over the shape of the curve since the curve always passes through each control point. To add a new control point, first position the cursor over the curve display where you want the new point to be and then shift-click (click the left mouse button while holding down the shift key with your other hand). To remove control points (except for the first and last point which cannot be removed), position the cursor over the control point and control click (click the left mouse button while holding down the Ctrl key).

Lightening or darkening an image

Let's continue using the same image and assume that we have just used the histogram to increase its dynamic range to go all the way from black to white, but we still want to lighten it up to help bring out some of the shadow detail.

The first way you might think of to brighten an image is to add a fixed amount to each of its brightness values. This corresponds to shifting the entire curve upward or shifting the entire histogram to the right. The problem with this approach is that it can cause the darkest shadow areas to lighten up and it can cause loss of highlight detail. This shows up in the histogram as a blank space at the left hand end of the histogram and spike at the right hand end indicating that a number of different brightness levels at the lighter end of the scale have all been changed to the same maximum level. A similar problem occurs when attempting to darken an image by shifting its histogram to the left. To avoid this difficulty, we need to lighten or darken the intermediate brightness levels while leaving the black and the white points unchanged. To do this we need curve that is not a straight line. Here's the procedure:




First add two intermediate control points in the middle of the curve and then position them as illustrated below on the right. Make sure you leave the first and last control points fixed as these determine the dynamic range. Notice that since the curve lies above the main diagonal, it is lightening the image. Also notice that because it slopes steeply on the left that it is increasing the contrast in the shadows; since the slope is flattened in the highlight area we are reducing the highlight contrast. The result is a brighter image with a lot more shadow detail. To darken an image, arc the curve downward (below the main diagonal) instead of upward.



Increasing or decreasing midtone contrast

When people talk about increasing or decreasing the contrast of an image, they are usually talking about increasing or decreasing its midtone contrast. Bear in mind however that you cannot increase contrast in the midtones without decreasing either the shadow or highlight contrast or both. Increasing the midtone contrast involves





lightening the highlights and darkening the shadows. Reducing the midtone contrast means lightening the shadows and darkening the highlights. The curves look like this:



Special effects -- solarization






Solarization is a special effect where part of the tonal range is inverted. This creates an image which is part positive and part negative. To solarize, you need to make part of the curve slope upward and part slope downward. Since there are many kinds of curves that pass through all the control points, Picture Window gives you three options: stair step, broken line, and smooth curve. You select these options from the 3-button vertical tool bar just below the one that selects curve/histogram. For solarizing, the broken line curve (center button) usually gives the best results. The following examples illustrate different types of solarization and the curves that produce them.




Special effects -- posterization

Posterization is the result of rendering an image using a limited number of gray levels. This creates an image which is made up of areas of solid color, something like a silk screened poster. To posterize, you need to make a curve that has a series of stair steps. To create this effect, click on the top icon on the 3-button vertical tool bar just below the one that selects curve/histogram. To create more gray levels, add more control points; the more levels you use, the more subtle the effect. If the stair step curve runs primarily above the main diagonal, the posterized image will be brighter than the original; if it runs below the diagonal, the result will be darker. You can combine posterization with solarization for additional effects. The following examples illustrate two different posterization curves and the effects they produce. The first example uses only three gray levels (black, white and a fairly light gray); the second uses five (black, white, and three intermediate grays).




Using other features of the Brightness Curve Transformation

Using the brightness probe

You may have noticed the small eyedropper icon  in the lower right corner of the Brightness Curve dialog box. Clicking on this button toggles the state of the brightness probe. When the button is depressed, the brightness probe is enabled; when the button is raised, the probe is disabled. Whenever the brightness probe is enabled and the Brightness Curve dialog box is the active window, you can click


and drag the mouse cursor over the input image window and a red marker line will appear in the curve or histogram display to tell you the brightness value of that part of the image. The marker line runs from the top to the bottom of the curve display or from the top to the bottom of the input histogram display as illustrated to the left. You may need to rearrange and/or resize some of your windows so you can see both the input image and the Brightness Curve dialog box at the same time. The point where the vertical red line intersects the gray scale at the bottom of the curve or histogram is an indication of the brightness of the input image at the probe's location.



The brightness probe is a very useful tool. For example, suppose there is one specific part of an image that you want to lighten or darken. You can start by activating the probe and clicking and dragging the cursor over the region in the input image you want to alter. By noting where the marker shows up on the curve or histogram display, this will show you what part of the tonal range needs to be changed. To lighten or darken that part of the tone scale, you can use either the curve or the histogram.

Either way, the first step is to create a new control point at the brightness level of the region you want to change. There are two ways to create a new control point:

1. Shift-clicking over a part of the input image representative of the brightness level you want to modify automatically inserts a new control point at the corresponding point in the curve or histogram.
2. Shift-clicking over the curve or upper (input) histogram where you want the new control point to appear inserts a new control point at the cursor location. You can base the location on the results of previous probing of the image.



The next step is to apply brightness corrections using the curve or histogram.

If you are using the curve, then drag the new control point up to lighten or down to darken -- the farther you move it, the brighter or darker that part of the tonal range will become. You can use the Preview button to check your progress.

If you are using the histogram, then drag the control point's bottom arrowhead to the right to brighten or to the left to darken.


If you start changing other parts of the image that you don't want to modify, you can probe them and create additional control points to control their brightness independently. Or, you can create a mask to restrict the effects of the transformation to the selected part of the input image.

If you end up creating a curve that is not fairly smooth, it will probably result in some kind of visual artifact in the result image since the local contrast will either be very low or very high in some parts of the brightness scale. If this happens, switch to the curve view and adjust the control points to make the curve smoother. If you can't get the effect you want on one region without adversely affecting other parts of the image, you will need to create a mask that isolates the region you want to adjust. This lets you restrict the transformation to a specific part of the input image.


Histogram Expansion -- what it is and when to use it.

When Picture Window displays histograms, it automatically scales them so the largest value fills the available space between the top and the bottom of the graph. If a histogram happens to contain one very tall spike, this can create a situation where the rest of the information in the histogram is scaled down so much that it disappears or becomes too small to interpret. When this happens, you can use Picture Window's Histogram Expansion options (available from the Options menu) to display histograms scaled in such a way as to exaggerate small values and compress large values. Histogram expansion is available in two strengths: Medium or High. This expanded scaling lets you see both the small histogram values and the large ones at the same time even if the large ones are huge and the small ones are tiny. However, when using histogram expansion, you must bear in mind that large values are really much larger than they appear and small values are smaller.

Working with color images



Everything we have discussed so far deals with black and white images, but what about color? For color images there are many different ways to compute histograms. We could compute separate histograms for the red, green, and blue components of the image or, what is often more useful, we can use the HSV or HSL color space and just histogram the brightness component, ignoring the hue and saturation. This lets us work with the brightness of the image without worrying about changing its color. Interpreting a color brightness histogram requires some understanding of the underlying color model. In the HSL model, a brightness level of zero corresponds to black and a maximum brightness value corresponds to white. In the HSV model, zero brightness still corresponds to black, but the a maximum brightness value corresponds to any color whose red, green, or blue component is maxed out at 255. This includes all the colors in the color hexagon displayed in Picture Window's



color picker dialog box. When you select the RGB color model, Picture Window displays a histogram computed as the largest of the three component histograms for each brightness levels.

The one thing to remember about the difference between HSV and HSL color spaces is that when you lighten an image using HSV, the colors get brighter while remaining fully saturated. As you lighten using HSL, colors get progressively more washed out as they get closer and closer to pure white. You must choose the color space according to the effect you want to achieve.



Lightened using HSV




Lightened using HSL

Normally the RGB color space is a poor choice for lightening or darkening an image. When you use this color space, the same curve is applied to each of the three components of the image (red, green, and blue). This can result in unwanted hue or saturation changes in different parts of the image. Working in HSV or HSL avoids this problem. On the other hand, when using the Brightness Curve transformation to produce special effects like solarization and especially posterization of color images, the best results are often obtained using the RGB color space instead of HSV or HSL. For example, creating a negative image using RGB yields an image where each color is replaced by its complement and light colors become dark and vice versa. The same effect applied in the HSV or HSL color space leaves the hue

and saturation of each pixel the same and just inverts the brightness -- the effect is very different:



Negative in RGB color space



Negative in HSV color space