

Local contrast adjustments:

There are three common methods, two of which I use with some frequency.

High radius, low amount (aka “hiraloom”); Curves with masking, and Overlay/Soft Light with adjusted Opacity.

A fourth technique presented, High Pass/Overlay, is really a sharpening method, but most (if not all) sharpening methods are really just local contrast enhancers. I talk about it last.

The one I don’t often use, “Hiraloom”, is just applying the Unsharp Mask filter (USM sharpening) using a high radius (e.g. 150 pixels) and a low amount (e.g. 35%). I’ll let you play with it and see what you think.

The one I use the most is Curves with masking.

If you’re already familiar with using Curves, you can skip down to just past the bold paragraph below starting: “This is a key concept.”

Curves can be intimidating at first, but a few little tips make them an intuitive and extremely powerful tool. In what follows, I’m assuming that you’re working with a black and white image (i.e. you’ve already converted it). There’s a bit more to it with color images, but not a whole lot.

Let’s look at the curves interface. The interface contains a chart and a variety of controls. The chart has a vertical axis on the left and a horizontal axis at the bottom. By default, the line in the chart is straight and at a 45 degree angle from bottom left to top right. I refer to this line as the “mapping curve” or simply the “curve”, even though it’s a straight line right now. Trust me, it’ll sometimes be plenty curvy.

The first thing to note is that the horizontal axis at the bottom of the chart window defines the “input” tone. The left axis of the chart defines the “output” tone. When the curve crosses the chart at a 45 degree angle (the default), each input tone on the input axis is being “mapped” to the same tone on the output axis (i.e. Curves isn’t doing anything).

Also note that along each axis there is a grayscale. This tells you what tone corresponds to that point on the axis.

You can determine where a tone is mapped by following a line vertically from the input axis to the curve, then left horizontally from that point to the output axis.

When you raise or lower a point on the curve, you change the output tone to which the input tone is mapped. In other words, if a tone in the input image is a 30% gray, and (don’t do it right now) you raise the curve at a point vertically above 30% gray on the input axis, the 30% gray in the input image is now lightened to a tone corresponding to the vertical position of the point along the output (vertical) axis.

To give some intuition, let's do a couple of things. First, take the anchor (the little square box) on the curve at the origin, and drag it to the top of the chart. Notice that the entire image is now white. What you've done is told Photoshop to map EVERY input tone to the white output tone.

Move the anchor point back to the origin. Now, grab the anchor point at the top right and move it all the way to the bottom of the chart. Your image will be black. You've just told Photoshop to map EVERY input tone to black.

Leave the right anchor point where it is, and drag the left (origin) anchor point up to the top of the chart. This inverts the image. Do you see why? You've told Photoshop to map each input tone to its tonal opposite. Note that mid gray is unaffected, because its input tone still maps to the same spot on the output axis.

Next, move the anchor points back to their original positions. The image should be back to normal.

Now, click on the center of the curve at the center of the chart. This creates a new anchor point. You can drag this anchor point up and down and change both the lightness/darkness and the contrast in parts of the image. Note that the curve adapts to smooth the transition, preventing abrupt changes to the tones in the image.

When you drag this new anchor point up, notice three things: the mid tones get lighter, the contrast in the shadows increases and the contrast in the highlights decreases. By lightening the mid tones, you are "stealing" tones from the highlight areas and "giving" tones to the shadow areas. You can see this visually by looking at the slope of the line in the highlight and shadow areas. The slope in the shadows, normally to the left of the midpoint, increases (becomes more vertical), while in the highlights, normally to the right of the midpoint, it decreases (becomes more horizontal).

This is a key concept: Where the slope of the curve gets steeper (more vertical), output contrast in the corresponding range of input tones increases. Where the slope of the curve gets gentler (more horizontal), the output contrast in the corresponding range of input tones decreases.

Those are the basics of curves. If you play with it for a while, it starts to become intuitive and you'll have a powerful new tool in your arsenal.

Now, about the mask...

All adjustment layers come equipped with a mask by default. The mask allows you to control the amount of application of the adjustment at any place in the image. The amount by which the adjustment is applied depends on the percentage of white in the mask's tone (grayscale), with full white being 100% and black, 0%. The adjustment layer's effect is multiplied by this percentage to determine the strength. So, if you paint the mask all black, the adjustment layer has no effect. If it's all white, the mask has no effect and the adjustment is 100% active on every part of the image. If the mask is 50% gray, the adjustment is applied half-way.

You can paint on the mask with a paint brush. Just make sure that the mask is active before you start painting. I've accidentally painted over areas in images many times. There's "undo" of course, but it interrupts your train of thought and is generally annoying. It's usually a good idea to paint with a soft edged brush, unless you have a good reason to do otherwise. This smoothes the transition at the mask edge and makes it appear more "natural".

Method 3: Overlay or Soft Light image with adjusted Opacity.

This one is simplicity itself. Simply take a copy of the image, move it to the top of the layers stack and set its blend mode to "Overlay" or "Soft Light". Soft Light can be thought of here as a more subtle version of Overlay. Your image probably looks terrible: chalk and charcoal.

Begin backing off the opacity of the new layer until the image starts to look good. I've found that often, this doesn't occur until I reach 15-20% opacity. That's okay. If you reach 0% opacity, you can discard the layer altogether. Sometimes, as in the example image I presented, the percentage is high (51%).

Last but not least:

The sharpening method I presented as a texture enhancer is sometimes referred to as the "High Pass sharpener". To do this, begin by making a copy of the image. Place this on top of the layer stack. Apply a High Pass filter (the radius is commonly set to about 2-3 pixels, but it's up to you). For a moment, you'll just see only mid gray, but if you watch for a second, your eyes will adjust and you'll see a very fine trace of the edges in your image. Next, set this layer's blend mode to Overlay (Oops! I omitted this step in my presentation. Thanks to Sam Wilson for pointing this out).

That's it! You can do all the usual stuff, of course, like mask it, adjust opacity, etc. One book I read suggests taking a mid gray brush and painting over the areas where you don't want the enhancement, but I prefer a mask, that way I can adjust it anytime I want.

Well, that's it for local contrast adjustment. There are other methods, I'm sure. But these are pretty standard.

A piece of advice: All local contrast adjustments should be checked again after you've been away for a while. Our visual system is extremely adaptive and will mislead you. Before printing, close your file and go do something else for at least a half hour, preferably overnight. Then reopen the file and see what you think. It's amazing how different it can look.