(H) Now, just for laughs, I'm going to use my "Gamut Warning" trick. View>Proof Setup>Custom and set the proofing to AdobeRGB 1998. Then, as we did on page XXX, go to View>Gamut Warning, and take a look.

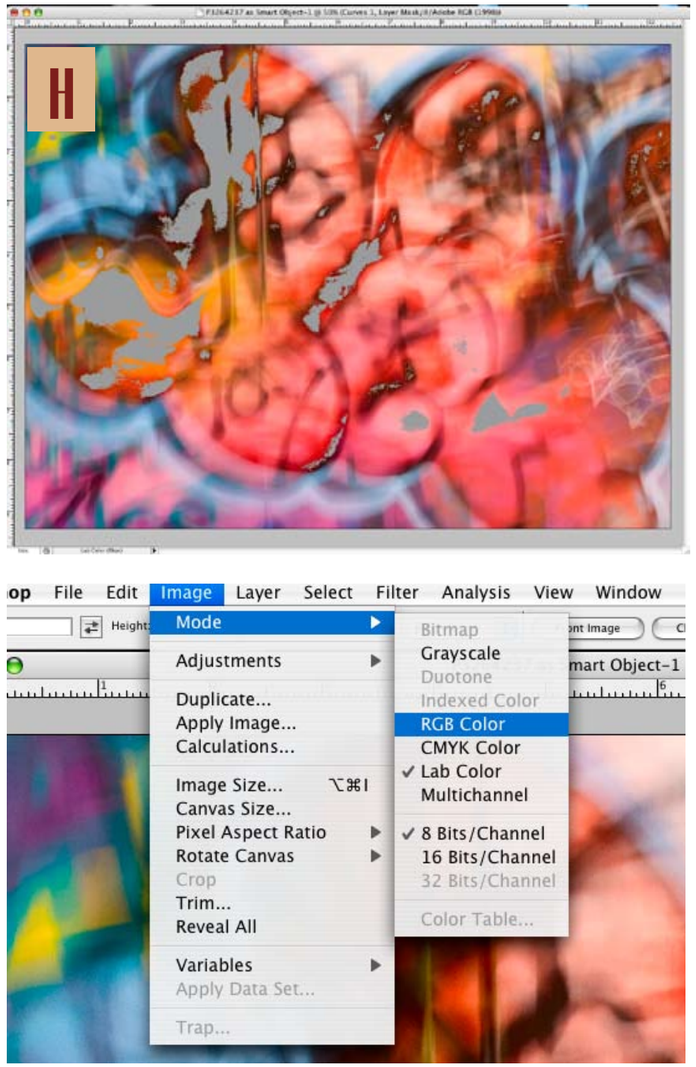

Without even trying too hard, I have pushed some colors outside of AdobeRGB; this does not even mention where they've gone relative to my printer and paper, which is probably a much further jump out of gamut. Dangerous, indeed.

I liken LAB editing to a nail gun. Bear with me; this is a really good analogy.

A nail gun allows you to work really fast. Roofers not only love them, but if you're trying to stay in business as a roofer, you have to have one. LAB color, because it allows you to darken and lighten without shifting color, and make color adjustments using two powerful channels, allows you to work fast, too.

A nail gun is, however, a really dangerous tool. You need to know where you're shooting nails, and what they are going into. Shoot a nail into rotten wood and it will come out the other side like a bullet. LAB color editing can shoot way past your mark, even into colors you can't see with your eyes, much less the colors you can see on your display.

Back to the example at hand—to deal with what the gamut warning has just told me, I simply go back and convert the file to AdobeRGB 1998 again. Here is that dialog, Image>Mode>RGB Color. Remember, it will convert to the RGB space you've set up as default in your Color Settings (Edit>Color Settings).

There are a couple of things I want to make you aware of. First, you should be working in 16-bit files (which, if you're observant, you may have noticed from this example that I'm not). Also, remember that you are definitely changing the file by going back and forth to the LAB space, so you need to be as precise as possible about using it, keep the conversion to only one time, and avoid too much change.

Another interesting advantage of working this way is in the masking and selections you can make in LAB. The very briefest of examples I can offer here is in sharpening. If I select only the L channel, and apply Unsharp Mask, I'm working on a very important and influential part of the image information for determining edges—the dark-to-light range. By adjusting the sharpness of just the image file's luminance, I can avoid any unpleasant side effects of Unsharp Masking, like weird color haloing. (This is a bottomless rabbit hole of very interesting and very specialized techniques, and one that you may understand a little better by understanding the context of the colors and their little journey through the system.)

LAB Editing in Raw Developer

Just for laughs, let's look back at Raw Developer, by Iridient Digital. Here we have access to the LAB controls. Since we are working with a RAW file in the RAW processor, the conversion to and from LAB mode is not as destructive and therefore not as much of an issue.

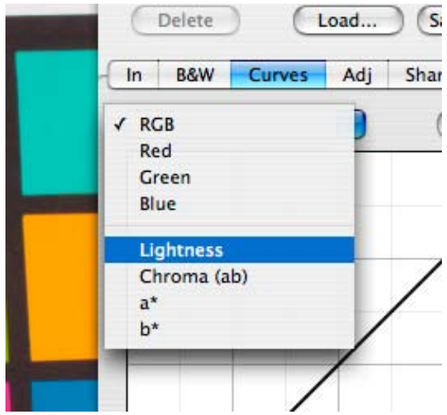

We start with our RAW file, connect to the processor with the camera profile, push it around with LAB "handles," and process it out to our working color space. Because we are still processing the file in RAW, the conversion issue goes away. We can still make swings in LAB controls that are out of our range, but they get processed right to AdobeRGB (or whatever working color space you have selected). LAB editing in Raw Developer is just as powerful and a little less erosive.

19 GRAYSCALE CONVERSION

Grayscale Conversion in the Color Journey

Black and White Are Colors, Too!

At one point in my career, my studio was in a building next to an electrician's office. The guys were great; they had no idea what I was all about, except that they figured I was a photographer, so I must have models floating in and out of the studio all the time—how could that be a bad thing? After they realized that most of what floated in and out of my studio were weird looking electronic widget products and overweight CEOs, most of them lost interest except for a couple of the guys—one of whom was James, who had some curiosity about taking pictures.

James came in one day and started chatting to me about a show he had seen on TV. He said it was about this old guy—"Ansel" or something—who took pictures of mountains in black and white. (Keep in mind that A. Adams is my main man.)

"Mountains, in black and white!" James said. "Jeesh, what a scam!"

I was curious. "Scam?" I said.

"Yeah—seriously—mountains in black and white? How hard could THAT be?"

Indeed, James, this is a good question.

Black-and-White Photography

Let's start from the beginning and look at what we are actually doing when we take a black-and-white photograph.

Grayscale Conversion

First, and essentially, we are making a grayscale image from color. That is, we are converting that color to a shade of another color—gray. Back in the days of film, we didn't have much control over the value—or brightness—of that shade of gray because the film had its own way of interpreting and translating colors. We had some controls, though, and they were based on filtering the light coming through the lens. If you take a color and restrict it, then that will translate as darker, right? If you take a blue sky, and put a red filter over the lens, and "read" its value (with a black-and-white film, in this case), the red filter restricts the blue color. The effect is to make the color read darker.

Grayscale conversion in Photoshop uses the same idea, except instead of filtering the lens (which you can still do, but most people don't), you are selecting what part of the three main channels you want to use to make the color gray. If you use primarily the red channel, blue areas will read as darker. Same idea.

Channel Mixer

As usual, in digital photography, there are about sixteen different ways to convert an image into grayscale, and for each way there is a cult-like following to claim their way is best. I start with Channel Mixer in Photoshop, because it best illustrates what we are doing and how this all links back to our color journey.

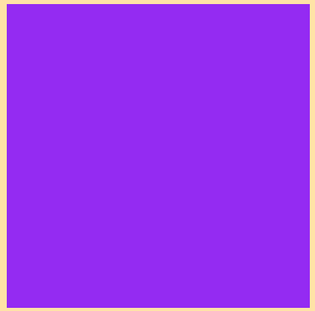

First, let's look at our little burple again. There's the little guy.

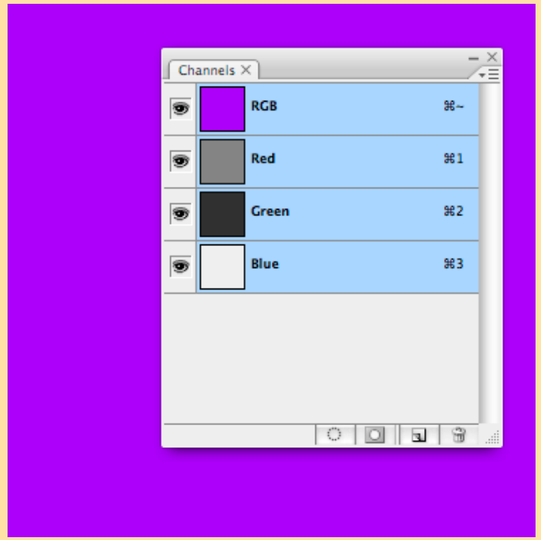

Now, I'm going to open up the Channels pallet. This pallet shows me each separate channel or RGB in the image, and in the case of burple, we can see that each channel is very different. The top channel is our RGB master, with all three combined. The Red channel is a mid-gray, Green is almost black, and Blue is almost white.

Just for illustration, I'm going to select the Green channel. Our little burple goes almost pure black with that channel selected. I've taken the available color channels, that read as very different values of luminance, and decided to render my grayscale—or determine my gray color—based only on the Green channel.

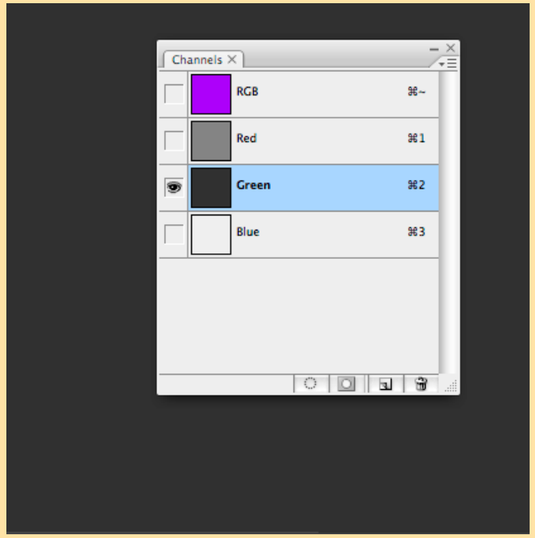

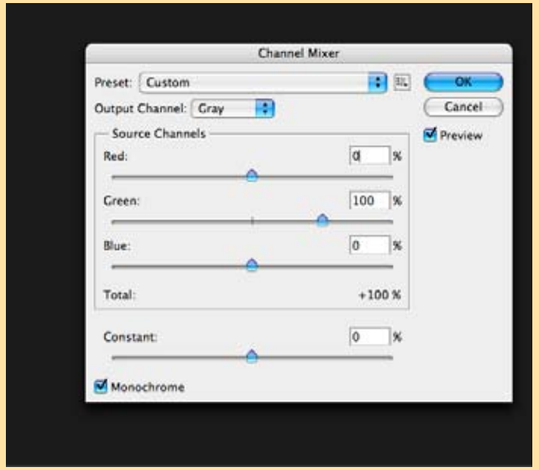

Channel Mixer lets you do that, but to certain degrees. In the Channel Mixer adjustment layer I select the "Monochrome" box. I zero out the Red and Blue channels, and hit 100% on the Green. Bingo—same value of gray.

The cool part is that I can go in and balance this to render my little burple color in any value I want. If I want white, I can use 100% Blue, and zero the others. If I want gray, then I take the Red channel to 100% and zero others. By using a selective mix of my basic color channels, I can control exactly what value I want to give my color.

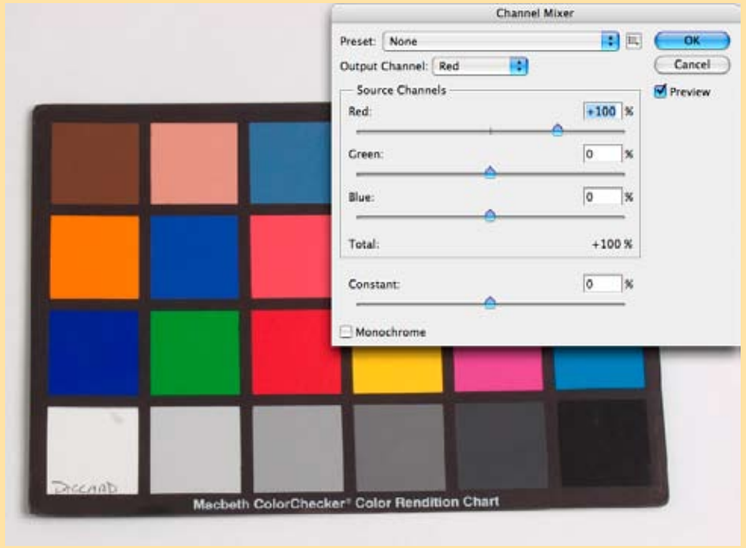

Let's look at a few other colors. I'm going to show you a shot of the ColorChecker here, and a few different Channel Mixer settings.

This first shot is just the target in color, so you can compare the color renderings.

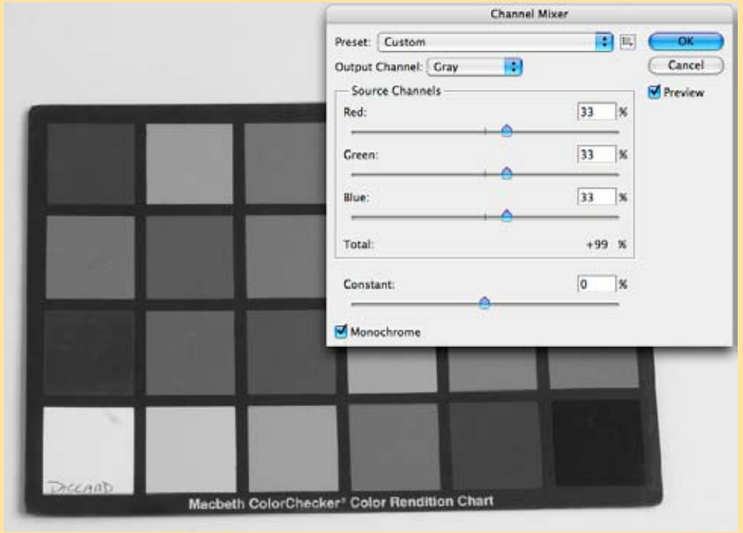

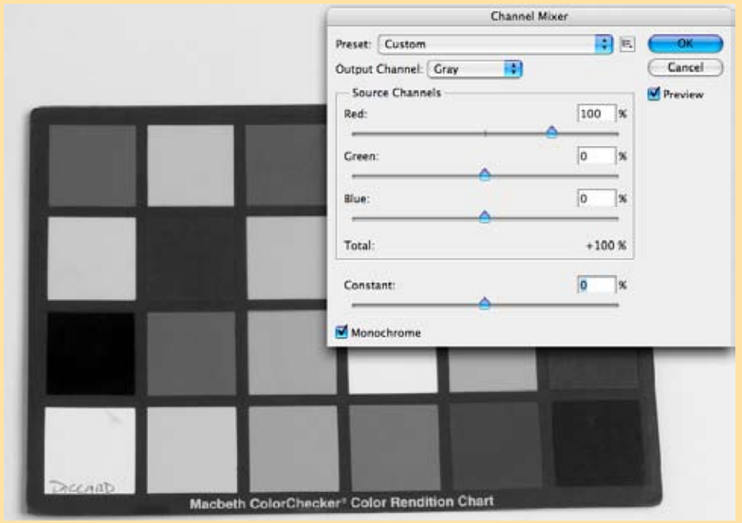

to gray equally. When you switch your "Mode" to "Grayscale" (Image>Mode>Grayscale), this is essentially, if not exactly, what Photoshop does to the file. (Photoshop does apply some adjustments to make the conversion more like what your eye sees.)

Next, I'm showing what it looks like if I use just my Red channel at 100%, and everything else at zero. Look at the red patch, the third from the left and second from the bottom. It has rendered as a very light gray. Now look at the blue patch on the left side, third down from the top—it is almost pure black (and actually looks blacker than the black patch on the bottom, far right).

Just to bring the point home, the next screen shows you some presets in Channel Mixer that are used to simulate film shooting with the filtering I mentioned above. Here is the "Black & White with Yellow Filter (RGB)" setting, and you can see that to simulate that effect, we need 34% of the Red channel, 66% of the Green, and 0% of the Blue. (The only real hard and fast rule I suggest when using Channel Mixer to convert to grayscale is to make sure the three channels add up to about 100%.)

The next shot shows I've selected "Monochrome" in the Channel Mixer dialog and set all three channels to 33%—I'm using all three channels equally. This is what is called a "Linear Desaturation," and this ramps the colors down

So far in our little burple's color journey, we have grabbed the three channels of color with our sensor, put them through our RAW processor, and taken the three separate, single channels of R, G, and B pixels and made one pixel with three channels of R, G, and B. Then, in Photoshop, we asked Channel Mixer to select precisely what percentage of each channel we want to use to create a single channel of gray. (Note that in the Channel Mixer window, our "Output Channel" selection is gray.)

This is a little confusing, because if you look at the channels pallet you still see three channels; the difference is the top channel is no longer in color—it is in gray. If you measure the image with "Info," you will still get three values for each channel of R, G, and B, but they will all read as equal and all three will have the same number on the luminance scale (0-255). We are still working in our AdobeRGB space and with three channels, but all three channels have been rendered to a gray value according to the mix we've asked for.

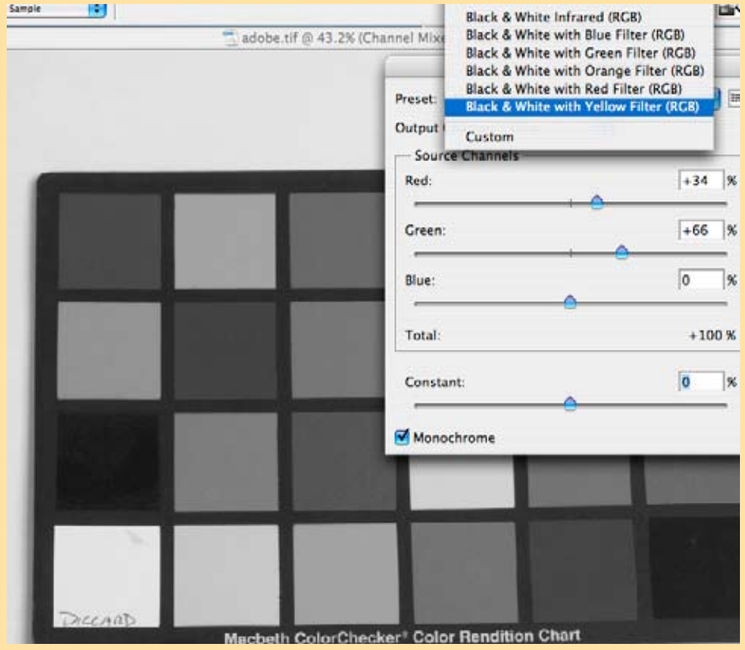

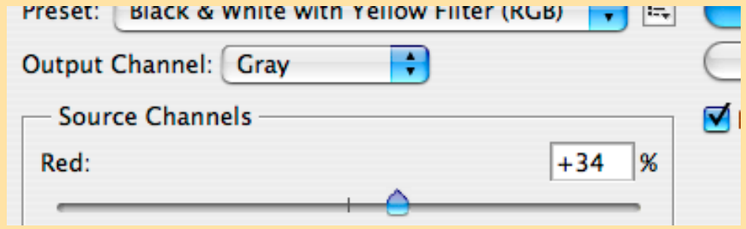

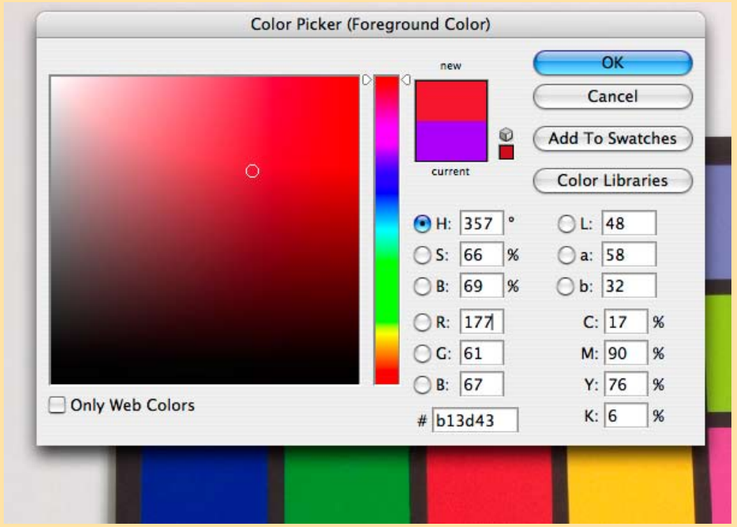

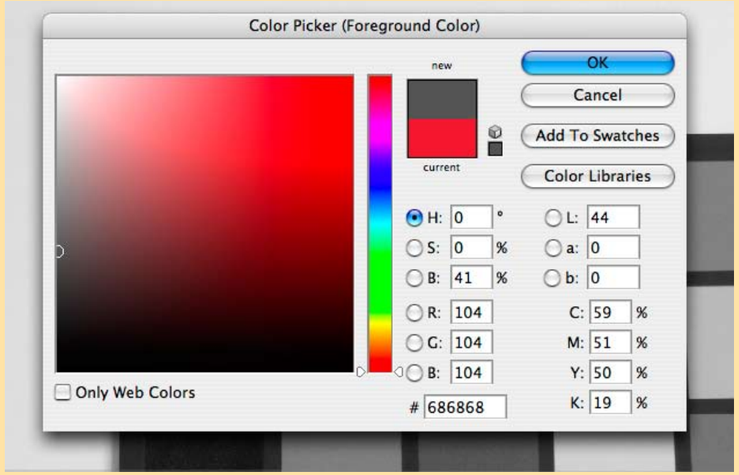

Maybe this will help. Here is the color of the red patch, shown on the Color Picker, before I apply Channel Mixer. The RGB values (seen at the bottom of the left column) are 177, 61, and 67.

When I apply the "Orange Filter" preset, it takes percentages of those numbers, and does the math to make them all equal. Here's what that looks like.

We've rendered the red color to gray, averaging it out to RGB channel values of 104, 104, and 104.

We are still working with our dear little burple, and we are still working in the AdobeRGB color space. We have remapped it—based on its original values—to have all three channels read as equal, and forced it to read a value of gray on a scale of white to black.

Black-and-White Conversion Example

This comes into focus when you see how it actually works in a photograph. Here's a shot of Barbara, processed with default settings in Photoshop. (Barbara will never trust me again when I tell her I'm "just testing a camera.")

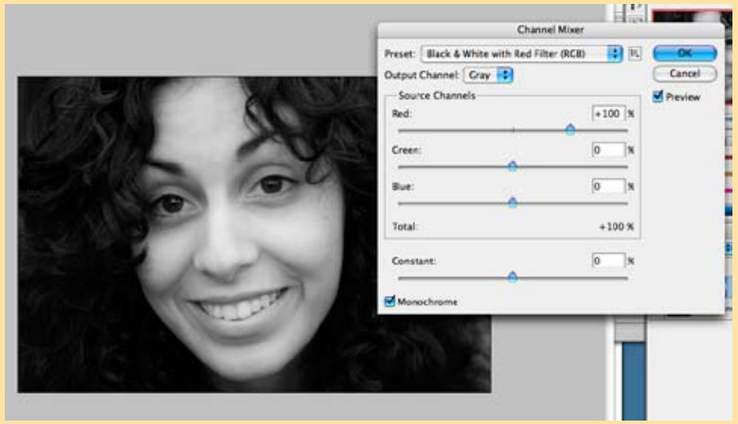

Here is Barbara with Channel Mixer using the G and B channels equally, and none of the R channel. I don't particularly like to render skin tones like this, especially on a young woman. It might work on a crusty old man—the kind of portraits that people describe as having "character"—but these channels darken lighter skin tones and accentuate any marks, freckles, and blemishes.

Generally, for fashion or portrait shoots, I would put a red filter on the lens. Here is what that looks like, done digitally with Channel Mixer. I find that much more flattering, for my purposes.

Again, what we are doing here is taking our basic RGB channels and deciding what color is going to go dark and what color is going to go light. In the second example, I've mapped the red values to go very light, and for this image it accomplishes what I want to accomplish.

Grayscale Conversion in RAW

This is all fine—the Channel Mixer grayscale conversion option—but as you may have guessed by now I'm all about the RAW processing. Adobe calls it a "non-destructive editing" workflow, as opposed to conventional adjustment editing, which is destructive. (Destructive editing in this means you start with a set amount of data, or information, and all you can do is push it around or toss parts of it out. You cannot add any information to the file with the general adjustments in Photoshop.) I like to take that concept to another level. Editing in RAW is actually "constructive" or additive editing, not just "not destructive."

Thankfully, one of the improvements in CS3 was the addition of a Channel Mixer-type grayscale conversion right in Camera RAW. Here's what it looks like. (Click the HSL/Grayscale tab to find it.)

The way I use this is to first look at the three main channels, R, G, and B, and set those where I need them. Then I use the intermediate channels—or more correctly, controls—and smooth the settings to give it a nice transition from one main channel to the other. The settings shown are the standard Photoshop setting, and are a basic film-type response.

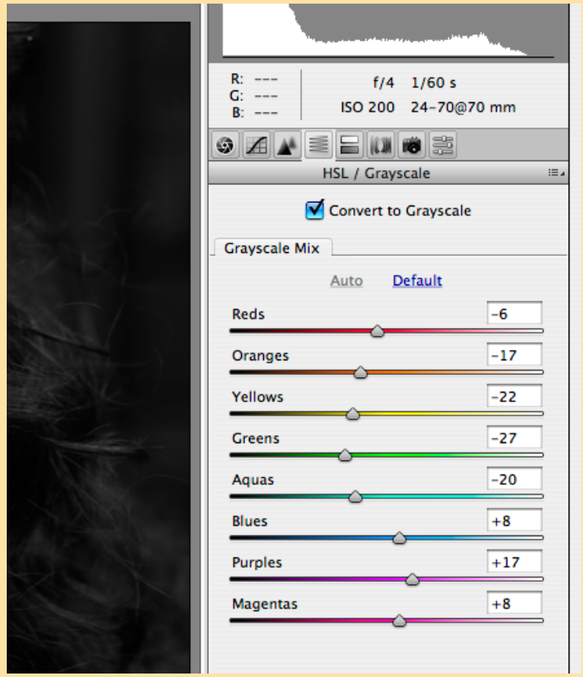

Grayscale Conversion in Adobe LightRoom

Adobe's LightRoom uses the same processor as Photoshop and Adobe Camera RAW, but has some very nice tools for controlling it. Here is a shot of LightRoom's HSL/Grayscale tab—check out the tiny button in the upper left corner.

Click that little icon, place your curser on any part of the image, and drag it up or down to make it darker or lighter.

In this next screenshot I'm showing the color target where I've clicked the red patch and dragged it up—mapping the red tones brighter and subsequently dragging the green tones

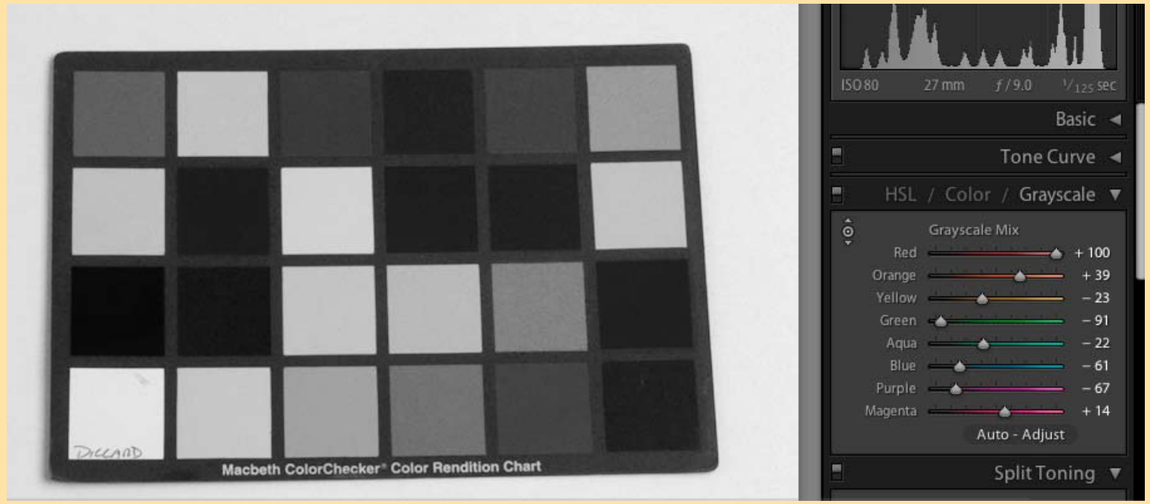

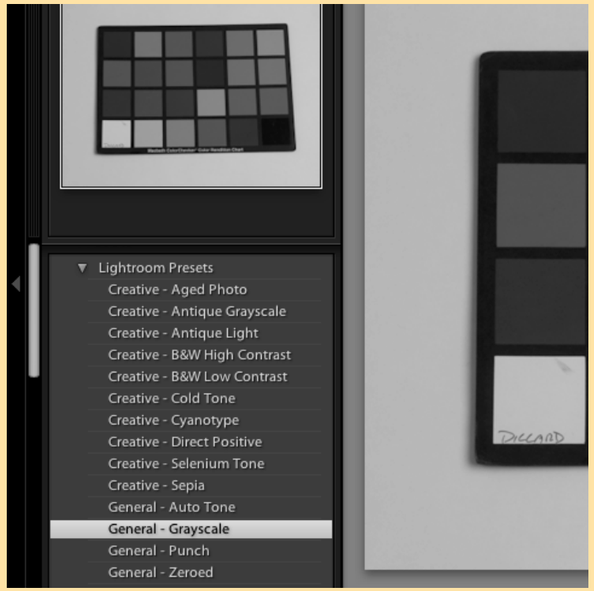

down. Compare the settings for that one with the default settings I showed you above. This is one of the many reasons LightRoom is getting such a great reception from photographers who need to work fast, but need to work well.

LightRoom also has some handy presets, on the far left of the "Develop" screen. Here I've selected the "General Grayscale" setting.

It also has User Settings, where you can start with one of their presets, change them, and save them as your own, or you can just start from scratch.

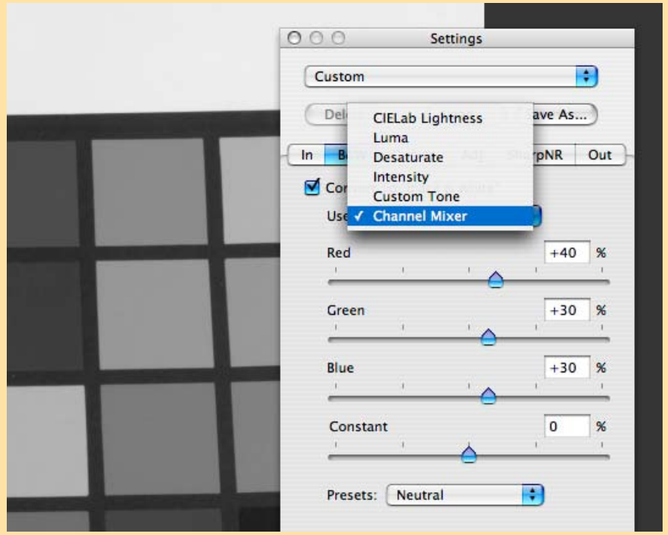

At this point I really don't know how things are going to go with LightRoom. Adobe typically "cherry-picks" nice features from other software and then builds the cool stuff into Photoshop. My guess is that you're going to see some of these features built into new versions of Camera RAW and have them available in Photoshop.

Grayscale Conversion in Iridient's Raw Developer

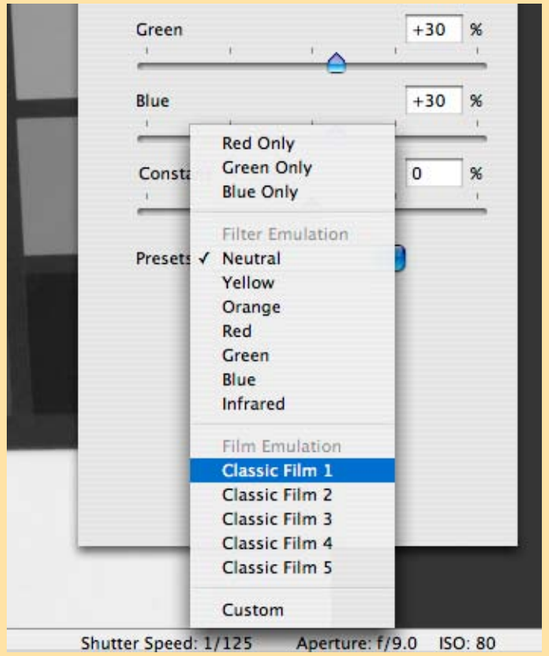

I am throwing this out there to support Iridient Digital's hard work, and taunt Adobe with a powerful and professional set of features in RAW processing. Raw Developer has a nice array of options for grayscale conversion in RAW. Here is a shot of the controls.

Under the Black-and-White tab is a great selection of options for converting the RAW RGB information into grayscale values, including using the "Lightness" channel from LAB. I still think the Channel Mixer option is by far the most powerful and gives me the most control, but it's nice to work with a package that gives me more to choose from.

You also get a pleasing little array of presets—some mimicking film renderings—as well as the ability to save your own settings like most other packages.

There are a couple of other RAW processing packages out there—one in particular—that has film "profiles" which again mimic not only the color response of a black-and-white film, but also the grain. I'd show you this, but it would turn into a listing of independent RAW processors and a feature shootout pretty quickly.

Suffice to say, they are all doing the same thing, but with a few different knobs and switches. They are all taking that basic information from the sensor, in red, green, and blue, and figuring out how bright to map that color. The file is then built into the working color space you select and sits on the same three dimensions that any of our RGB colors sit on.