



By Dan Margulis

COLOR, CONTRAST AND L*A*B* (PART I)

Tradeoffs, tradeoffs. Our life is full of them. Tradeoffs between style and substance, between price and performance, safety and practicality. In graphic arts, or more specifically, in the presentation and handling of photographic images, the most haunting tradeoff is the one between color and contrast, two concepts that are so tightly intertwined that it seems impossible to alter one and not the other.

Among many artists, there is a sneaking suspicion that this tradeoff is not inevitable, and that part of the solution lies in the exploitation of a frightening item known to cognoscenti as CIELAB, L*a*b*, or just plain LAB.

Being that this column is devoted to practical production and not frivolous vendor hype, it usually draws a fair response from readers. But of all the pieces appearing in this space over the years, the most popular by far, to judge by an overwhelming favorable reaction, was one that I did almost as a throw-away a year ago, giving a quick and dirty method of improving almost any Photo CD image. That column's successful recipe (The Five-Minute Photo CD Gourmet, *Computer Artist*, June-July, 1995) largely depended on LAB maneuvering, a topic that, as far as I know, had never been addressed elsewhere.

The L channel of LAB, which contains no color information, is ideal for sharpening, since it avoids the weirdly-colored halos that characterize over-sharpening in other color-spaces. Left to right: two images with unsharp masking (highly exaggerated) applied in the L channels only; the unsharpened images; the same USM settings applied to RGB documents.

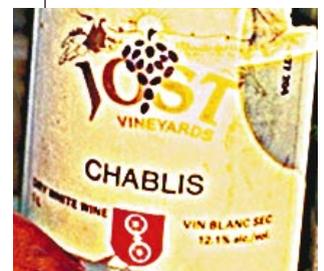
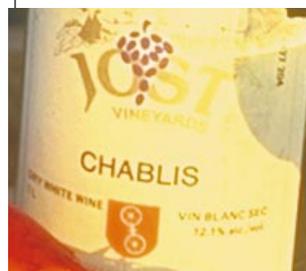
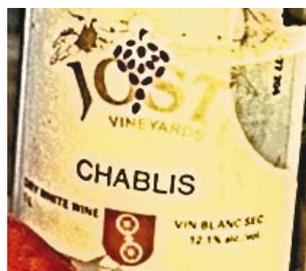
The recipe featured two things LAB does extraordinarily well—sharpening and setting overall range—but these are merely the beginning of the LAB story.

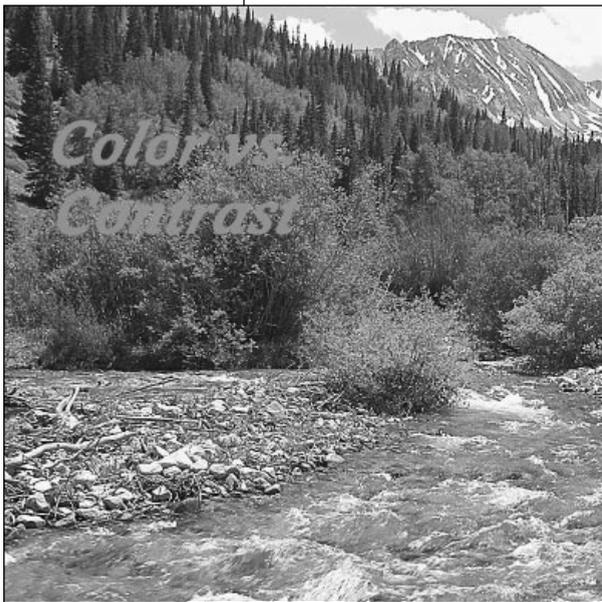
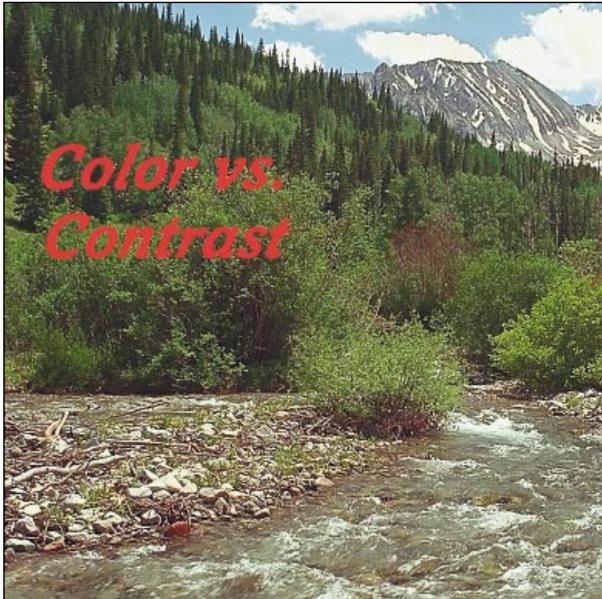
Virtually everyone, giving the tricky LAB alternative a wide berth, retouches and color corrects either

in RGB or CMYK. There are rabid partisans who favor each one, but unless you are engaged in high-end color-correction, either is workable (The Natural Superiority of CMYK, *Computer Artist*, Oct.-Nov. 1995). The two are in fact quite similar, and it is pretty easy for anyone familiar with one to work in the other.

LAB, however, is radically different. Like many radical approaches, it is very good in certain ways and very bad in others. We will now explore the good part, of which there is more than can fit in one column. This time we will deal with straightforward retouching and range-setting. Next time, we will steam straight into the uncharted waters of curve-based correction in LAB.

LAB is radically different. Like many radical approaches, it is very good in certain ways and very bad in others.





The characteristic gray blurs that constitute the A (left) and B channels are shown below. There is no need for detail, which is carried in the L channel (above). The red-green A here has a little more action than the yellow-blue B, especially in the type.: the lighter color represents red in the A, yellow in the B.

COLOR AND CONTRAST

LAB can be exploited precisely because it separates color from detail, solving many problems that would be extremely difficult in either RGB or CMYK. If you customarily work in RGB, you can move in and out of LAB freely to do these things, because there is no quality loss associated with the transition. If you are already in CMYK, there is a minor quality hit in going to LAB, but this may well be outweighed by the benefits of being there.

LAB has three channels, but, as you can see below, the A and B are scarcely recognizable. They define color only, not shape, and are thus unlike anything found in RGB or CMYK. Not only are they therefore blobs, they are *gray* blobs. The LAB color gamut is far wider than anything we are accustomed to. As a result, only about half of the possible tones in these two channels are in use. There are no whites and blacks—only grays, grays, grays.

The two can nevertheless be understood with a little effort. They are, as a scientist would put it, interactive opponent-color channels. The English translation is as follows. In the A, lightness represents an area that is more red than it is green, darkness something that is more green than red. The lighter (or darker), the more intense the color. A pure gray represents, well, a pure gray. The B works in the same way, except lightness means more yellow than blue and darkness more blue than yellow. This explains why, whereas the A in our sample image is appallingly flat, the B is blood-curdlingly so. At least in the A there is some variation, because there are

objects in the image that are obviously red or obviously green. The B is a blur because the type is only modestly more yellow than blue and the trees only slightly more blue than yellow.

The L is far more mundane. It can be seen as a crisp black-and-white version of the image.

A recurring theme is this: Without LAB, whenever you apply a filter, a curve, or a tool, there is a good chance that whatever the weakest color is will take it on the chin. That is a general recipe for making whatever the dominant color is more brilliant, more vivid. Often enough that is the last thing you want to have happen.

If you work in the colorless L channel only, it won't.

To start with, unsharp masking is more effective when applied to the L than to any combination of channels in RGB or CMYK. On the opening page of this column, you can see why. The idea of USM is to exaggerate transition areas to make the image seem more in focus. It does so by placing dramatic-looking borders around objects that are butting: the darker one gets an even darker border; and the lighter one gets a light border; so that instead of going light to dark, we go light-lighter-darker-dark, in most convincing fashion.

If USM is applied to an entire document, this is equivalent to applying it separately to each channel. In the examples, I sharpened in RGB, but the results would have been very much the same in CMYK.

The yellow wine label, when sharpened normally, darkens the type as it should, but yields a revolting white halo around it, the result of an exaggerated transition in the blue channel, which gets replicated in the yellow plate. A different disaster takes place in the second image, where two contrasting, complex colors meet at the edge of the red uniform. In the RGB file, there is not much transition to exaggerate in the blue channel. The red and green channels have transitions, but they go from light to dark in opposite directions.



Vivid colors are wonderful things, no doubt, but they are not exactly what we want from USM. Notice how the woman's neck has become green where it touches the red uniform, and how the edge of the uniform has become such a brilliant red that it has lost detail. The yellow star on the uniform, and the red logo on the wine label, have also become much more intense.

These color-related problems go away if we sharpen the L instead. On the wine label, the yellow becomes lighter where it butts the type, but it clearly is still *yellow*, not that disgusting white. Where the red uniform hits the background, it is again surrounded by a light halo, but the halo is not the eye-popping green of the RGB sharpening.

AB STANDS FOR ALWAYS BLUR

Photoshop's extensive defocusing filters, Despeckle, Median, Dust & Scratches, and the various Blurs, are the digital equivalent of going to the dentist. You do it, I do it, we have to do it sometimes in emergencies, and it almost always hurts.

Barring the case where we are trying to get an out-of-focus look, these filters always damage detail, and can only be justified on a lesser-of-evils basis, as when an image is extremely noisy. But, certainly, they should be avoided whenever possible.

Is it not a happy thing, then, that LAB, unlike other colorspaces, has two channels that carry no detail at all, but only color? We can blur the A and B almost to our heart's content in a typical picture, for exactly the same reason that we

can apply heavy sharpening to the L. If there *is* detail in the A and B, it is probably detail we don't want, namely, noise or damage. Therefore, Gaussian Blur settings of as high as 3.0 radius, which would normally destroy an image, may be acceptable in the A and B.

Quite apart from the noise and graininess introduced by fast films and/or imperfect shooting conditions, anyone who has ever scanned knows that the blue channel in RGB, which becomes the yellow in CMYK, usually is home to more static than might be hoped for. Yellow being a weak printing plate, a lot of us blithely apply a Gaussian blur to it, accepting the modest loss in detail, and never thinking that we could have solved the problem with no detail loss at all by the simple expedient of blurring the B.

Grayish noise is a lot less obtrusive than noise with a color component. Eliminating noise in the yellow-blue B channel (and in extreme cases, the red-green A) is therefore highly effective, but there is an even more obvious example of when to use this technique.

Prescreened color originals, such as those shot from a printed page, are notoriously difficult. There is no known way of getting really good results; about the best that can be hoped for is something south of satisfactory. The traditional best method is to scan slightly out of focus at an extremely high resolution, and then down-sample, hoping to lose the offensive dot pattern on the way down.

The fact that a pattern is detectable in a file does not mean that

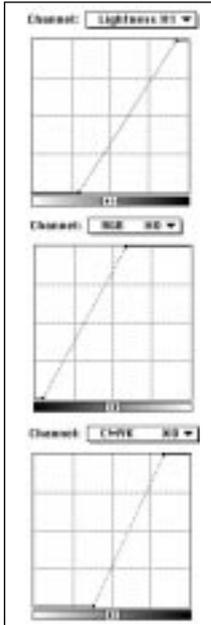
there will be a moiré in print. But if the patterns happen to be magenta, yellow and cyan, moiré becomes highly likely. If we blur the A and B, they won't be. The pattern will be strictly monochrome, and on a good day, we may be able to retouch some or all of it out of the L channel.

Another delightful use of this idea is in collaging. When we merge an object into an existing image, usually we have to indulge in some kind of anti-aliasing to avoid jagged edges where the two meet. The unfortunate thing is that, depending on how good our selection was, the backgrounds of the two images and a couple of other things, we may get funky-looking transition colors right at the fringe where the two objects meet.

In such a case, if retouching the weird colors out looks at all difficult, save the effort. Go into LAB and blur the A and B.

With LAB, there can be a ceasefire in the battle between accurate color and good detail. The bottom left version is correctly neutral in color, but dull and lifeless. The bottom right rendition has nice detail, but lousy color. The version below therefore combines the A and B from the bottom left with the L from the bottom right, letting this feline have its cat food and eat it too.





RETOUCHING AND REPAIRS

In the outdoor image on the second page, please make the following irritating assumption. Suppose that some knucklehead has merged and saved the green type into the image, without doing so in layers, and that we now find out we have to delete the type and restore whatever was beneath it. Knuckleheads, as you know, do not have backups.

At such a time, the last thing you probably want to hear is some other knucklehead suggesting that you might want to repair the damage not by cloning, but with Photoshop's smudge tool.

That seems absurd—unless you happen to be in LAB.

However it is that you propose to make this repair, you will surely concede that it is much easier to fix

the L channel shown here than the entire color document.

All you have to do to exchange the tough assignment for the easy one is to hire a four-year-old to erase the type from the A and B. Doing so is a snap. Most four-year-olds will have no trouble once you explain to them how the smudge tool works. If you care to do the work yourself using a more sophisticated method so much the better. The resulting color variation will be quite sufficient to hide the repairs needed in the L.

The L is ideal for retouching. The sharpen tool works better in the L for the same reason that the unsharp masking filter does. Dodging and burning is more effective in the L as well, since it lightens or darkens the image without introducing a new basic color.

I suggest that when someone asks you whether the best color-space for general retouching, repair and collaging is CMYK or RGB, that you answer, "none of the above."

THE BEST WAY TO SET RANGE

The theme of unwanted brilliance also pops up in that most important of color-correction techniques, setting the highlight and shadow. It can once again be avoided neatly by employing the colorless L.

A full tonal range is absolutely the most critical element of color reproduction. If you do not have one, regardless of your creative abilities, channel-maneuvering cognition, etc., you will never be able to beat even a modestly talented person who does no more than set a proper highlight and shadow.

Scans ordinarily fall somewhat short of the ideal, so it is up to us to fix them up. This can be done in RGB or CMYK with curves or, for the timorous, levels. To do it right, unfortunately, you have to do it separately in each channel. Photoshop tempts us to fall from grace, by including a wicked master-curve function that seems ever so much easier.

Setting range with a CMYK master curve will be a disaster, because, even if it gets the three color plates

correct, it will ruin the black. An RGB master curve move does not have this drawback, but it has the nasty habit of taking little problems and making big ones out of them.

Although in principle the master RGB curve treats all three channels equally, in practice it only works that way if the three were equal to begin with. If we are trying to lighten the highlight with it, whatever color was weakest to begin with will be hurt the most. That, as we have already seen, guarantees a cleaner, more vivid, more pronounced dominant color, exactly what we don't want in a highlight.

If you are dead set and determined to apply only one curve and not three, the answer is not the RGB master curve but the L of LAB, which can set range without altering the prevailing color.

This is easy enough even for the curve-hater. To do it in the simplest possible way, convert to LAB, and locate both the lightest and darkest significant areas of the image. In Photoshop's Info Palette, set the options to read LAB color. The L channel will be defined as 100 (white) to 0 (black). The extreme numbers will not show detail on the printed page. For average printing, the minimum and maximum values should be, roughly, 95 and 10. Move the end-points of the curve in toward the center until the known light and dark points reach these values.

Sharpening the L, blurring the A and B, using the L channel to establish range by shortening the curve. The lion of LAB does not seem nearly so fearsome now, does he? For pretty much *any* image that starts in RGB and is headed for CMYK, it pays to take an intermediate stop in LAB for sharpening, range correction, and, if needed, reduction of noise.

Dan Margulis is a Computer Artist contributing editor and author of Professional Photoshop, published by John Wiley & Sons. Readers may reach him by fax at 201/763-2835 or by e-mail at 76270.1033@compuserve.com. For information on Dan's three-day, small-group color-correction tutorials, offered in Atlanta, contact Sterling Ledet or Ledet & Associates at 404-325-3338.



If you don't like making channel-by-channel color corrections, you are far better off setting overall range in LAB. Extending range with the master curve in RGB will exaggerate any color imbalance, since the weaker, contaminating colors get hurt the most. In CMYK, using the master curve will lead to certain disaster. Using the L channel instead will not alter color balance. The flatter the original, the more you need LAB, as this possibly recognizable image shows. At left, top to bottom, the results when the correction was made in LAB, RGB, and CMYK.