

The Power and the Peril

LAB has a reputation as somewhat of a wild animal: very powerful, very dangerous. The power is evident: spectacular, almost instantaneous revivifications of color. Transformations of extremely washed-out images into ones that are professionally usable. The best space for unsharp masking. A much faster way to transform colors, as when a car is red but the client wants it green. A more flexible way of retouching and repairing damage.

The dangers of working with LAB mostly come from using it where it isn't appropriate. LAB can make bigger and faster changes than any other colorspace, at the cost of a considerable amount of precision. If your workflow calls for fixing a lot of images in 30 seconds to a minute apiece, LAB is probably the best choice for correction. For anybody with more time for finetuning than that, LAB should be a sometimes thing, rarely the final move.

In addition to retouching, LAB is at its best with certain types of grossly defective images and with ones that are excellent to begin with, particularly those that need the introduction of subtle color variation. Some of the curving in the A and B channels is best reserved for experts but there are several powerful curves and blends that anybody can learn to use.

These two sessions presuppose a general knowledge of unsharp masking, blurring, basic retouching, and the use of curves. We will not be discussing third-party plugins, color management concepts, or specific output devices.

All curves shown will have the white point as the origin and the dark points to the top and right. Thus, an upward move in the curve will darken the image.

An Unlikely Way to Spec Color

LAB's strength is that it completely separates color from contrast. All the detail, but no color, is contained in the L, also known as the Lightness or Luminosity channel. The L is easy to understand: consider it to be a black-and-white version of the color image. Curves and retouching work in the L exactly as they would in a grayscale image. The numbering system is strange: 100 is pure white, 0 pure black. Because the L contains all the detail, manipulating it offers more bang for the buck than any other method of adding contrast. As against that, it's a lot easier to plug shadows and/or blow out highlights in the L channel than any other method.

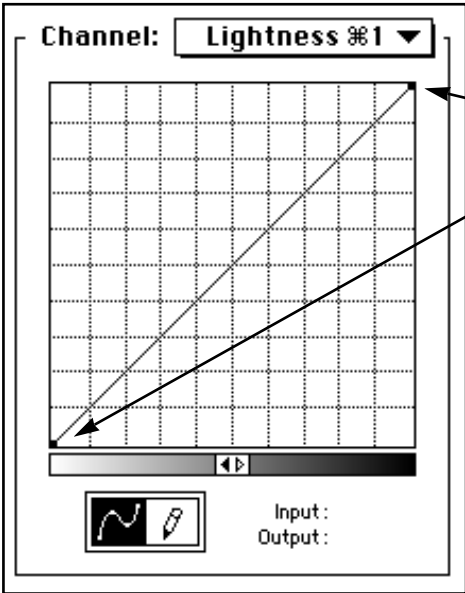
All color information, but no luminosity, lives in the A and B, which don't stand for anything and which are opponent-color channels. The A runs between magenta and green; the B between yellow and blue. The baffling numbering system runs from -128 to +127. Negative numbers are green in the A and blue in the B; positive numbers magenta and yellow respectively. However, only about half these AB numbers are realistically usable: anything more than 75 points removed from zero is a color that's likely to be out of not just the CMYK gamut but that of RGB devices such as a monitor as well.

In both channels, a zero represents neutrality. If an object that's supposed to be neutral reads zero in either the A or the B

The L channel contains all of the contrast, but none of the color information. It's easy to understand: just consider it as a black-and-white version of the color picture.



Dan's
**Handy
L*a*b*
Cheat
Sheet**




Zero = Completely Black

100 = Completely White

Typical range should be white = 95, black = 3-5.

If you are worried about losing shadow detail, choose a lighter black.

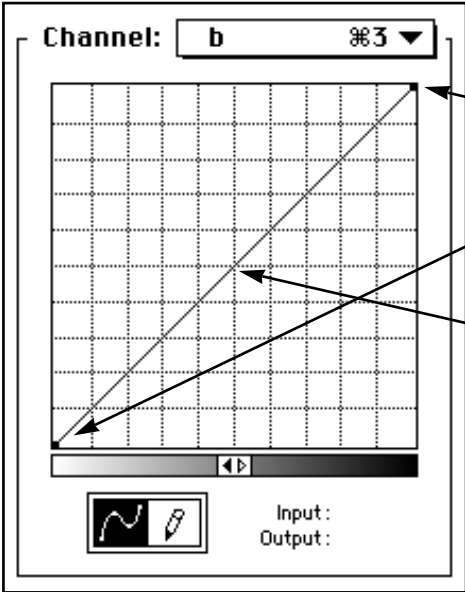


Minus 128 = Absolutely green (unprintably so)

Plus 127 = Absolutely magenta (ditto)

Zero = Neither magenta nor green (neutral)

Values above 75 and below -65 are generally out of CMYK gamut



Minus 128 = Absolutely blue (unprintably so)

Plus 127 = Absolutely yellow (ditto)

Zero = Neither yellow nor blue (neutral)

Values above 85 and below -55 are generally out of CMYK gamut

Toward
GREEN


Toward
MAGENTA

Toward
BLUE

Toward
YELLOW

The Magic of LAB

notes



The A and B channels are blurry because all detail is contained in the L. And they're gray because anything white or black would be so wildly colored that no monitor could display it, let alone any output device. That said, lighter grays represent warmer colors (magenta in the A, left, and yellow in the B); darker grays are cooler colors (green in A, blue in B).

channel, that channel is correct regardless of what's happening in the other. This differentiates LAB from RGB or CMY, where all the channels depend on one another: you can't tell whether any channel is neutrally correct without hearing the values in the other two.


The A and B channels look like blurs, as they contain no contrast information. Furthermore, they look like gray blurs, as anything white or black would be impossibly out of gamut for any realistic use. Lighter grays represent magentas in the A and yellows in the B; darker grays are greens in the A and blues in the B.

The Simplest Curve Is a Straight Line

Adjusting the L channel is easy for anyone who knows how to handle a grayscale image, which is all that the L is. The A and B are nastier, because so many of their colors are so brilliant as not to be of any use in output. Any adjustment applied to them is a big one. Even changes that would be small in RGB or CMYK can alter colors radically. Accordingly, standard curve shapes become experts-only moves in LAB.

There is, however, one simple and nearly foolproof way to use AB's power to enhance color. It takes advantage of LAB's unique definition of neutrality. Rotate the curve counter-clockwise around the midpoint—that is, keeping it a straight line and maintaining the



center point in its original position, with horizontal areas at the foot and top of the curve. Since the midpoint of the A and B curves is neutral, anything that starts out gray will stay that way. However, there will be more variation in both hue and saturation in all colored areas. There's no analog in RGB or CMYK to this maneuver.



A standard way of increasing variation in colors with little risk of complications: steepening the A and B curves by pivoting them around the midpoint. As the midpoint is neutral, if the original picture has no cast, doing this won't introduce one.



When the original has a color cast, the same steepening technique can still be used. However, the curves can no longer pass through the original center point, as this would not just retain the cast, but strengthen it. Instead, it must pass to the right of the center point if a warmer color is needed and to the left to produce a cooler one. Measuring typical points in the snow demonstrates that the above image starts with a slight green cast from the A (negative number) and a strong yellow cast in the B (positive number). To correct this, the A curve moves slightly to the right, or warm side, and the B strongly to the left.



The Straight Line and the Color Cast

Using straight-line AB curves as described above enhances all colors, especially those that weren't nearly neutral to begin with. But it also will enhance any undesirable color cast that may be found in the original. To solve this problem while retaining the simplicity of the straight-line curve, identify the cast by checking out known neutral colors. Zero is the desirable value in both channels. Negative numbers indicate a bias toward coolness; positive toward warmth.


If the channel shows a cool cast, use a steep straight line that passes to the right (warm) side of the center point. If the channel was originally too warm, pass the curve to the left of the center point.

Dividing the AB Curve in Half

The A and B channels each have a warm and a cool half. Often one wants to enhance one half and not the other. For example, shots containing a lot of greenery usually can benefit from a steeper A curve. Any red or purple objects in the same shot, such as roofs of houses, may cause problems if they get brighter.



Steepening the A and/or B channels uniformly brings out variation in all colors. Sometimes we need to enhance certain colors but not others. The lawn and trees at this tropical resort could stand to be more green. But if we steepen the A channel in the way shown previously, it will help the greens but also create candy colors for the roof and columns, which are heavily magenta. In such cases it's helpful to consider the curve as two halves. The top half affects cool colors (green in the A, blue in the B) and the bottom half warm colors (magenta and yellow). Here, the bottom half is locked while the top half gets steeper.



By treating the curve as two halves, usually dividing it with a point near the middle, the change can be limited to cool or warm colors. In fact, one can even suppress the colors of one half while exaggerating those in the other.

When to Avoid LAB

The volatility of the A and B channels makes it possible to crush any uniform color cast, no matter how pronounced, without changing the detailing of the image. Some people, impressed by spectacular demonstrations of cast elimination, head for LAB any time they see an image with a bad cast.

While LAB may be needed, or at least valuable, in smashing a cast, before starting on an adventure with LAB curves consider the possibility of a simple channel blend. If an image has a serious color cast, the chances are that one channel is so seriously out of whack that attacking it with a blend can reduce the cast and increase detailing at the same time.

For example, images with strong yellow color casts are too dark in the blue (RGB) or yellow (CMYK) channels. Blending some other channel into them (using Image: Apply Image or one of several other commands) is likely to be easier than a gross curve in LAB.

A second peril: LAB cast reduction only works when the cast is uniform—that is, it exists at all levels of lightness in the image. If there's a yellow cast in the highlights and no cast (or

a cast of a radically different color) in the shadows, LAB can only make matters worse. Unfortunately, these non-uniform casts have become more common in the age of cheap digital cameras, which often have automatic routines that attempt to balance the white and dark points.

Those who correct in LAB for eventual conversion to CMYK know that it is ridiculously easy to create colors that are out of the CMYK gamut. This may not be a bad thing, provided that they're slightly, and not preposterously, far from what CMYK is capable of reproducing.

But many people fall into the trap of forgetting that the LAB file may be out of the RGB gamut as well. Photoshop gives no warning when this happens. Oddly, this phenomenon is more dangerous to those whose destination is CMYK than RGB. If parts of the LAB file are outside of the RGB gamut, the monitor won't be able to display them accurately. What the monitor shows is nevertheless a good approximation of what will happen if the file is converted to RGB. A direct conversion to CMYK, however, can be a horse of a *really* different color.

Ordinarily, converting an LAB file to CMYK will produce almost exactly the same file as converting it to RGB and then to CMYK. When out-of-RGB-gamut colors are in play, though, the results can be wildly different, creating very unexpected CMYK results. For example, 100^L43^A48^B is a fluorescent, electric orange-yellow that a monitor has zero chance of emulating. Instead, it shows what I would describe as a pale lemon yellow.

A copy of this LAB file, when converted to RGB, will look almost identical. But when these two apparently identical files are converted to CMYK, they become very different indeed. On my system (your numbers may vary slightly depending on your Color Settings) the RGB version converts to 20^M44^Y, a reasonable approximation of the yellow that the monitor was showing. But when the LAB file is converted, it goes to 34^M38^Y, which is not even a yellow at all, but a pink.

Solution: if you think you may have produced LAB colors too intense for the monitor to display, don't convert it directly to CMYK. Convert it to RGB first. This will lock in colors that approximate what the monitor, however erroneously, is currently showing.

Sharpening the L Channel

Probably the most widespread use of LAB is as a sharpening space. Indeed, applying unsharp masking to the L channel only is one of the most effective means of adding snap to an image. Since the L contains no color information, certain disagreeable artifacts of conventional sharpening are avoided.

This, however, isn't the only advantage, or even the major one. Some—including Adobe itself—have mistakenly suggested that applying USM in RGB or CMYK in Luminosity mode gives the same result as sharpening the L channel by an equivalent amount. It doesn't. The L channel does a better job in differentiating darker shades, which is a major objective of USM. There's a particular advantage in the deepest shadows.

There's no sense in converting to LAB for sharpening in certain cases, but in others there can be an advantage. For example, a model with blond hair against a background that isn't important will sharpen just as well in RGB/Luminosity as in the L channel of LAB. Give the model dark brown hair, however, and an LAB sharpen will give it an extra highlight that can be very attractive.



Blurring the A and B channels, especially the B, is the best way of getting rid of digital noise. Above left: the blue channel of an original RGB digital capture. Above right: after converting to LAB, the B channel, being yellow vs. blue, shows the noise, which can then be blurred out without affecting detail. Right: after the blur to the B and reconvert the image to RGB, this is the new blue channel.



Should AB Mean Always Blur?

Excessive noise in certain channels is the hallmark of much digital photography—as well as sports photography, any other kind of photography involving rapid exposures, shots taken under poor lighting conditions, underwater photography, scans of older originals with aging emulsion layers, etc., etc., etc. The problem is especially bad in the blue (RGB) or yellow (CMYK) channels, although it can occur anywhere.

Confronted by this, many users automatically start applying Despeckle, Dust & Scratches, Median, or Gaussian Blur to the offending channels. This may get rid of the noise, but it will also affect contrast—unless you do it in a channel that has no contrast, namely the A and/or B of LAB. It's usually easy to pick a high enough setting to kill the noise without making a noticeable difference to color transition areas. Some noise may remain in the image, but it will be grayish noise, which is much less obtrusive than colored noise.





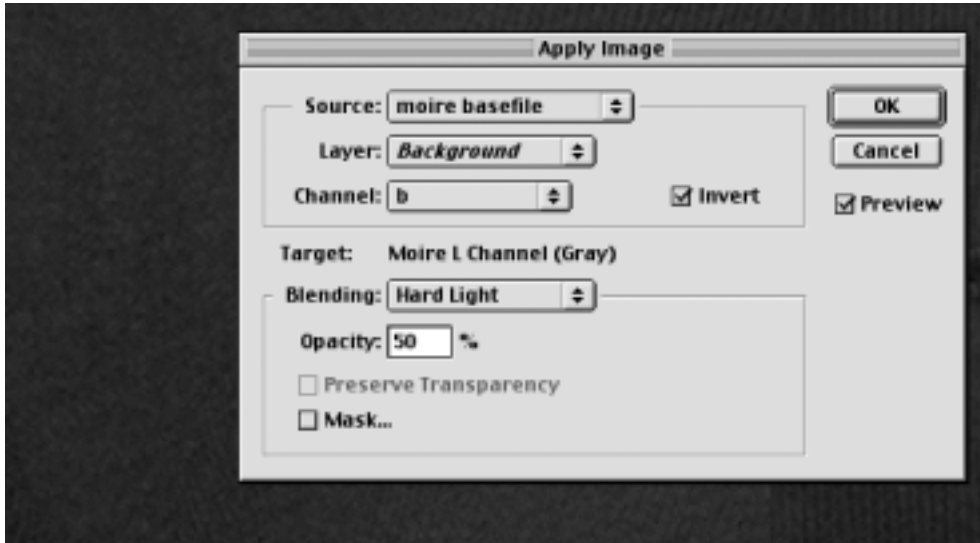
Conquering Moiré

Blurring the A and B channels is never more effective than when dealing with that curse, the prescreened original. The original file always has a pattern of cyan, magenta, yellow, and black dots, which is sure to conflict with the CMYK screens that need to be imposed when the job is reprinted. This makes moiré inevitable without some kind of operator intervention. But the pattern is also the only detail in the image. The mistaken idea that one has to blur contrast-containing channels (in other words, any RGBCMYK channel) has led to some very mushy rescreened images.

Blurring the A and B channels only is a

Although colored moirés can easily be eliminated in the A and B channels, when the L channel is also damaged (left) it gets tougher. Here, the moiré is more pronounced in the B channel than in the A. The B is therefore the best choice for repairing the L.





Applying the B channel to the L in Hard Light/Invert mode eliminates the moiré in the L without damaging the texture of the fabric, as most other methods would. Soft Light and Overlay modes would also work.

clear improvement. There will still be a pattern of dots but it won't be dots of any particular color, as the color of each will blend with that of its neighbor. This eliminates press moiré. The prescreened original may still look grainy, but if it prints at a small size or isn't especially important, one can get by with the AB blur alone.

Moiré has also become more prevalent in the age of digital captures, particularly of people wearing clothing with distinct patterns. If the moiré only shows up as bad color, it can be eliminated in the A and B channels. In the more complicated situation where the moiré also affects luminosity, it can often be wiped out by an unusual blend usually involving the B channel, occasionally the A.



An Unlikely Source for a Mask

Experts in mask generation always suggest looking for the start of the mask in one of the color channels. The last place one would expect to find such a thing is in the formless, gray, blurry A and B channels. And yet it's often there, although it usually has to be augmented by a contrast-enhancing curve.

Strongly colored objects are usually well defined in one or more of the RGB channels. So far, so good—unless there are neutral objects that are of approximately the same darkness.

The Magic of LAB

But in A and B, neutral areas, regardless of how dark they are, are distinct from colored areas, making the selection easier in many cases.

One important category is in images that are cross-lit—that is, that have a different cast from one side of the image to another. Selecting the casted area only requires major effort in RGB or CMYK, but is often very easy in the A and B channels.

AB Curves for Desperados

The L channel is so intuitive that curving it shouldn't present a problem. The A and B, however, are so volatile that one usually wants to stick with the simple curve shapes described above.

Some times, though, the situation is so grim that one has to throw caution to the winds. There are (at least) three occasions where this might happen.

- If there's a catastrophic overall color cast (but first, see if a channel blend can ameliorate it.)
- Where the color of a certain object has to be livened up but the image features something else of that same color that's on the verge of being out of gamut.
- Where a certain color is already out of the output device's gamut and there's a closely similar in-gamut color elsewhere in the image.

In any of these cases the likely approach is to place points on the AB curves that correspond to the positions of the interest objects, and spread them apart or bring them closer together as the case may be. In doing these, great care must be taken with the center point of the curve, unless you're sure you can correct any damage once you've left LAB.


Using LAB Without Being There

LAB maneuvering has enough advantages that it's sometimes helpful to use them without being in LAB at all. The two blending modes that allow this are Luminosity and Color.

For example, if you have two versions of an image, one of them with good detail and the other with good color, the pristine way of merging the two is in LAB,



Eccentric AB curves are sometimes needed to spread specific colors apart. The woman is wearing a brilliant blue sweater, which will be muted severely in CMYK. The chair she's seated in is slightly blue, and will compete with the sweater. The idea is to make the chair more gray, the sweater more blue, and yet not destroy the overall neutrality of the image. The curve to the B channel shown below does all these things by placing points that define both chair and sweater.



The Magic of LAB

using the L channel from one and the AB from the other. The same result is available in RGB or CMYK, however, by layering one over the other. If the detail-rich version is on top, choose Luminosity mode; if the correct-for-color version is, choose Color mode.

Similarly: one often finds images with reasonable color that lack weight. Rather than apply some correction that might add the punch but fry the color, consider using a Layer or Adjustment Layer set to Luminosity mode. Then use curves that are aimed only at increasing contrast, just as you would do if you were working in the L channel of LAB directly.

Un Unlikely Blend Source

The blurry A and B channels seem like the last things that one would want to use for blending. But properly used, they can solve the thorny problem of backlit faces, among many others.

If a face is backlit, it will appear too dark, and likely not colorful enough as well. But it will still have a magenta (light) component in the A channel. By blending the A channel into the L in Overlay, Soft Light or Hard Light mode, the face can be lightened without a selection. Neutral areas won't be affected. Even the lightest highlights will survive, as long as they aren't reddish.

Similarly A/B overlay blends can be used to pep up color in a somewhat more intuitive way than by using AB curves. Make a duplicate layer and, on the top layer, apply the A and B channels to themselves in Overlay mode, intentionally making the picture too colorful by using high opacities for the blend. If desired, use different opacities

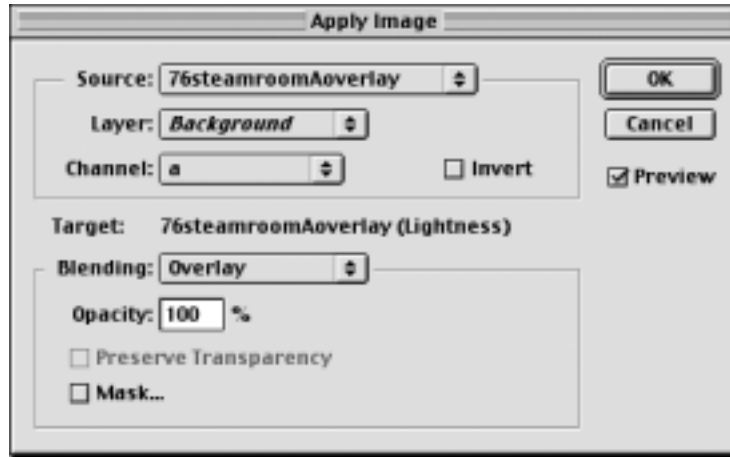


When the light source is above or behind the subject, many people resort to selections or masks in a frantic effort to remove the resulting shadow. No such thing was needed to generate the bottom version from the top one. In effect, the photograph was relit from the front, with the aid of an unusual form of LAB blending.

The Magic of LAB

notes

The men in the A channel are light because flesh is more magenta than it is green. The neutral background, however, is close to a pure gray. As pure grays are ignored in overlay blends, we can lighten the fleshtones in the image by applying the A channel to the L in Overlay (or Soft Light, or Hard Light) mode.



for the A and B blends. To make faces more golden-looking, for example, use a higher opacity for the B overlay.

Once the blends are complete to create the overly colorful image, it's easy to dial in whatever layer opacity is desired to produce livelier colors in the original.