

Nobody likes to look tasteless by presenting an image that's grotesquely too loud, any more than they would want to show something terribly oversharpened. Anybody can use LAB to create colors that are too intense, if not well outside the final output gamut.

The question is, would you prefer the opposite problem? Would you prefer to work with something not colorful enough, and add, or something too colorful, and back off?

The second problem, in my opinion, is a better one to have. This is particularly true in the context of the Picture Postcard Workflow, where files are often bland when they emerge from the luminosity step. But the principle is a good one no matter what workflow one uses.

The Picture Postcard Workflow does not worry about color intensity until near the end of the correction process. Its first color pass is merely to eliminate obviously incorrect color. The second step is contrast only. After these two steps the color is usually tepid.

The next step is always this Color Boost action, sometimes preceded by the Modern Man from Mars (MMM) action. A third action that combines the two, MMM + Color Boost, is also included in the download set, and this is the one I recommend to experienced PPW users. But this document will make matters easier by discussing only the Color Boost.

The action is a simple one and so is its objective.


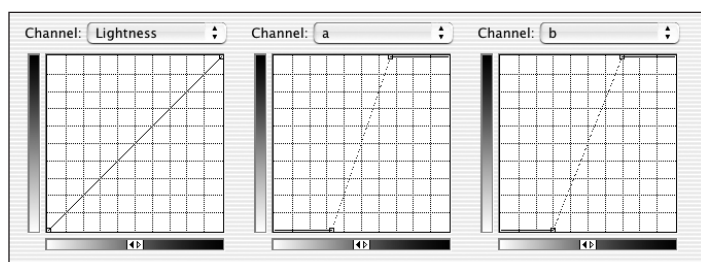


Figure 1. The Layers palette after the Color Boost action runs. Below, the curves of the Color Boost layer. The curves on the Endpoint Adjustment layer are defaults and do not change the image unless you choose to alter them.



It seeks to produce colors that are much too intense, leaving it to us to figure out the most graceful way to reduce them. Also, it gives a final opportunity to adjust contrast, if needed. Therefore, users don't have to have perfect highlight and shadow values before running this action, nor need they worry if the picture looks slightly too light or too dark.

The counterintuitive procedure of deliberately producing something too colorful is a logical response to the common request for happy, vivid colors—but not too much. In practice *too much* is very difficult to define. Parts of a picture can be too much and others not, or the whole thing can be too much.

Therefore, the Color Boost action deliberately provokes too much, so that we can decide whether the entire picture is too much, or only parts of it. Because anything that is not too much is good.


The PPW package contains some complicated actions, with fifty or even a hundred steps. Color Boost is not one of them. Its steps are as follows:

- Convert file to LAB if necessary.
- Flatten image if already in LAB, otherwise a prompt to flatten before converting.
- Add a curves adjustment layer, called Color Boost. Use the curves shown in Figure 1. Change layer opacity to 75%.
- Add a second curves adjustment layer, called Endpoint Adjustment. Close the layer without altering the default curves.

The Endpoint Adjustment layer is optional, self-explanatory, and will not be discussed further here. The 75% opacity on the Color Boost is enough to make all but the blandest pictures radioactive, but nothing stops you from raising the opacity, or even duplicating the layer, if you want more.

LAB aficionados will note that the AB curves do not use the same angles; the A curve is steeper. This contradicts the recommendation in my book *Photoshop LAB Color*, however, experience has shown an improvement this way. If you feel differently, as we'll see later, you can override this decision case by case, or as a permanent change in behavior.

Although it might seem that one could go through these steps manually almost as fast as by running



Figures 2 and 3. The original and a default application of the Color Boost action.

an action, in practice the action is needed. These AB curves are so steep that any imperfection would result in a visible cast. They must pass exactly through the center point. Photoshop's curve dialog isn't big enough to see whether this is the case if you do it on the fly. The curves in the action, however, have been tested for accuracy.

Before running the action you need to be sure that there is no color imbalance, otherwise it will be made much worse. This is why the first step of the PPW is to eliminate such minor problems, before they become major. All three of the images shown in this document were measured for color issues before applying the action and corrected if need be. Similarly, all these originals have adequate contrast, so there is no need to alter the Endpoint Adjustment layer.

The PPW suggests that contrast should be enhanced before entering the Color Boost phase. There is no known danger in ignoring this advice.

Once the action has been applied and the picture is too colorful, the straightforward move is to reduce the opacity of the Color Boost layer to taste. This is in fact the best move about a third of the time. Somewhat more frequently a better result is obtained by loading

the L channel, which resembles a grayscale version of the image, as a layer mask. This works because strong color in dark areas is often more obnoxious than in lighter regions. Less frequently, one uses a different channel, such as the A or B of LAB, as the layer mask, or edits a layer mask after loading it. For example, you could load the L as the mask, and then adjust the mask's own opacity.

Figure 2 is the original; Figure 3 immediately after running Color Boost. I have some clients who would think this beautiful, but in my opinion it is too much. The question now becomes how to back off. I offer three of the thousands of alternatives. I've tried to adjust opacities so that none of the three alternatives seem grossly more "colorful" than the others.

Anyway, Figure 3, the default, uses no mask on the Color Boost layer, which is set to 75% opacity. The simplest way to tone things down is to reduce that opacity, that is, moving everything back toward Figure 2. In Figure 4, the opacity has come down to 37%, meaning that it is in effect a 50–50 split between Figures 2 and 3.

Figure 5 uses the method I think is best for most pictures. The L channel is loaded as a layer mask. This reduces the impact of the color boost in darker areas



more than in lighter ones. A higher opacity is needed because of the dampening effect of the mask. It's 66% here.


The impact, as you can see, is subtle—you may have to look twice to see that Figures 4 and 5 aren't identical. Nevertheless, there's a cleaner break in Figure 5 between the enhanced light greens of the background and the darker trees in the foreground that the mask is protecting.

As far as I'm concerned these two can be your only methods if you just want to adopt a quick workflow. For the sake of argument I show an esoteric option as well. Figure 6 also uses a layer mask. This one is based on an inverted copy of the B channel of Figure 3. The B channel is light in areas more yellow than blue, dark in those more blue than yellow. Inverting it as a mask blocks the color boost in yellower areas (the background) while permitting it in bluer ones (the water). So if that's the effect you're looking for here, it's a reasonable option, and yields something different from Figures 4 and 5.

An A or B channel is always very flat. When loaded as a mask, its range commonly has to be increased, and layer

Figures 4, 5, and 6. Top, opacity of the Color Boost layer is reduced to 37% from the default 75% of Figure 3. Bottom left, a copy of the L channel is loaded as a layer mask, and the opacity increased to 66%. Bottom right, an inverted copy of the B channel, to which Auto Tone has been applied, is used as the layer mask. Opacity is increased to 88%.





Figures 7 and 8. Left, the original. Right, the Color Boost action is applied and layer opacity increased to 100%.

opacity increased, too. Here, being too lazy to apply a proper curve, I hit the layer mask with Auto Tone to expand its range. Layer opacity was set to 88%.

Curiosity, Cats, and Masking

This first example was fairly typical in that the choice of how to proceed wasn't obvious. The case could be made that the darker trees should be suppressed, as in Figure 5, or that the yellower background should, as in Figure 4, but it isn't clearcut. I don't know how an audience would feel about it.

In the next example, I do know, because I've shown it to about a thousand people in seminars and asked how they felt.

The cat of Figure 7 is so bland that the Color Boost action needs a boost itself. Figure 8 is not the default, but rather the action with the Color Boost layer increased to the maximum 100% opacity. Now, I trust, we can agree that it's too much.

But you can't object to the white fur. These LAB curves don't alter neutrals, so that fur hasn't changed from the original. In fact, only four things have.

1) The background has become more blue in Figure 8. The audience's verdict is unanimous: either they like it, or they don't care. Nobody wants to cut it back.

2) The darker fur isn't quite neutral in the original. Therefore it becomes somewhat more noticeably brown in Figure 8. The audience feels the same way as about the background. Generally they like the added color. Those who don't aren't offended by it, and would leave it alone.



3) The eyes have become a more vivid green. This is the only area of non-unanimity. Certain extremists vote for either Figure 7 or 8 just as they stand. The majority, however, feels that the eyes should be somewhere in the middle. From what I can gather from the responses the sentiment seems to be for something closer to Figure 8, but not quite that strong.

4) Back to unanimity. The nose is a shocking pink in Figure 8. Everybody strongly objects to it.

In short, anyone who says that this whole picture is too much is mistaken. The nose, which makes the cat look like it's gotten drunk, is too much. The rest of the picture is fine.

* * *

We therefore should find a mask that discriminates against pinks. Before showing the two likeliest candidates, let's clear the two usual suspects. Figure 9 is the straightforward opacity reduction. Figure 10 loads the L as a layer mask, with the layer set to a higher opacity. The difference between the two is more apparent



Figures 9–12. Top left, a straight reduction in opacity without a mask. Top right, the L channel is used as a layer mask on the Color Boost layer. Bottom left, the green channel from an RGB copy is used instead. Bottom right, a contrast-enhanced inverted copy of the A channel is used to try to suppress the pinkness of the nose.

than between Figures 4 and 5. The dark fur in Figure 10 is being suppressed by the mask, so it is not as pronounced a brown as in Figure 9. I'm not saying that either is better; that would depend on your agenda.

If the objective is to tone down the nose but not the eyes, two channels come to mind. If a copy is taken to RGB then the green channel should work, because it is light in green objects and dark in magenta ones.

Or, we could use an inverted copy of the A of LAB. That channel is, when uninverted, light in magentas and dark in greens. The major difference between it and the green channel is that it takes no account of darkness. Therefore, it is likely to permit the dark fur to get more colorful than a green channel mask would.

Figure 11, therefore, is the green mask version. It is similar to Figure 10, the L mask, but the nose is distinctly less obtrusive.

As we mentioned in the discussion of Figure 6, the AB channels are very flat. To be an effective mask, their contrast must be enhanced by a curve, Auto Tone, or some other method. Doing this here produced Figure 12. This version seems to cater to most objections. The darker fur and the background have been permitted to gain saturation. The nose is much less pink than in Figures 9 and 10. And the greenness of the eyes has been maintained from Figure 8.

For these reasons, if forced to pick one of the four, unaltered, as the best version, I'd vote for Figure 12, with Figure 11 in second place. But what if you aren't so limited, if you can go further in refining the color? For example, suppose that you feel that the eyes of Figure 12 are too intense but you approve of the rest of the picture.

In that case, you would apply a further curve to the layer mask. Where the mask gets darker, the color is lessened. So, with the layer mask active, you would open the curves dialog and Command-click a suitably green area of the eyes. That places a point on the curve indicating how dark the eyes are in the mask. Raise that point to darken it, and you're home.

To illustrate how to edit a layer mask we turn to a final example.

Figures 13, 14, and 15. Top, the original. Middle, a default application of the Color Boost action. Bottom, the L channel is loaded as a layer mask on the Color Boost layer.



A Curve in the Shape of a U

Figure 13 is the alpine original. Figure 14 applies the default Color Boost action. Do you see particular problem areas? Or would you say that everything is equally bad? In the latter case, we reduce opacity, in the former, we use a layer mask.

According to me, the greenery and the jacket are only slightly too much. The sky is a little more colorful than I would like, and the reddish areas much more colorful.

None of the methods we have seen so far are ideal. Reducing the opacity works but we will have to tone down the green more than we'd like. Applying the L as a layer mask favors the lighter areas, which are already the problem. Inverting the L as a layer mask would result in lurid colors that are currently dark. Nor are the

A and B good options. Brown is a species of red. Red is positive in both A and B. Any move with the B would adversely effect the greenery, which is also B-positive. A move in the A would damage the fleshtone.

Instead, we load the L, but only as a starting point. Figure 15 shows the result. As expected, the darker areas are muted and the lighter ones too colorful.

The layer mask itself, a copy of the L channel, is shown in Figure 16. Where light, it allows strengthening of color. Where dark, it suppresses it.

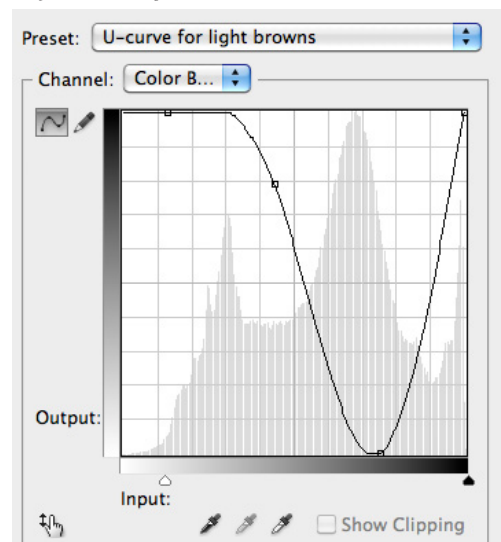
The objective is to transform this mask into something that blocks much of the left side, which is relatively light, while permitting color gains in the midtone and three-quarter tone. Fortunately, since the move involves color only, we can produce quite a radical mask and nobody will know—there will be no telltale arti-

facting. Sure, the new mask, Figure 17, looks absurd—but can you see anything strange about the contrast in the result, Figure 18?

The U-shaped curve that did this is shown below. The bottom of the U was found by Command-clicking on what I took to be the most important greens.

The original, Figure 13, and the unaltered L channel mask version, Figure 15, are repeated for convenience on the next page. I prefer Figure 18, but if you would like to modify it there are any number of ways, such as reducing its opacity,

Figures 16 and 17. Top, the original L channel, which was loaded as a layer mask in Figure 15. Bottom, the layer mask is modified by the U-shaped curve shown below.



or putting a gentler curve on the mask. The point of the whole exercise is to suggest first that you can do just about whatever you like with the color, and second that it is difficult to visualize what you would like to do without having something extremely colorful like Figure 14 to guide the way.

* * *

The Color Boost action works in tandem with other pieces of the workflow. Whenever we see skin, we think of applying the Skin Desaturation action, which was not done in this last exercise. It would have made life a bit easier because the action would have taken parts of the ground at left as being fleshtone values and toned them down. But the basic concept remains.

Also, a reminder that we generally do not run the Color Boost action alone, but a more complicated one called MMM + Color Boost. The MMM part adds color variation and does not necessarily produce more brilliant colors, the way Color Boost alone does. The MMM action has its own separate PDF documentation, as does MMM + CB.

The joint action produces at least four different layers to play with, two for color and two for luminosity. After all, one can never have enough options!

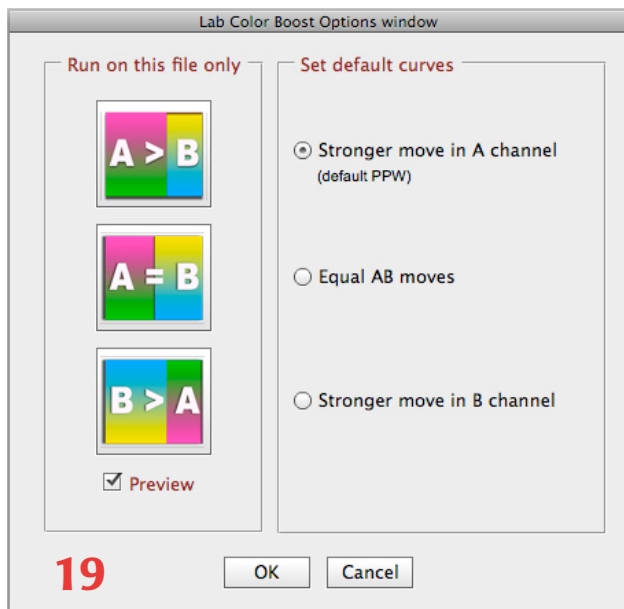
A New Angle on Curving

As mentioned earlier, the default Color Boost curves are steeper in the A channel than in the B. This means that the magenta and green components of the image intensify faster than yellows and blues.

The PPW panel permits you to override this decision, either for a

Figures 13 and 15, repeated for convenience, and 18. Top, the original. Middle the original L shown in Figure 16 is used as a layer mask on the Color Boost layer. Bottom, the modified layer mask of Figure 17 is substituted.





Figures 19, 20, and 21. Top, the Color Boost action options window. Middle, an original photo. Bottom, the original goes through the entire PPW, including the MMM action, but not the Color Boost.



single image or as a permanent preference. Option-clicking CB in the panel brings up the dialog shown in Figure 19. Its left side refers to this image only; the right side changes future behavior.

Why is the default the way it is, instead of the more logical way of equal treatment of A and B? Because years of experience have shown that normally this is the correct thing to do.

Why this should be true is less important than the fact that it *is* true. A surmise: we live in a yellowish world. Average every pixel in a huge number of images (and I have) and you'll get a value very close to 0a—no bias toward either magenta or green. The B value, however, is slightly positive, around 5b, meaning that the “average” picture is yellowish. I speculate that further exaggerating the natural yellowness may be irritating.

When would this approach be undesirable? I find it so when the image features one or more persons who appear rather pink already; this is particularly true of individuals with blond or similar hair. Also, if the image showcases yellowish browns, the defaults may turn them overly orange.

The impact isn't huge, but it can be noticeable. Figure 20 shows restored art from Huaca de la Luna, constructed around 500 AD by the Moche people of what is now Peru. Their ceramic work was of better quality than could be found anywhere else in the world of that time, and many of the colors have held. In the original photographed, the dim lighting suppresses them. How far to go in our own restoration?

Figure 21 shows the progression of the PPW through and including the MMM action. The next step is Color Boost. Figure 22 shows the default angles of A>B, with layer opacity reduced arbitrarily to 40%.

Figure 23 uses the options dialog

to substitute $B > A$, also at 40% opacity. The action doesn't change neutral areas, so the near-whites can't be distinguished from one version to the other. The reds and yellows, however, are clearly different. You get to make the call on which is better. A reminder that you can also choose $A=B$, which isn't shown.

In the accompanying Layers palette, note the gentle reminder in the Color Boost layer of the way chosen to enhance color. Refer back to Figure 1, which was done at the action's default, and you'll see an $a > b$ in the Color Boost layer where this one shows $b > a$. This handy information is displayed automatically.

If you use the MMM + CB action through the PPW panel, you'll find a host of other intriguing options.

Figures 22 and 23. Top, the Color Boost default angles, 40% opacity. Bottom, $b > a$ instead. The panel automatically inserts the proper indication in the Layers palette (below).

