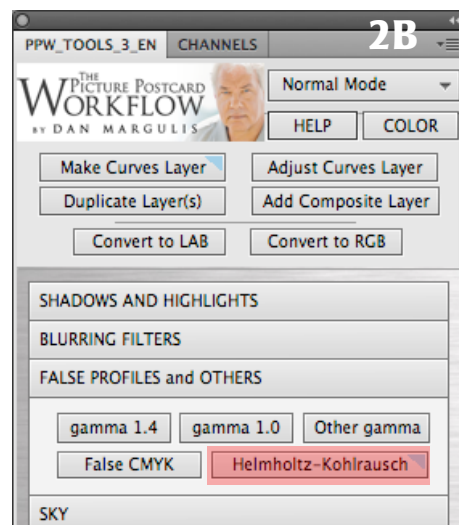
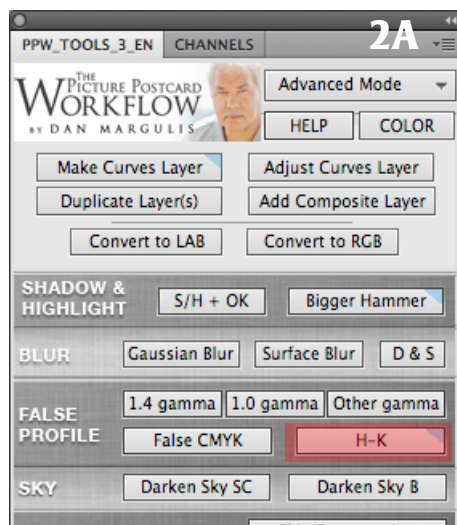
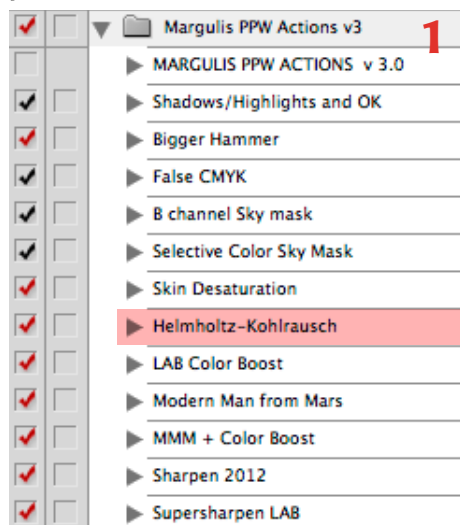


The main objective of this essay is to document the Helmholtz-Kohlrausch action, henceforth H-K, contained in the Margulis PPW Actions set. The action's goal is to darken neutral and near-neutral areas in the quartertone and midtone. Doing this often emphasizes more strongly colored areas in a pleasing way. A second, independently adjustable layer desaturates most colors, but near-neutrals more than others. This can pave the way for needed color variation down the line, or be worse than worthless, all depending on the character of the image.

The action gains no new capability in PPW panel v3, although it now employs a layer group to toggle the entire effect. The name derives from the Helmholtz-Kohlrausch effect, which states that artificial measurement devices such as cameras become confused when they evaluate the darkness of a colored area as opposed to one that is only slightly lighter, but neutral. The artificial device is apt to declare that they are both of the same darkness. In fact, they aren't – at least not for our visual system: the more saturated a color is, the brighter it will appear to our eyes. This essay will discuss when it is useful to employ the Helmholtz-Kohlrausch action, and attempt to give a set of directions and rules.

In order to use the action you must, of course, have the set loaded. In your Actions Palette, it should look like Figure 1. Alternatively, if you have the PPW

Figures 1, 2A and 2B. The Actions Palette and the PPW TOOLS panel.




TOOLS panel installed, you will find a button connected to the action in the FALSE PROFILE section (Figure 2A and 2B).

The action can be invoked in any colorspace, but it requires a flattened RGB file to run. If you're using the panel and you're not already in RGB, you have the option to continue working on the converted original file or on a duplicate.

The action creates an alpha channel called Ersatz Black. As the name suggests, it resembles the black channel of CMYK, in that it is rather light wherever strong color appears. This channel is treated in a different way on each of the action's two layers, which are meant to control luminosity and color separately. They are named "H-K Effect" and "Color only" respectively. The split is valuable because the user can choose to darken neutral areas without changing their color, or neutralize them further without changing their darkness, or both.

The H-K Effect layer, which is in Luminosity mode, is masked to prevent shadows from plugging. The Color only layer, which is in Color mode, does not use a mask. However, the neutralization process works with a copy of the Ersatz Black channel that has been modi-

Figure 3. The Layers palette after the action has run.





Figures 4 and 5. The original image and its desaturated version. Notice how darker the trees appear in the second.

fied to affect mostly the quartertone-midtone range.

After the action has run, your Layers Palette will look like Figure 3. The temporary alpha channel called Ersatz Black is deleted by default, so the Channels palette will remain unchanged. If you are working with the PPW panel, however, you can option-click H-K to bring up a dialog that allows retention of Ersatz Black, either for this file only or as a permanent change in behavior. Some of the many other uses of such a channel are described in the second half of Chapter 13 of *Modern Photoshop Color Workflow*.

The treated image will therefore look slightly darker and desaturated in the areas where the Ersatz Black is at work, depending on how we choose to blend the “H-K Effect” and the “Color only” layers. In order to understand why we should want to dampen colors this way, we must examine the nature of the H-K effect.

Something about the H-K Effect

I have a question about Figure 4: which is lighter – the trees or the mountains? You don’t have Photoshop at hand, so no Info Palette is available. Only your eyes can tell.

I proposed this test to five non-experienced people and their reply was unanimous: the trees are lighter. One of them objected that luminosity is not constant in either area, which is true; but when asked to judge the average, she chose the trees.

Yet this is the wrong reply, and one may easily prove it by exploring the image with the eyedropper. While both color and brightness vary from point to point, the average Lab value is 25L(7)a14b in the trees and the (averaged) mountains measure 46L0a(12)b.

Let’s forget the hue, for a moment. Could you tell that the average brightness in the mountains is roughly twice the brightness in the trees? I couldn’t. I think nobody could: unless they looked at the de-

saturated version of the image – Figure 5. Notice how darker the trees look in this B/W version obtained via Image -> Adjustments -> Desaturate. If you tried Image -> Mode -> Grayscale the result would be quite different, because the Green channel, the lightest in the trees, would contribute a lot to the final result.

Bottom line: should we believe our eyes or the Info Palette? My flippant reply is: both. From a numerical point of view, the Info Palette is right; from a perceptual point of view, we care about what our eyes have to say, and our eyes say that the trees are brighter.

Like it or not, we should take this into account.

Where Saturation Lies

You may have noticed that I disregarded the hue and discussed brightness (L, to us) – but didn’t mention saturation: I didn’t, because that’s where the key lies. The saturation in the trees is higher than in the mountains – both locally and globally, and you’re staring at the so-called Helmholtz-Kohlrausch effect in the face. The Wikipedia entry for this German-sounding phenomenon is laconic to say the least: “The Helmholtz-Kohlrausch effect is an entoptic phenomenon wherein the intense saturation of spectral hue is perceived as part of the color’s luminance.”

In simpler words: if you have two objects which are equally bright, but one is more saturated than the other, saturation will be perceived as a luminance component and the former will appear brighter: to your eyes, of course, not to a colorimeter. Entoptic is a term used to describe a visual effect whose cause lies within the eye: this is, therefore, a perceptual phenomenon.

It’s also wildly non-linear: the increase in the perceived brightness depends not only on the saturation of the color but on its hue as well. In practice, blue and red areas tend to appear brighter than yellow and

green equiluminous ones.

The real problem of the H-K effect is that it remains largely unexplained although it was discovered in the 19th century. If you're interested, google "The Brightness of Colour" and read the article written in 2009 by Corney et al on www.plosone.org. The idea is that "natural visual systems evolved to encode the past empirical significance of stimuli". That is, blame it on our evolution.

What we are most interested in, though, is a by-product of the H-K effect: while it is true that we tend to perceive more saturated colors as brighter, the opposite holds as well – we are less sensitive to saturation in darker areas. Your camera sensor doesn't subscribe to this point of view, neither does Photoshop: it is well known that a slight bluish cast in a shadow may be acceptable and even pleasant, but if one needs to perform heavy color boosts such shadow may turn out deep blue – and this is not in agreement with our vision.

It therefore seems sensible to preventively desaturate the darkest and most desaturated areas of an image when one needs to boost colors, in order to



tame the undesirable effect described above.

Paint It, Black

At least one concept is easily summarized: if a picture is too colorful, then each color will fight with every other. That's why we need some chromatic contrast in order to make the really bright colors stand out, and in the light of the H-K effect this suggests that we should put those colors which do not need to be excessively boosted on a leash. This basically means we should desaturate a bit whatever is already relatively dark and dull.

The previous sentence implies that we are dealing with a picture where both vivid and less vivid colors are present. Moreover, we're assuming that the most vivid objects in the picture are those we want to emphasize – which is not necessarily always true. Therefore, if you have a photograph where either the most important objects are near-neutral or the shadows contain some absolutely critical detail, the usage of H-K is discouraged. While the action won't plug your shadows, you may find yourself in trouble if their dynamic range is compressed, as it happens using the

Figures 6 and 7. Left, a Bogus Black channel produced via the False CMYK Profile method. Right, the same, aptly curved.



action. Therefore, careful with that channel: H-K'ing a picture of a black cat inside a coal mine may not be the brightest idea on the table.

Back to our task – taming less saturated colors. There are several ways to accomplish this, but before explaining them, a small reminder about an aspect of CMYK is necessary.

One of the most powerful characteristics of CMYK is also the most bewildering for beginners and is connected to the dark matter (pun intended) of GCR.

While in a given RGB space and in Lab a color can only be represented by one triplet of numbers, in CMYK colors can often be represented in different ways. While there is just one way to write a color like, say, $0C80M70Y0K$, things change as soon as a third component kicks into action. Therefore, two formulas like $21C35M100Y9K$ and $16C32M99Y15K$ actually describe the same color, give or take rounding errors. This is called GCR, for Gray Component Replacement. In CMYK, a mixture of C, M and Y inks can form gray; K is neutral, and it therefore can substitute an amount of colored inks. Of course, the opposite is also true: C, M and Y can replace K if needed. The ramifications of this fact extend beyond the Solar System, but let's just say it can be done and it makes sense. Light GCR therefore identifies a CMYK separation where K is gentle and there is comparatively more colored ink than in, say, Medium GCR, where K is slightly more aggressive but there are less colored inks.

Every CMYK separation has a Black Ink Limit, which is the maximum density allowed in the black ink. A typical value is around 85% or even 90%, but in at least one particular case this rule is broken.

This case is called Bogus Black or, alternatively, False CMYK Profile. The origin of the name is so interesting that I've decided to put the story at the very end, in the Credits' section, but its meaning is just what the word means: it is a counterfeit black channel that we create, use and then mercilessly discard. I won't explain how you can make one, and I'll just say it can be done in Photoshop in the Custom CMYK window. But you won't need to do it on your own, because a False CMYK Profile is already available in the Margulis PPW Actions set. I'll show you how it looks, though. Figure 6 is a Bogus Black channel produced by converting an RGB image to False CMYK Profile via the False CMYK action in the PPW set. You should notice two things: first, even the darkest areas are not so dark, because K is kept within an unusually low black limit (around 60%) in this conversion. You may see this in the histogram of the Curves Palette inside

the image, as well.

Second, and the key point: whatever kind of GCR you pick when you convert to CMYK, except for the extreme conversion which doesn't involve GCR and has no black at all, black ink will invariably play the same game: it will appear in dark and desaturated areas. Look at the ribbons in the lower part of the picture: you do not know their color, but you can bet your head that they are either very light or very colored – or both things together. You can say so because black appears only in their shadows, not in well-lit parts. If they were darkish and desaturated, black would reign. This is what happens, for instance, in the darkest leaf just above the ribbons: you can surmise it is green, but it can't be a bright green, because the amount of black present prevents such color.

Supposing that our original image doesn't have out-of-gamut colors, the conversion to a False CMYK Profile won't hurt it too much. You can expect to see somewhat cheesier shadows due to the low black ink limit, but that should be it.

Figure 8. The Ersatz Black channel produced by the H-K action.



Now, the interesting move. We need to curve the black channel so that depth is restored in the shadows. Also, we don't want the areas where K is light to darken. So, what we need is a curve similar to the one in Figure 7, which is pasted inside the resulting image. Result: deeper shadows, more contrast in dark and desaturated areas, which become even darker and more desaturated.

At this point one can enhance the colors: the preventive manoeuvre performed with the Bogus Black will prevent excessive saturation in the darkest areas of the image, and will therefore make the bright and colorful ones stand out more – in accord with the H-K phenomenon. Two birds with one stone.

Mind you: the aim of this technique is not to tone down every color; vivid areas should remain as they are. To quote the title of this section, when the Rolling Stones sang 'I see a red door and I want it painted black / No colors anymore I want them to turn black' they were not just sporting an obvious juvenile spleen, but also making poor judgment on what saturation is about. We just need to make the dark darker and the desaturated more desaturated, and stop there.

The RGB Way

There is a different path available, though, which doesn't require a diversion into CMYK and remains inside RGB. The idea is the same: find anything that qualifies as dark and desaturated and hold it down, so it won't explode later. In order to do this we need some kind of black channel, or at least something that may be comparable to a black channel.

This is built by the H-K action, and it is, you guessed it, the Ersatz Black alpha channel, pictured in Figure 8. The channel is built through a certain manipulation which involves the three RGB channels, but the bottom line is that it resembles closely some sort of CMYK black channel. The Ersatz Black is applied to the original image in Multiply mode, and the resulting image – darker and desaturated where needed – is applied on top of the original through a mask which restricts the effect of multiplication in order to avoid

Figures 9 and 10. An original and a version processed with the H-K action.



that shadows may plug and, in general, an excessive darkening of the image.

You can compare a detail of the original picture (Figure 9) and a version processed with the H-K action (Figure 10). Concentrate on the darker and less saturated parts, like the leaves and the inside of the pomegranate in order to see the effect of Ersatz Black at work. It is a fascinating, albeit subtle result: Ersatz Black gives it darkness, darkness yields depth, and our eyes give it shape.

H-K in the Picture Postcard Workflow

Our main goal is, of course, to maximize and optimize the effect of the operations we perform, especially when we deal with color variation.

In the PPW you may decide at some point to follow the path of assigning a low-gamma RGB false profile to your image. This procedure is explained in the False Profile documentation and has the effect of lightening the image while maintaining the original numbers: only the appearance of the image is altered, not the values of its pixels.

The main reason why you may want to do this is twofold: either you're about to perform a straightforward multiplication in either RGB or Lab and you need a lighter image to start with; or you want to use some kind of Bogus Black in order to maximize color variation in images which do not need a huge color boost. I'll concentrate on the second case, here.

Figures 11-13. Descriptions are in the text.



Figure 11 is the original picture; Figure 12 is the same picture with a 1.5 gamma profile assigned (the original is standard sRGB, gamma being therefore 2.2), and then multiplied in RGB with the standard technique; Figure 13: as Figure 12, but the H-K action was applied immediately after assigning the low gamma profile – multiplication following.

To me, Figure 12 is better than the original due to the enhanced local contrast brought in by the multiplication: I see more detail and structure in some critical areas (see the leftmost pomegranate, for instance). If you think it's a bit too light with respect to the original, a quick curve will do. Yet the version treated with the H-K action is even more interesting: more color variation is present due to the fact that some areas were desaturated and darkened before multiplication took place. This is particularly evident in the yellow and green leaves at the extreme right.

One of the key questions is how does the H-K method compare to the more standard Bogus Black procedure? The cleanest possible answer at this stage is: it depends on the image, without further specification.

Figures 14 and 15 attempt to compare H-K and Bogus Black, respectively. They both have strengths and weaknesses. The Bogus Black version (Figure 15) is lighter, as expected, because the black channel is curved, not applied to the image. The H-K version (Figure 14, identical to Figure 12) has more color variation where it's more needed (for instance, the pomegranate at the center has more structure, whereas I'm happy not to have too much going on in the darkest leaves, which is exactly what I would like to see; the ribbons have more shape, too.)

On critical images it may be worthwhile to try both approaches, but this can be slightly time-consuming. My personal feeling is that H-K be more likely to work on very "busy" images with a lot of subtle detail. The problem is – how?

Ersatz Black Serves MMM

It would be wrong to think that the only use of H-K is in connection with the RGB false profile move sometimes involved in the PPW. There are times when we simply wish to induce more color variation, but we don't need to actually boost the color which would probably become unbearable if we did.

A typical example is shown in the crop of a garden image presented here in four versions. Figure 16 is the original. Figures 17, 18 and 19 have one thing in common and a difference. The common thing: they have

all been treated with the MMM (Modern Man from Mars) technique, which again can be performed via the homonymous Action in the set.

The principle behind MMM is complex, but can be summarized as follows.

A certain selection is made in the image: it can be anything, from a small area to the whole image and even a multiple selection. Whatever it is, it should contain a range of colors we wish to enhance. Enhancement, here, actually means variation, not strong saturation. The image is then converted to Lab and equalized: this means that the levels in each channel are re-arranged in order to maximize the contrast in the values contained in the selected area. All contrast, of course: which means - luminosity (channel L) as well as color (channels a and b).

The resulting version is unusable per se, but when it is blended at low opacities, it can work wonders. An ancestor of MMM was the Man from Mars move (not "Modern", that's where the third M comes from) which would produce similar results exploiting extremely steep curves in Lab, but the new version is perceptually more pleasant and, in my experience, better tuned. It is also a lot more complex to implement, and that's why it will remain "just an action" to us. The selection, as we will see, is critical. This would be a long discussion, therefore I suggest you refer to the MMM documentation for details.

Back to our image and the differences between Figures 17, 180 and 19. Being dominated by green, this image calls for a serious variation in such color and MMM is almost mandatory in these cases. Now, in order: Figure 18 is a straightforward MMM based on a square selection containing leaves only; Figure 18 was processed with the H-K action first, and then treated ex-

Figures 14 and 15. Comparison between H-K (top) and False CMYK Profile (bottom). Both methods have strengths and shortcomings.





Figures 16-19. Figure 16 is the original. Figures 17-19 are treated with MMM as described in the text, with or without H-K (18 & 19 vs 17) and with different selections (18 vs 19).

actly as Figure 17 (same selection); Figure 19 is like Figure 18 but based on a different selection, i.e. the intersection of the original selection with the inverted selection based on the Ersatz Black channel. If the latter is difficult to grasp, imagine this: select a square containing leaves only. The Ersatz Black channel will be what it is in such area, darker where the color is darker and lighter elsewhere. If you invert it, it will become light in less saturated areas, and if you turn this into a selection you will actually select the less saturated parts of the image. Now, by intersecting this with the aforementioned square you're simply limiting the selection to what is inside the square, which was fully selected by definition. In simpler words, the selection used in Figure 19 favours the less saturated objects in the originally selected area, and basically

means "the selection of the leaves inside the square, less the most saturated and lighter ones".

The idea of making this somewhat difficult selection came from a suggestion by Dan Margulis. This has a serious effect on the equalization adjustment: the new selection contains fewer luminosity levels and colors than the original. Best case, it might contain the same amount of information, but not more. Since the equalization is based on the content of the selection and then applied to the whole image, the smaller the palette we are working on, the more pronounced the result will be. In other words: select fewer colors, and the equalization will perform a stronger stretch on the channels. The effect may be striking, especially in Lab.

The Verdict

The one thing we can safely state is that the original image is the flattest among the four. To determine which one of the three versions is the best is a horse of a different, er, color.

There is certainly more color variation in Figure 17 than the original, and this is expected because that's exactly what MMM attempts to do. It is even more apparent in Figure 18, because the dark areas are more desaturated. In Figure 19, things change: the color variation is surprisingly less pronounced, but the H-K manoeuvre is visible; the dark areas are better defined than in Figure 17. It is difficult to declare a winner, though.

When in doubt, I summon up a jury and let them judge: eight people in this case, four men and four women, with varying degrees of experience in the field of image evaluation. I asked three questions:

1. Which version among 17, 18 and 19 has a more natural color variation in the leaves?
2. Which version among 17, 18 and 19 has more shape and detail in the leaves?
3. Which version among 17, 18 and 19 do you consider the worst?

I explicitly instructed them not to take the original into account--it was to be a reference version only, to check the others against a common starting point.

Eight people may not be much of a crowd, but they're more than enough to at least show a tendency. The results were interesting.

- Question 1: Figure 17 got 3 votes out of 8, and Figure 18 got 5 votes. Figure 19 got no votes.
- Question 2: Figure 18 got 7 votes out of 8, and 1 vote went to Figure 19. Figure 17 got no votes.
- Question 3: Figure 17 and Figure 19 got 4 votes out of 8 each. Figure 18 got no votes.

This means that the version previously treated with H-K (Figure 18) looks more realistic and natural to the striking majority of the observers: this is sensible, because deeper shadows add depth and therefore define shape. It also means that nobody would con-


sider Figure 18 a bad version. The surprise was that I thought Figure 19 should be in the place of Figure 18, but something went wrong – at least with this picture. By reducing the original selection through the intersection with the Ersatz Black channel, I would expect a more pronounced color variation. Well, it didn't happen: the green is indeed greener, but perceptually flatter than Figure 19. It is also lighter, and that's correct (can you tell why?). The real problem is that the result is strongly dependent on the selection we make. In this particular crop, the green is about everything you can select, but in the whole picture you may have a wider choice.

Select, Select, Select

The dependency of the result on the selection is not easy to grasp without a visual example. Enter therefore Figures 20-22: Figure 20 is a simple gradient going from red to white, in Lab; Figure 21 is the result of Image -> Adjustment -> Equalize performed on Figure 20 with the outlined area selected and the option "Equalize entire image based on selected area" active – which is what MMM basically does; same holds for Figure 22. The geometry of the selection is not important: the only important thing is what the selection contains in terms of luminosity and color, something I will loosely call a "palette" from now on.

If the palette contains a large variation (imagine the rectangle of Figure 21 superimposed over Figure 20), the result is a contrast enhancement which forces the originally selected values to stretch as much as possible. Anything else is clipped. The reason why you're not seeing black in Figure 21 is that the color at the left is Photoshop's attempt to represent an impossible Lab color: 0 in the L channel but strongly positive in a and b. Figure 22 is exactly the same, only the original selection is a lot smaller and therefore defines a smaller palette, with less variation: as a result, the maxed-out area becomes much bigger. The bottom line is that we have enormous variation inside the original palette, but outside everything is dead flat. If the structure of the picture is such that the relevant objects fall outside the palette defined by the selec-

Figures 20-22. Left to right: a gradient in Lab, the same equalized through a selection containing a wide "palette", the same – but through a smaller selection. The emphasis is on the size of maxed-out areas.





Figures 23-28. Figure 23: the near-monochromatic result out of Adobe Camera Raw. Figure 24: iterative application of the R channel in Luminosity mode yields some contrast. Figure 25: straightforward MMM. Figure 26: H-K followed by MMM, same selection as Figure 25. Figure 27: H-K followed by MMM, with the complex selection described in the text. Figure 28: 100% Color blend of Figure 27 on top of Figure 26, and the final result.

tion, enhancement won't happen in them. A change, for sure; variation, not.

A counterexample follows. I've included the original version from the .dng file (Figure 23) in order to show that this is not a picture – rather an exercise in mimicry. Interestingly enough, the photographer claims that the scene he saw was very similar to this: the insect, apparently a *Tettigonia viridissima*, was so similar to the leaf that he failed to see it in the first place.



Figure 24 is a very standard correction: blend of the R channel into RGB in Luminosity mode, Flatten Image, blend again R as before – and so on, five times. Luminosity was then corrected in Lab, but color remained untouched.

Figure 25 is a straightforward MMM with a rough lasso selection of the insect. Big improvement in color variation!

Figure 26: H-K and then MMM exactly as in Figure 27. Much better luminosity and depth, to my eyes.

Figure 27: exactly as Figure 26, but MMM was performed with the complex selection described above and based on the intersection of the inverted Ersatz Black channel with the original selection. Better color than any other version in terms of variation, but not as good as Figure 26 as far as luminosity is concerned.

Figure 28: a blend of Figure 27 applied onto Figure 26 in Color Mode, 100%. It therefore exploits the color of Figure 27 and the luminosity of Figure 26.



Figures 29 and 30. Left, a grayscale representation of the original selection yielding Figure 27. Right, surprise: the leaves are not selected as one would expect.


The latter wouldn't have been a clever move in the previous case: the leaves had less variation in Figure 19 than 18 and a similar blend would flatten the whole picture. In this case the selection suggested by Dan Margulis works a lot better. Why? The subjects are apparently very similar: leaves in one example, an insect almost molten into a leaf in the second. Moreover, in both cases everything is green. Yet the result couldn't be more different.

The key question is – what does the selection prepare for the Equalize command to chew on?

We need to go all the way, at this point, and find out. Figure 29 is the grayscale representation of the selection yielding Figure 27. The black area means that nothing outside the lassoed part is selected. Inside, the lighter the image, the more that area is selected. Bingo: there is some kind of selection in the whole subject, and especially in the dark green areas; but not only in those areas. Compare it to Figure 30, though: that's the grayscale representation of the selection yielding Figure 19, the loser in our contest. The black areas you see are green bright leaves: which means the brighter green objects are not selected in this case and (see Figure 22) the whole enhancement of the equalization actually flattens them, because they fall outside the palette we're working on.

You may wonder, at this point, whether all this is

Figure 32. The proof that Luminosity may not work as well as the inverted Ersatz Black as a selection.



worth the effort. Couldn't we simply forget about Ersatz Black and load Luminosity instead, in order to try and force some more color variation in our image via MMM? In the leaves' picture, for sure. Or, even better, we may not load anything: whatever we load might as well reduce the variation in something already flat, color-wise, and it's difficult to tell in advance what may happen. In the insect's one, the opposite holds. See Figure 31: that's the equivalent of Figure 29, only the original selection was intersected with RGB luminosity. Do you notice how less selected the subject is? In a picture like this, we definitely need a selection like the one in Figure 29 to squeeze the last breath of color from flatness. This is not pure speculation, because it will yield a useful rule of thumb suggesting when to do or not do this operation.

The key point is: H-K seems to be a valuable option before MMM, and the Ersatz Black Channel may turn in useful to modify the selection used in MMM. This is slightly awkward, I realize, so...

Bogus Said – Play It Again, Ersatz

Let's go back to Figure 22 and try to grasp the whole process in one go. The first correction (Figure 24) is certainly a key passage, because the R channel of the original development contains enough information to actually allow us to show the insect, and it induces a mild color variation. From then on it's just color boosting via MMM in different flavours, and in this case Dan Margulis' selection works well because it contains reasonable variation (see Figure 29): there already is some color variation in the selected parts. In the case of the leaves we had a surface made of small green objects, instead, whose most interesting parts would fall out of the selection (see Figure 30). Color variation would not happen because we would blend something which would not vary too much, basically it was all uniform green – and the keyword here is "all": we needed more contrast, and



Figure 32 and 33. Top, a straightforward RGB-to-Grayscale conversion. Bottom, the same after a previous run of the H-K action.

the method yielding more contrast is H-K followed by straightforward MMM. In the insect's case, it's the opposite: the image has such color and luminosity that the selection picks enough variation for the channels to become seriously stretched, but not enough to become almost flat in the areas of interest (see Figure 22). So, we want variation, but it must have a degree of smoothness – and this is why in Figure 19 color was boosted (everything is greener) but variation didn't occur: the leaves would lie in the “flattened area” of our palette. Finally, I mentioned that Figure 27 has better color, but Figure 26 is superior in luminosity: join the two as described, enter Figure 28, and you're done.

One remark: as long as the leaf doesn't turn icy blue or the insect starts to exhibit pink polka dots, the color is fine for me. This case is extreme enough and almost anything goes, also because we can't really bet on any color here, except for the fact that the green in the leaves should fall inside a known yet very wide range. But as long as the viridissima rears its head, I'm happy – and I guess it does.

Ersatz Black in Grayscale Conversions

An open field of study is connected to the conversion of RGB images to Grayscale. As you certainly know, there are several ways to accomplish this task, Image -> Mode -> Grayscale being probably the worst. Yet suppose you need to do exactly this, and quickly. Does it make sense to run H-K before you cross the bridge? It often does, for the very same reason why the idea may work in color. A grayscale picture doesn't have color, but you're trying to convince the viewer that some color exists, and the only weapon you have to do that is luminosity.

The idea of darkening the darker and less vivid

parts still holds and translates well into Grayscale. Compare the two images in Figures 32 and 33, for instance. Figure 32 is a straightforward grayscale conversion of the original. Figure 33, the same - after H-K. I would choose the second: it shows more variation and shape exactly as the corresponding color version, and also for the very same reason and mechanism. No colors are

involved, except for those in our imagination – and as you know there are quite a few in there.

As a general remark, this is just an example, and this technique may come in handy even when you're not using a straightforward conversion to Grayscale. Feel free to experiment.

Conclusions

Human vision is sometimes weird, and certainly weird enough for some phenomena to remain unexplained for over a century. The Helmholtz-Kohlrausch effect is among these: yet it is clear that we tend to consider saturation as a non-linear, additional component of perceived brightness. Photoshop doesn't: therefore, clamping the darkest areas a bit and desaturating darkish colors makes sense, because this actually rearranges the dynamic range in such a way that a subsequent saturation won't probably make them too colorful. If it did, this would contradict our sensory experience.

The two weapons at our disposal are a conversion to a False CMYK Profile and H-K. The two methods yield different results, and it's awkward to say which is better, generally speaking. In other words, they're very image-dependent.

Yet the Ersatz Black Channel is just an imitation of some real Black Channel, and we know that K has several uses. In a very rough draft of his notes on H-K, Dan Margulis would list the usual possible uses of a black channel, which are worth repeating:

- *As a source of multiply blends into individual channels (usually on a luminosity layer), with the idea of increasing contrast.*
- *To multiply into an entire RGB document. This emphasizes the purer colors because no black would be found there. Anything else would get darker*

and more neutral.

- *To force neutrality into all but the brightest colors, without adding weight to them. The procedure is: on a duplicate layer of an RGB document, apply the artificial black channel at 100% opacity, Multiply mode. Then change the mode of the entire layer to Color and adjust opacity to taste. (Note: if modest darkening is also desired, duplicate the Color layer, set its mode to Luminosity, and again adjust opacity to taste.)*
- *As an outright replacement for part or all of an RGB channel when the main focus of attention is an object that would be described as black. In such cases CMYK gains an advantage over RGB because the black object migrates to the black channel and the CMY channels don't show much detail. This makes concentrating on improving the black object much easier.*
- *An artificial maximum-GCR black makes an excellent mask for cases where we wish to emphasize colored areas, but not so much more neutral ones or darker ones.*
- *An artificial heavy-GCR black channel, inverted, makes an excellent sharpening mask. We usually want more sharpening as the image gets darker, but we prefer to minimize sharpening of strongly colored objects.*
- *This same inverted, blurred mask can give outstanding results when we are attempting to lighten and create contrast in shadow areas without damaging colored regions of the image.*

In reference to what was discussed in this article, we can add the following suggestions, given that in general H-K is more aggressive than the False CMYK Profile. This may lead to some rough guidelines, which should not be taken as Gospel anyway.

- *If you're dealing with an image which has a lot of fine detail, and such detail is essentially built on luminosity variations (e.g. the leaves' picture), H-K probably works better than Bogus Black. Otherwise, if the detail relies on color more than luminosity, my best bet would be Bogus Black.*
- *If contrast is built on color as well (e.g. the pomegranates picture) the two methods will produce different results which have both strengths and weaknesses. If your image isn't critical in the shadows, H-K is probably a better choice. If you need a light, bright image instead or are afraid that shadows may either plug or lose detail, False CMYK Profile is likely a safer way to go.*

- *If the critical areas where you're trying to enhance color are rather brilliant, MMM is the way to go, and previous H-K may bring more shape to the subject. The selection necessary to run MMM should of course encompass the areas of interest.*
- *If the critical areas are rather dull instead, it may be worth trying MMM with a selection which, again, encompasses the areas of interest but is intersected with a selection based on the inverse of the Ersatz Black channel generated by H-K.*
- *Care should be taken when converting to False CMYK Profile if the image has strong colors or, on the opposite, very light, washed out tones which may suffer from going to CMYK. Color is not a problem in itself, because it can be restored easily; luminosity is, though, because clipping may wipe out essential detail. H-K is a much safer choice.*
- *H-K may re-arrange the dynamic range of images in such a way that a subsequent conversion to grayscale may look more convincing.*

Appendix

Creating The Ersatz Black Channel

I've previously mentioned that the Ersatz Black alpha channel is built through a certain manipulation which involves the three RGB channels. It is a simple manipulation indeed: take a duplicate of the R channel, apply the G onto it in Screen blend mode, and apply the B onto the resulting channel, again in Screen blend mode. You've already seen (figures 7 and 8) that the result closely resembles a Bogus Black channel. The bottom line we need to remember is this: both such channels, when multiplied with the image which generated them, will darken and desaturate the darker and more desaturated areas. The reason why we need this is discussed in the article.

New features

The most relevant new feature in the new release of the H-K action is that it splits luminosity and color. This gives us one more handle to control the process: basically, it means that we can control how much we want to darken and how much we want to desaturate the original image, independently.

Consider figure 34. Take this as an original, although it is actually a picture which already arrived at the Luminosity step in the PPW. The colors are fine, but we may need some more saturation and variation, although I wouldn't personally want to over-saturate a picture like this. Easier done than said: the selection

of three areas (in the right part of the tree, in the water and in the grass at the bottom of the picture) and subsequent application of the standard CB + MMM action with its original parameters, yield figure 35. It might be too colorful for your taste, but the interesting point is the variation in the greenery on the right and in the trees – especially the leftmost.

Figure 36 is the version obtained by CB + MMM on an image previously processed with the H-K action and using both the H-K layers with standard settings, while figure 37 is the version obtained in a similar fashion but without using the H-K Effect layer (which only affects luminosity) in favour of the Color Only layer. Whether you prefer the first or the second is mainly a matter of taste: figure 37 is lighter than figure 36, so pick your favourite. (My vote would go to figure 37, if you ask.)

Both the H-K and the Color Only layer are generated by means of the Ersatz channel, and both are masked (albeit differently) in order to restrict their effect to areas of interest.


The main reason why one should want to use only one of the layers is of course connected to the nature of the image. A major goal of the whole H-K action is to prevent shadows from hyper-saturating when we go through manoeuvres like the Modern Man from Mars and Color Boost. It is worth remembering that the first of these uses a neutrality mask which prevents neutral areas from becoming too vivid; the flip side is that a colored shadow may as well undergo a variation just for this reason, not being neutral enough for the effect to be clamped by the mask. Color Boost isn't masked by default, instead, so everything will be boosted.

We can't force anything that disturbs us to neutrality, of course. In the original (figure 34) the sunny strip on the path is substantially neutral, whereas the shadow measures about (1)a(8)b. I can live with it: an attempt to remove some of the blueness in RGB without masks taken from Lab and similar tricks seems to make everything else too yellow, so I'll keep the cold-ish shadow.

The MMM version (figure 35) pushes the value of the shadow to (2)a(16)b, which is a bit too much for my

Figures 34-37. In order, original, MMM, MMM with H-K, MMM with H-K's Color Only layer. Full description in the text.





Figures 38-41. In order: original, straightforward Lab multiplication, same with full H-K previously run, same with H-K only in Color mode. Full description in the text.

taste. Both H-K'ed versions (figures 36-37) reduce this to about (1)a(11)b, which goes back to an acceptable degree of saturation.

Yet some images are of dark nature originally, and they either live or die by what happens below the midtone area, so maybe making them darker (actually compressing the shadows' dynamic range) would not be a recommendable option.

Look at figure 38 for an example: while the image is not "dark" strictly speaking, it is still a peculiar image in that it has a colored highlight and it is relatively flat, color-wise. No value in the a and b channel ever exceeds 15.

Lab multiplication is an option, and a straightforward version (with a tweak on the highlight value, which gets too dark otherwise) is in figure 39. Quite an enhancement in color, but maybe too uniformly saturated. Figure 40 uses full H-K (both layers) to tame the effect, and darker parts become less saturated, which is

better to me. But, also, the picture looks darker – which is probably not what such picture needs: yet any significant tweak to the overall lightness makes it too light for my taste, and moreover the perceived saturation decreases, somehow contradicting the process we've followed so far. Figure 41 uses H-K in Color Only mode, and is the right compromise in my opinion.

CREDITS

The pictures of the mountain and woods and the pomegranates are mine. The leaves are a crop from a photograph taken by Roberto Bigano. The insect is Marco Diodato's work, and the cow comes from Davide Barranca. The Swiss Lake is Dan Margulis' picture, while the last one, Capitol Reef, is Alessandro Bernardi's.

The jury in the leaves' evaluation was made of people who either attended my CCC classes or participate in the CCC facebook group. They are, in alphabetical order: Claudia Baldus, Andrea Bardi, Fulvio Bosco, Angelo Dau, Romy Doronjo, Andrea Iacca, Beatrice Pirani, Annalisa Santoro.

Special thanks to Giuliana Abbiati and Davide Barranca (RBG) for their constant support and help, and, last but by no means least, to Dan Margulis for getting me involved in this and for his constant invaluable suggestions and encouragement and his contributions to this article.

Dan Margulis suggested that I should explain what "Bogus Black" means, and you've read that. But the whole story behind the expression is very interesting, so here it comes – in his own words.

«The term has a historical meaning in US graphic arts. Back when newspapers were set in hot metal, their advertising typography was so poor that many agencies decided that they would purchase their own machines and submit their own mats ready for letterpress reproduction. The typesetting unions at the newspapers were very opposed to this, and refused to accept the supplied materials. A job-saving agreement was reached. The agencies' materials would be accepted and run in the newspaper, however, the newspaper's own typesetters would re-typeset everything--and then throw their work away. This practice, which persisted for decades, was known as "setting bogus".» (Dan Margulis)

Finally, I would like to dedicate this essay to my son Simone and to the inspiration of his deep dark eyes – which no Bogus Black could ever imitate.