

Defanging the Curves Vampire

Topic: An introduction to the curves command.

Column first appeared: December 1996, *Computer Artist* magazine.

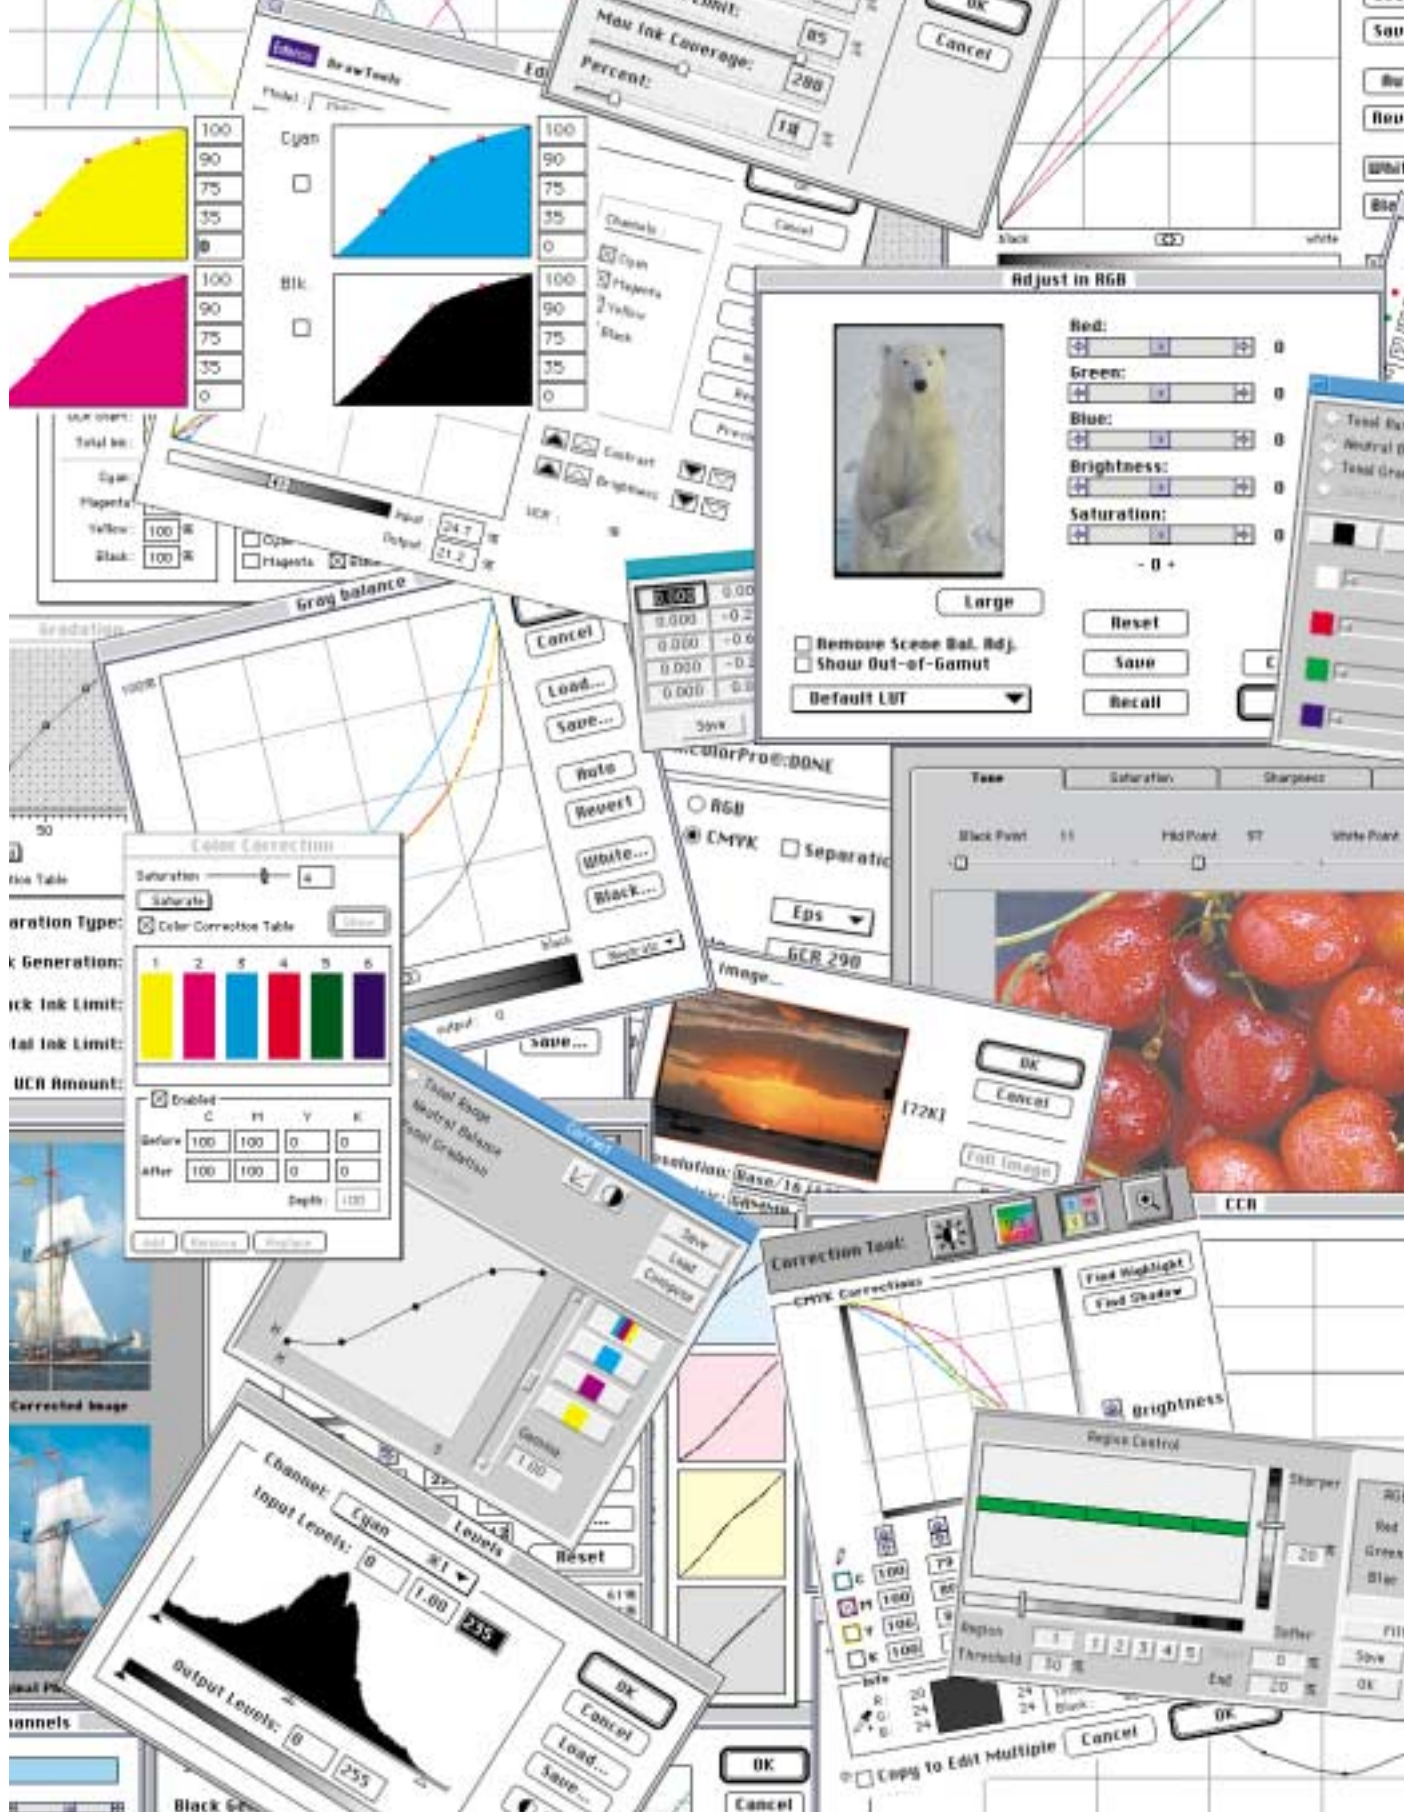
Source of this file: A revised and expanded version of the column, as it appeared in Dan's 1996 book *Makeready*.

This archive, to be released over several years, collects the columns that Dan Margulis wrote under the *Makeready* title between 1993 and 2006. In some cases the columns appear as written; in others the archive contains revised versions that appeared in later books.

Makeready in principle could cover anything related to graphic arts production, but it is best known for its contributions to Photoshop technique, particularly in the field of color correction. In its final years, the column was appearing in six different magazines worldwide (two in the United States).

Dan Margulis teaches small-group master classes in color correction. Information is available at <http://www.ledet.com/margulis>, which also has a selection of other articles and chapters from Dan's books, and more than a hundred edited threads from Dan's Applied Color Theory e-mail list.

Copyright© 1996, 2007 Dan Margulis. All rights reserved.



5

Defanging the Curves Vampire

Curves are the key to effective color correction. Everyone knows this, yet some are terrorized into using cruder methods. This phobia can be cured!



erious graphic artists, as a rule, are not buffaloes by complexity, and try to get full value from the most powerful features of the programs they use. Only the rankest novice, for example, will neglect to make use of style sheets in a long document. In illustration programs,

we all manufacture named custom colors when many objects share one color, and few of us will fail to group objects that belong together.

The glaring exception to this rule, the obvious case where many of us fail, from sheer terror, to adopt a method that the entire world knows to be superior, is the use of curves in color correction.

As the first to document this particular disease, I get to name it, and I believe it should be called kampyliophobia: the irrational fear of mathematical curves. Like the fear of other bogeymen, its elimination is highly desirable, and this is what I will try to accomplish here.

A distinguishing feature of the syndrome is the patients' ability to rationalize their fear, which they do in one of two ways. They may say,

Figure 5.1. *Input–output curves take original image data and rearrange it for another purpose: to take advantage of the shape of a new colorspace, to compensate for dot gain, to emphasize a certain color, or simply to make the image look better. Far from being an esoteric and difficult proprietary feature of Photoshop, as some believe, curves are ubiquitous in the color world. The facing page shows a host of color-altering routines from many vendors. They may not all look like curves, but maybe that is to avoid the fear factor: there's at least one curve in every one.*

first, “My color is just fine already, thank you very much,”—and they’re right, too, because they are missing the point of curves, which is *not* so much to establish color, as to enhance detail.

Every image-manipulation program has ample tools to change overall color. Color at any rate is highly subjective. You and I may differ on what color is right. We will definitely agree, however, on the desirability of something known variously as bite, snap, life, crispness, sharpness, good detail. It’s hard to put into words, but we know it when we see it. And you want it; I want it, our clients want it—and we won’t get all we need of it without using curves effectively.

The true kamyliophobic will now resort to the second rationalization: “This may all be true, but I have no experience with curves at all, and who knows what damage this inexperience may lead to if I tackle such a fearsome unknown.”

The unknown is *always* fearsome; but here the patient in fact knows the monster well. If you own a sound system, and you ever twiddle any dial other than the volume, you are undoubtedly applying some kind of curve to digital data. If you ever attempt to adjust color or contrast on your TV, you are already heavily into curves. If you ever use Photoshop’s Levels command; if you have a color management system; if you ever convert images into CMYK, use an imagesetter, make digital contract proofs, or attempt to calibrate your monitor; there is no need to fear the vampire of curves: you have already survived its kiss.

Let me be the first to admit that getting the hang of applying a set of curves to a color image isn’t easy. We will therefore start with black and white, and with a question that *is* easy, yet critical to the later application of curves: in the four monochrome images that constitute the top halves of Figure 5.2, are there certain objects, certain shades, that are more important than others?

In A, there really are not. The little girl’s collar is light; her dress is dark; and her hair and skin fall somewhere in the middle. All are important; none can be shortchanged.

The other three images aren’t like that. In B, the rabbit is clearly the most important object, and the rabbit is light. C’s kitten is dark, and the lettuce in D falls in the middle. You would presumably agree with me that if we can engineer an improvement in the rabbit, or the kitten, or the lettuce, we should do it, even if that improvement happens to hurt the background.

The next two paragraphs offer a technical explanation of

what a correction curve does. If you do not care why vampires can't be seen in mirrors so much as the fact that they can't, you can skip over them.

A curve remaps every pixel of the parts of an image it is applied to. The interface's horizontal axis represents tonal values as they currently are; the vertical axis what they will become. The default curve is a straight line at a 45° angle, since every horizontal value is the same as every vertical one. Photoshop allows the user to define whether the upper right point of the curve represents lightness or darkness; most professionals prefer that it be darkness, and this is the way it is shown throughout this book, although there is no harm in doing it the other way around.

As anchor points are added and moved away from the default line, the entire curve changes shape, affecting every tonal value, but particularly those that are close to the changed anchor points. Those areas of the image covered by parts of the curve that now fall above the default line will become darker; those covered by areas falling below the default line will get lighter.

The Steeper the Curve, the More the Contrast

Back to plain English. The whole secret is this: *the steeper the curve, the more the contrast.*

Now, the bad news. If you make the curve steeper in a certain area, you have to make it flatter somewhere else to compensate. Objects that fall in the steeper areas of the curve gain contrast. Objects that fall in the flatter areas lose contrast. If you manage to corral all important objects of the image into steeper areas, you have a winner.

The simplest application of this simple concept is found in A of Figure 5.2. The original at top is flat. The white areas of the little girl's collar are too dark and the shadow areas in the background are too light. The conventional wisdom is to change this to a proper range by measuring highlight and shadow values and correcting them, either with levels or curves. The conventional wisdom is quite correct, but *why?*

Moving the lower left point of the curve to the right lightens the highlights, and moving the upper right point to the left darkens the shadows. Everything between these two points now falls in an area of the curve that is steeper than the default 45°. Since we have made one area of the curve steeper certain other areas must be made flatter, and here, those areas are *totally* flat.







Figure 5.2. One emphasizes objects of interest by steepening the curve in the regions that affect them. In A, the image is flat. By shortening both ends of the curve, the middle gets steepened. In B, the light half of the image is clearly more important than the dark half. Raising the curve's midpoint, as at right, accounts for the added life in the bottom version.





C, being primarily dark, is handled exactly the opposite of B. As B and C are both midtone moves, they could have been done with Photoshop's Levels command. D, however, cannot. Its important area is in the middle, and to steepen the middle the curve must downplay both the light and dark ranges, without blowing them out altogether.



Everything that falls in the steeper part of the curve—the entire picture, in this case—gains contrast. Everything in the flatter areas—namely, everything either lighter than the lightest area of the image or darker than the darkest, in other words, nada, nothing, nichevo, niente—loses out.

I t's obvious that when we have a chance to help everything and hurt nothing, we should do it. From there, it's a short step to start improving important things at the expense of lesser ones.

It's pretty obvious that when we have a chance to improve everything and hurt nothing, as in A, we should do it. It is a short step from that realization to B, where we can improve the important things at the expense of the unimportant ones.

The original in B already has a full range. The whitest areas of the rabbit can't be made whiter, and the darkest areas of the background are already at around 85 percent, which is as high as our press can probably handle. A correction like that of A will therefore not work, but there is still room for improvement.

Instead of (or perhaps in addition to) moving the endpoints, we should consider grabbing a point in the middle of the curve and raising it. Let us assume that we use the exact midpoint, although this is not a requirement.

As we raise the midpoint, the picture gets darker. But what is more pertinent is that the lower half of the curve becomes steeper. The upper half, therefore, gets flatter. The lower half is the lighter half of the image. The rabbit falls in the lighter half. Since steepness equals contrast, the rabbit gains contrast. Small parts of the background lose, but we, as image technicians, win.

The kitten in C is the opposite case. If we *drop* the midtone, the curve will become steeper in its upper half and everything darker than a midtone will gain contrast. Everything in the light half of the picture will lose contrast. Few tears need be shed over that particular development.

The Limitations of Levels

Color professionals can do wonderful things in making images come to life, but far and away the most important is what we have just described: moving the endpoints of the image to eliminate contrast in areas that do not exist in favor of giving it to areas that do; plus, some kind of maneuver to allow lightening or darkening the middle of the image.

Photoshop's Image: Adjust>Levels command lets us do exactly those things, which explains why such effective corrections are possible with it. For fear of kampyliaphobic users draping ropes of garlic around their necks, Adobe does not advertise levels as being a curve approach, but a curve it is nonetheless: a curve with only three adjustable points.

The Levels command is definitely more intuitive, and for this reason most Photoshop books and instructors teach people to correct with it. Vampires, vampires! Personally, I favor charging directly into the darkened castle. Levels is a subset of the curves function and has no independent worth. Take ResEdit and hack it off your system. If you are too frightened to use curves for the simple corrections that *can* be done with levels, what on earth will you do when faced with fangs that only curves can conquer?

Before considering D, which is exactly such a case, let's review the other three. A can be corrected with levels just as accurately as with a curve. B and C can be improved with levels, but not with the same precision. The problem is, when we say that, for example, the light half of the image is more important, we don't ordinarily mean exactly the lightest 50 percent. The interest area is more likely to be centered in the lightest 58 or 37 percent, or somewhere else just as inconvenient. If it's 58 percent, then a move of the midpoint, which is what levels does, will help most of the interest area, but it will actually hurt some of it, because everything above 50 percent loses contrast.

If it's 37 percent, then a move of the midpoint has two drawbacks: first, it exaggerates contrast in some of the less important areas of the image, which may look unnatural; and, more importantly, we may not be able to make the curve as steep as if we put the point at 37 percent where it belongs. The improvement in the important areas will be great, but there will come a time when the damage we are doing to the rest of the image will be unacceptable. That time will come a lot sooner if our curve anchor point is 13 points too high.

D, with the interest concentrated in the middle range of the image, can't be fixed effectively unless you take stake, crucifix, and curves in hand and prepare to do battle. In A, only two control points were needed, and in B and C we needed three, but D really needs *four*, knocking levels into the ranks of the undead. To steepen the middle range of the curve, we have to lower the lightest point of the lettuce, and raise its upper point.

We cannot, however, afford to use the endpoints of the curve to do this. In the upper left of the background, some objects of significance are lighter than the lettuce and others are darker. We are willing to lose detail in them to a certain extent, or at least I am. I am not, however, willing to blow them out completely, which would be the only choice with levels.

You may be thinking that if three out of four images can be improved nearly as well with levels as with curves, is it really

The one-size-fits-all approach of the master curve sometimes is correct in RGB, but never in CMYK.

necessary to go to the trouble of mastering them? Perhaps not—if your practice is limited to black and white. In a color image, any object that has a color other than white, light gray, or black will have at least one channel that behaves like D. So, having established that the reputation of curves is

somewhat worse than their bite, let's consider what happens when the lettuce becomes green rather than gray.

Forget the Master Curve!

Color images differ from B/W in having not one but three channels to correct, plus a black if we are in CMYK. Nevertheless, the same goals are there: full range in each channel, and in each channel as much range as reasonable allocated to the important areas of the image, if these are identifiable. Plus, there is now a third goal, color fidelity, that was not applicable before.

A temptation now appears, in the form of the inclusion of a master curve in both RGB and CMYK modes. The simplicity is seductive, the apparent improvement palpable, and the sinister nature almost undetectable.

The master curve is a one-size-fits-all approach, affecting each channel in the same way. Occasionally, that is right in RGB. In CMYK it is *never* right, since black will always behave differently than the color channels.

If we are dealing with an RGB version of the image of the black kitten, a lightening of midtone value via the master curve will actually be acceptable. The kitten will be just about as dark in one channel as any other.

But when the main subject is not neutral, not black, but rather the green head of lettuce in Figure 5.3, this approach will fail. The channels in such an RGB image look no more like one another than Bela Lugosi looks like Tom Cruise. The green

obviously is strongest, and red nearly as heavy, since red plus green equals yellow. The blue channel will be weak. The lettuce falls in different ranges in each one. Only a separate curve for each can be successful.

Writing them requires a little preliminary investigation, for which I will switch to CMYK. A quick look indicates that the strength order of the channels is yellow, cyan, magenta, black. We need, however, more specific values in four areas: the light and dark point for the image as a whole, and also, assuming that we will have a go at the foreground lettuce, the lightest and darkest point of that.

The lightest point anywhere is plainly the bowl. It starts at 5C4M14Y. In my opinion, this bowl should be white (we aren't always lucky enough to find a white object), so a standard highlight value for this will do. I generally like 5C2M2Y.

The darkest point, just to the right of the bowl, measures 81C56M82Y56K. An appropriate shadow value for magazine work is 80C70M70Y70K.

Within the lettuce itself, the lightest point other than a reflection reads 45C11M95Y1K, and the darkest that still seems green is 75C47M100Y31K. That's roughly 30 points of range for all colors except yellow.

The Three Big Objectives

Armed with this information, we proceed to create four curves, simultaneously trying for the three big objectives:

- To extend the range of the entire image as much as possible, consistent with retaining detail
- To get colors that we like, particularly in neutral colors and fleshtones, if any
- Given the constraints of the first two, to allocate as much contrast to the most important areas of the image as possible

In the context of the values we have just discovered, the first objective is not too far off to begin with. The highlight is almost right, and the shadow is only slightly too light.

To achieve the second objective we need to do some work. The plate should be white, and the deepest shadow should be neutral also. Neither one is at present. Both are too yellow. The shadow also is deficient in magenta. We will be trying to extend the range of the lettuce in each channel, but in doing so, we will have to make sure that these problems are also solved.



Figure 5.3. Effective correction of a color image can best be done by assuming it is four grayscale images. In the lettuce opposite, the yellow plate is heaviest, and should be treated as C was in Figure 5.2. The cyan has its lettuce values in the middle and thus behaves like D; the black is light in the lettuce and should therefore be treated like B; and the magenta, being both light and flat in the original, behaves like a combination of A and B. The composite curves, below, transform the picture at top opposite to the much better one at bottom.

The four curves that do this appear in Figure 5.3. You could approximate the yellow and the black corrections using levels, but not the magenta and cyan, both of which, like D, absolutely need four points.

The additional life in the lettuce is all the more striking when compared to what went on elsewhere. If anything, the salad in the upper left of the background is worse than the original. That is the work of the curves. The price for the steepness that added contrast to the foreground lettuce was a flatter curve elsewhere. That's a very fair price to pay, I'd say.


The fadeout of the background illustrates, oddly, one of the reasons curve corrections are so effective. The curves emulate what our eyes would have done were we looking at the lettuce in person. The human visual system, in effect, applies curves to whatever it sees. If we are focusing on a green object, we become more sensitive to variations in green, and less to other colors. We emphasize differences in both color and brightness in closely similar colors, and we do so at the expense of detail in whatever is in the background. Cameras and scanners do not do this.

This is therefore a much more natural, not to mention easier, way to correct images than isolating the lettuce by making a selection and then working on it alone. Even if the selection can't be detected, there always seems to be something just a little bogus about such work, and the skilled eye rebels at it.

The Best Curving Tool

I hate to proselytize, but if you become comfortable with curves, you will make your life in the graphic arts ever so much easier, and not just in Photoshop.

Curves are, as Figure 5.1 shows, to be found throughout the graphic arts. If you are proficient with them you will understand






Figure 5.4. *The ultimate curving tool.*

everything you need to about how a file is separated into CMYK from RGB or LAB, about how to calibrate an imagesetter, digital proofer or film recorder, about how to compensate for adverse printing conditions, and a

whole gang of other useful things.

If you first began to use a computer as an adult, as I did, you will probably recall the frustrating experience of learning to use the mouse. Remember how it always kept sliding off the pad? Remember how we could never get the hang of clicking vs. double-clicking? How at first we weren't sure which direction to hold it in, and how we kept wrapping the wire around our wrists?


When you look back on that aggravating time, you will probably agree that that blasted mouse was a major obstacle to actually learning useful things about the programs. Now that we have presumably reached the point where the mouse feels just like an extension of our hand, we can get down to business.

The same thing can happen with curves. Some propellerheads spend hours trying to get the color right in a file that could be fixed up in seconds with a single set of curves. They ordinarily end up with 14 or so supplementary channels and five or six layers, each with a mask, not to mention a file size of a hundred megabytes. Each time I watch it happen I am reminded, in a way, of someone who has not yet grown comfortable with the mouse.

When you do get comfortable with curves, you will realize that all the previous terror was unnecessary. It will be as though a scar has vanished from your forehead.

If you can learn the mouse, you can learn curves. They are both easy once you get past the initial, er, learning curve. Reproduction curves become particularly intuitive when you consider that Figure 5.4 portrays the most powerful densitometer, the most powerful colorimeter, the most powerful curving tool in existence. Lucky you, to happen to own a pair of them. ●

Figure 5.5. *A curve emulation of what happens to an image when it is printed on poor-quality paper, discussed in the Afterword.*



Afterword

In real life in the magazine, one can't easily refer readers back to graphics in previous columns. Books, happily, are another story. I would like to suggest a reconsideration of Column 3, which discusses dot gain and its ramifications. Analyzing the problem in terms of curves is a useful illustration of the points just covered. In Figure 5.5, let's try to draw a curve of what happens when we expect high dot gain.

In Photoshop's simple-minded view of the world, dot gain is heaviest in the midtone, less apparent elsewhere, and nil in the highlights and shadows. If we expect higher than usual dot gain, it reasons, we should be able to simulate it by raising the midpoint of our hypothetical curve. Entering a higher dot gain value in Printing Inks setup does exactly this.

In reality, though, things are not nearly so straightforward, for two important reasons ignored by Photoshop. First, heavy dot gain implies poor paper, and poor papers are usually not as white as good ones. If the paper is grayish, this is more or less the equivalent of printing on white paper but adding five percentage points to each ink throughout the picture. That means we should revise our theoretical curve, raising the zero point at bottom left.

Second, poor papers have higher dot gains because they are usually more absorbent. But in deep shadows, darkness depends more on how much ink lives on the *surface* of the sheet, rather than on the apparent diameter of the

halftone dot. Translation: shadows are darker on good paper. Curve translation: drop the top right point. Also, because this effect starts to show up in the three-quarter tone as well as the shadow, push the previously drawn midtone point to the left.

Now, remembering that Photoshop's screen display only compensates for one of the three points we have drawn so far, let's consider to what extent the monitor is going to deceive us, even if we enter a properly high dot gain value in Printing Inks setup.

That adjustment will set the overall darkness of the image and probably do a pretty good job of it. It will not, however, compensate for the grayness of the paper. Therefore, all colors will seem to be cleaner and purer than they actually are, and we should compensate by increasing saturation generally.

More importantly, look at how flat the upper right of Figure 5.5 is. That flatness, of course, is a contrast-killer in the darker half of the image. In images such as the frogman of Figure 3.1, which depends on shadow detail, the monitor will give a hideously overoptimistic view of how much of it there will be on press.

You can drive a stake right through the heart of this, by *dropping* the midtone in the original image to steepen the upper half and thus regain some of the lost contrast. If you blithely, blindly trust the monitor, it probably won't occur to you to do so. If you remember the implications of curves, it *will*. ●

Jargon Watch

CMYK Cyan, magenta, yellow, black, the inks used in printing color images, but also a major storage space for digital images.

color management A much overhyped concept: the use of third-party hardware and/or software to attempt to make colors more or less match between two, and frequently more, different kinds of devices.

curves A remap of the tonal values of an image, for color correction, range extension, or enhancement of detail. Normally used to refer specifically to Photoshop's Curves command or the equivalent in other programs; however, many simpler routines, such as levels, or the remap applied during the separation process, use an approach that is technically one of applying curves.

digital contract proof A high-quality color proof suitable for use as press guidance, but produced on a ink-jet, dye-sublimation, or other digital printer, without the use of any film. This differentiates it from traditional contract proofs such as Cromalin or Matchprint. Digital contract proofers, such as Kodak's Signature or Scitex's Iris, have been available for some time but have been extremely expensive. In 1995, several vendors introduced dye-sublimation digital contract proofers for less than \$20,000, a considerable price break.

highlight value The lightest value that a given process can accommodate, consistent with retaining detail.

levels A Photoshop tonal-correction tool, simpler and somewhat friendlier than its curves function.

mask A means of confining an image correction to a single area without the possibility of altering the remainder. It is in principle the same as a selection except that *mask* tends to suggest a more complicated correction. Masks are frequently saved in separate channels.

ResEdit A powerful but dangerous Macintosh system utility that allows users to mess around with things like data and resource forks. It has some practical uses in font management among other things. I don't actually recommend that you use it to hack levels off your system, but you could if you wanted to.

RGB Red, green, and blue, the colors of light to which human eyes are most sensitive. Digital image files are frequently kept in files represented by RGB channels; these, however, must be converted to CMYK before printing.

selection In digital retouching, the isolation of a certain area for local correction. The software will not permit nonselected areas to be altered.

shadow value The darkest value that a given process can accommodate, consistent with retaining detail.

Dialog Box

Please excuse an amateur's question, but what exactly do you mean by "balancing highlight and shadow?"

In printing parlance, the "shadow" does not mean what it does in standard English. It means the darkest significant area of an image where the viewer is not supposed to notice color. In other words, something that is navy blue is not the shadow even if it is the darkest area of the image.

The "highlight" is an area that is *both* the lightest significant area and one that we want the viewer to perceive as being white. In evaluating whether an area is the lightest one in the image, we ignore areas that have to do with the actual transmission of light, such as a light bulb or a reflection of light off a mirror.

By the above definitions virtually every image has a shadow. The majority have a highlight as well, but plenty do not.

Balancing the highlight and shadow means setting these areas to neutral whites and grays, not favoring any color. This standard

Highlight and shadow defined, and why professionals balance them

is universal in the prepress industry, as is extending the range of the image so that the highlight and shadow are in fact as light and dark

as the process will allow, consistent with holding detail.

One of the major consequences of this is to bring the whole picture, not just the highlight and shadow, closer to what the viewer expects to see, reducing or eliminating any color cast. If the original picture contains a slightly green horse, balancing the highlights and shadows will make it less green, whether or not it ever occurred to us that we wanted it that way. There would be some debate in the industry as to whether, after balancing the highlight, we should go in and evaluate the horse to see if it was still somewhat green, and if so take action. But there is no debate at all as to the propriety of balancing highlight and shadow.

When you're making this correction, are you adjusting the highlights and shadows to neutral via some sort of quasi-colorimetry technique (scanning the highlight and shadow areas and looking at the relative values of R, G, & B)...

Well, in C, M, & Y, in my case, although you can do it either way. Yes. Whether doing this during a scan or in Photoshop afterwards, one checks what the values of the highlight/shadow currently are and changes them to what one would like them to be.

...or are you just twiddling the (virtual) knobs and looking at the screen in PhotoShop?

Leaping lizards, no!!!!!! The screen cannot be trusted for this purpose, unless you happen to be a spectrophotometer, not a human being. If you are a human, you have been gifted (handicapped?) with self-calibrating vision. That is, whatever light is hitting your eye is simultaneously adjusting your color perception for you. This is not so bad when you are looking at something with a pronounced color, such as a face or a leaf, but if what you are looking at is a highlight or shadow, it's a disaster. Your visual system moves in the direction of neutralizing any color imbalance. Therefore, the longer you look at a light source, such as a monitor, the grayer it will seem to you. This is highly awkward if you are trying to figure out whether the object is *actually* gray. You *must* trust the numbers and not the screen in evaluating highlights, shadows, and neutral colors. Trust the screen for overall appearance, perhaps, but for these three items, a "calibrated" monitor is a contradiction in terms.

And am I understanding correctly that your usual practice is to "stretch" the contrast range of the image to create "true" highlights and shadows where the original had no highlights? (I've shot one or two of those.)

Yes. Prepress differs from photography in that regard, an unhappy necessity in view of how little colorspace we have to play with. This is yet another reason that scans perfectly calibrated to an original piece of film make little difference: we have to change the range in virtually every case regardless.

There is a very interesting consequence of all this. Once they go through this process, amateur photographs that have been printed at the local quick photo booth look very much like professional chromes. Most people suppose that the better and more quality-conscious the photographer, the more he or she needs a drum scan. Just the opposite is true. If you have a

technically excellent photograph, just about any old scanner will be good enough. Where the high-end equipment really blows away its competition is where it has to get first-rate results out of third-rate photography.

Your term “where the original had no highlights” is a little ambiguous. Virtually every image has a shadow, but some don’t have a highlight. If what you mean is, there is a highlight but you shot it too dark, yes, we stretch it. If what you mean is, there is no natural highlight, no, we don’t force it to happen. Highlight must be *both* white and the lightest non-reflecting area in the image. If a man is wearing business attire, including a white shirt, somewhere in that shirt will be the highlight. Change the shirt to light blue (assuming he’s not wearing anything else white) and we do *not* set the highlight to be the whites of his eyes.

In such a case one can figure out what the highlight would logically have to be if it were there, and set the image accordingly. This is, of course, easier said than done, but we prepress types have to justify our salaries somehow. ●