A Rock and a Hard Place

Topics: Separation methods, gamut limitations.

Column first appeared: June 1998, *Electronic Publishing* magazine.

Source of this file: The author's draft as submitted to the magazine.

Author's comment: Until a few years before this column was written, most work intended for press was drum-scanned directly into CMYK, so the question of how to convert RGB into CMYK wasn't relevant. The age of desktop publishing changed that, and provoked several vendors to announce that they had found foolproof solutions to the problem. Here, I debunked those claims, pointing out that there is no one-size-fits-all conversion method, unless one is willing to settle for the equality of mediocrity.

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.

A Rock and a Hard Place

Color separation algorithms leave us in a dilemma: do we want good results most of the time, or bad ones none of the time?

By Dan Margulis

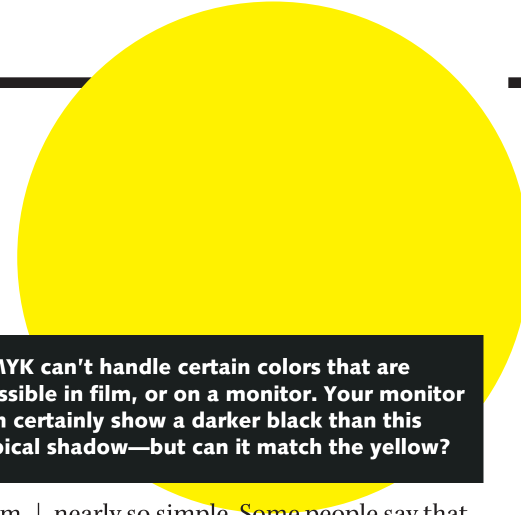

CMYK can't handle certain colors that are possible in film, or on a monitor. Your monitor can certainly show a darker black than this typical shadow—but can it match the yellow?

American football is played on a rectangular field roughly 50 by 100 yards, exclusive of the end zones. The field in the Canadian version of the game is about 10 yards wider in each direction.

The larger field mandates certain rule changes. For one thing, Canadian teams have a twelfth player. Strategies change a little as well, but the same plays work in both games, and the same skills differentiate the star player from the mediocre one.

A football fan therefore adjusts easily to watching either version of the sport. But suppose that the differences were much greater. Imagine a kind of football played on the side of a hill, rather than on a flat surface, and on a trapezoidal, rather than rectangular, field, with a brook and a few trees in the middle of it.

Once you realize that in such a game the set plays and strategies of conventional football no longer necessarily work, you are well on the way to understanding why so many people have trouble making decent color separations. To be more precise, they are having trouble translating RGB scans (or digital captures) into CMYK.

For that matter, we are starting to need new types of separations that involve different flavors of CMYK (as for both a newspaper and an annual report) or devices that use more than four inks or toners in an effort to get snappier color. Several desktop printers have this, and we'll probably see more of it in large-format printers as well. So building our skill in separation will be critical to our future success.

Prepress professionals don't have a whole lot of experience in solving this problem. Up until about five years ago, most separations were done on drum scanners that converted to CMYK on the fly. An RGB file never existed, so the question of whether the CMYK file looked like the RGB never came up.

Jumping to the conclusion that making the conversion must be easy, if only one has enough expensive color-measurement devices and software, various parties have hyped "color management" solutions which, quite correctly and quite predictably, the market has emphatically rejected.

The reason these products haven't flown is not a lack of sophistication, or inadequate computing power. The reason is, the whole concept is wrong. Those whose quest is the perfect separation algorithm are chasing rainbows, setting traps for unicorns.

Indeed, the perfect separation method is a mythical creature, but one with a substantial sting: the closer one tries to get to it, the farther away it seems to be. This column will try to explain why.

Decisions and damage control

As with football, trying to translate between colorspaces with is only hard when the rules they play by are *radically* different. A monitor and a chrome have slightly different gamuts (the monitor can get a richer blue; the film can achieve better secondary colors) but these differences in what colors can be had are small in the overall scheme of things. So it isn't difficult to create RGB files that more or less match the chrome. It is also easy to adjust one professional digital proofing system, such as Iris, to match another, such as Approval.

Going from RGB to CMYK is not nearly so simple. Some people say that this is because the playing field is smaller, naïvely ignoring that it is tilted as well. Let's take a quick survey of what RGB can portray that CMYK can't—and vice versa. For this, you will need your imagination, because while I can *tell* you what colors aren't possible in CMYK, for obvious reasons, I can't *show* you.

The differences can (and should) be divided into two categories: color and contrast. Contrast is mostly a matter of how bright and how dark the white and black points are. In this area, CMYK is pretty lame. The black in the box above is about as deep as we dare get in a printed image—darker blacks are possible, but would cost detail in shadow areas. Monitors can display a much richer black. Also, we can't make a white any brighter than the paper we are printing on.

The non-level playing field: since CMYK produces great yellows and terrible blues, how are we supposed to handle an image where the two colors are in balance in RGB?

Because there is less of a darkness range available, the CMYK practitioner needs to emphasize contrast. In football or hockey, a larger playing surface rewards speed and finesse, and a smaller one favors physical strength. CMYK is much the same thing.

But the playing field, in addition to being smaller, is also weirdly shaped. Those infatuated with third-party color management dismiss CMYK as simply not having the color range of RGB. In some respects, it indeed doesn't. In others, it has *more.* Let's contrast the capabilities of a monitor to those of commercial printing.

First of all, the building blocks of each are different. The phosphors of a monitor are red, green, and blue, highly convenient if pure red, green, or blue appears in the image. On press, red, green, and blue are each mixtures of two inks, which is a disadvantage. On the other hand, CMYK is well equipped to produce pure cyan, magenta, and yellow.

Especially yellow. It's the purest ink, and under good printing conditions, a stronger yellow is available than can be had in real-world RGB. Can your monitor match the intensity of the patch of yellow on page ··? Solid magenta and cyan also can be as intense on paper as they are on a screen.

As these colors get lighter, however, CMYK has more trouble with them. Bubble-gum pink being a shade of magenta, you might think that you could portray it as well in print as on a monitor. No way. As colors get lighter, they get represented by smaller and smaller dots, and accordingly, larger and larger quantities of blank, featureless paper. The monitor has no such dot structure, and can create much more appetizing-looking bubble gum.

And the notorious weakness of print work is that cyan ink does not mix well with magenta. Therefore, although reds and greens in CMYK are somewhat worse than those available in RGB, the blues of CMYK are *far* worse—the clearest example of the tilted playing field.

To summarize the rule differences: in CMYK we have better yellows, about the same magentas and cyans, but lousy reds, worse greens, and disastrous blues, in comparison to RGB. As the colors get lighter, everything changes in favor of RGB, except that the CMYK disadvantage in blues is minimized. Also, CMYK lacks contrast generally.

That's a mouthful. Properly understood, however, it explains why third-party color management has thus far been such a flop.

Almost every possible color is portrayed in this staged image, and only a sophisticated separation algorithm can get it right on the first shot. Ignoring out-of-gamut colors, as the EIAM described in the column would, might result in a version like the one at bottom.

A question of aesthetics

The aspen forest on page ·· speaks starkly about the injustices of CMYK. Taken from Digital Stock's new Nature & Landscapes set, this is an outstanding image—in RGB. It loses a lot in translation to the printed page.

In the original, the sky is lighter than what you see here, but also much bluer—a nearly luminous, icy, brilliant light blue. I have more chance of playing tackle for the Edmonton Eskimos than of reproducing that color accurately in *Electronic Publishing.*

On the other hand, part of the reason the blue is so striking in the original is that it plays off against the bright yellows of the leaves, a CMYK strength. We are not at the limit of yellow ink yet. I can make those leaves yellower still, brighter than they were in the original. But should I?

Granted that we can't match the original or even come close, there are many ways to try to make the best out of this bad situation. Should we

•Tone down the yellow, in order to keep the relative balance with the blue?

•Ratchet the yellow up, so as to accentuate the *contrast* between yellow and blue?

•Wipe out any yellow or black ink in the sky, which will wipe out detail as well, yet make the sky seem bluer?

•Increase cyan ink in the sky, to make it bluer, albeit darker?

•Or, is it possible that the answer is none of the above, but rather the image just as it appears here? Not too likely: this separation was done in Photoshop using the application's default setting, which is not good for this type of image.

I suggest that this is a problem without a solution. Some images will look better if we make the colors more vivid. Others will not. Human beings make such aesthetic decisions routinely, and accurately. Profiles and other algorithms are rather bad at it.

The EIAM and the PCCM

Now, let's consider a general approach to converting a document from RGB into CMYK. I will commence hostilities by proposing a method so absurd that you may have trouble recognizing its intrinsic logic.

Here it is: for every RGB color that *can* be faithfully reproduced in CMYK, do it. For every color that *can't,* do something completely random, such as make the CMYK color lime green. Because of this uncertainty, I dub this approach EIAM, which stands for, Every Image an Adventure Method.

EIAM has, to put it mildly, distinct disadvantages. For example, if it is used to convert the aspen forest image, the sky will become lime green, which is unlikely to please our client.

If this sounds very radical and unreasonable, it is, *but no more so than the current politically correct way of doing things.*

The Politically Correct Calibrationist Method, or PCCM, often involves the use of some artificial color-measurement device measuring swatches rather than the obviously superior method of a human eye deciding whether two images look alike. In PCCM, to the extent that actual images are used to try to match the two colorspaces, they tend to be grossly atypical ones, such as the Scitex piece shown opposite.

This approach is the wishy-washy one of finesse and compromise, just as the EIAM is the blunderbuss method of brute force and hope for the best. On this spread, note how each handles *certain* images better than the other.

Before you write off EIAM for what it did with the fabric shot, remember that this particular image was *made* for PCCM, specifically composed to present virtually every kind of difficulty that the separation process can encounter. Photoshop's "Olé No Moiré" and Kodak's "Musicians" are similar. There is critical detail in both highlights and shadows, there are out-of-gamut colors everywhere, there are neutral colors that must be retained, and subtle shadings even in the most brilliant colors.

You could spend a lifetime in the graphic arts and never work a live image that has all of these characteristics at the same time. If you *do* encounter one of these monstrosities in a live job—don't hold your breath—PCCM is definitely the best way to separate it.

But it attains its quality at a considerable price. Think about the critical color, blue. In the original RGB, we have certain blues that are simply too brilliant to be duplicated in CMYK. We have some other bright blues that *could* be duplicated, if we made them as blue as possible, which is what EIAM, the steamroller, would do.

PCCM, the great compromiser, tries trying to retain some distinction between the two kinds of blues, tones both of them down—along with every other blue down the food chain.

That works well enough in this particular image, but what if there were *no* brilliant blue in the original? Then, the compromise would be pointless. We would be toning down our in-gamut blues for no reason. Here is where EIAM, which doesn't do any toning down of anything, would have a decided advantage.

Despite its calibrationist, match-the-art aura, PCCM guarantees that we will *never* match the art—*all* colors will be toned down, and all images will look flatter than the original.

EIAM, on the other hand, is for the high roller. If the art can't be matched, catastrophe! But if it *can*, EIAM will do it, and in those images it will do much, much better than PCCM.

A general law, sad but true

Which of our two proposed methods of separation is better depends very much on your definition of *better.* If the definition is, which method produces more *acceptable* images, obvi-

When an image has no out-of-gamut colors, as in Goya's famously ugly painting of the Queen of Spain, over-sophistication can give a famously ugly separation. Why suppress colors, as a finesse method must? Here, the EIAM (bottom) is a clear winner.

ously PCCM wins: it is stolid, stodgy, and free from ridiculous errors.

It is also a recipe for mediocrity.

Suppose, though, that the question is, which method works better *most of the time?*

Guess what! Most images don't contain out-of-gamut colors. And for all those that don't, that silly EIAM will kick butt, as it does in the Goya painting on the preceding page. How good can political correctness be, when an *absurd* method gets palpably better results on the majority of images?

EIAM *is* absurd. In real life we don't deliberately sabotage images, the way EIAM would to anything that contains an out-of-gamut color. So, if forced to choose one or the other, we have to pick PCCM, because even if we are dissatisfied with its results, we can perhaps fix it, which is more than can be said if EIAM starts dispensing lime-green pixels all over the place.

But intermediate approaches are possible—and very, very practical. Confronted by an out-of-gamut color, EIAM drops back fifteen yards and punts. It is content with a totally unacceptable image. We can, however, visualize a smarter method with all the advantages of EIAM. Such a method would fake it by substituting for out-of-gamut colors not lime green, but something close to what the colors actually were. Granted, a lot of detail might vanish in these areas.

With PCCM, we have *acceptable* color 100% of the time, but if its color isn't bad, neither is it good. With pure EIAM, we have a better image than PCCM maybe 60% of the time. The other 40% of cases are disastrous, full of lime green and totally unacceptable.

The less ridiculous version of EIAM described above does better. Now, we may beat PCCM 70% of the time. An additional 10% of the time the image will be acceptable, yet not so good as PCCM's. The remaining 20% remains, well, unacceptable.

The time has now come to state the law that governs all transformations from one colorspace to another. It is a sad law, a rock-and-a-hard-place law, but an uncompromising, invariable one. Here it is:

The better the algorithm does on the typical image, the more prone it is to do something really objectionable to ones that are not typical.

Interpretation: our choice really depends on whether we want something to be good, or whether we want it not to be bad. The difference explains a lot about what people constantly say about the separation process.

For one thing, it explains why so many people accuse Photoshop of having a "bad" separation engine. What is meant by that, of course, is that Photoshop, like its relative EIAM, sometimes lets loose with a real howler, changing blues to purples with great elan and losing detail in out-of-gamut colors. If you see enough of these stinkers, admittedly you may think that Photoshop itself is what stinks. But all it is doing is following my law: since it generally makes *good* separations, it frequently makes *bad* ones.

The law also explains why, after ten years of endless hype, third-party color management remains such a colossal bust, at least as far as the conversion into CMYK goes. All such programs are PCCM-based—which means that for the majority of images, they will not do even as well as the idiotic EIAM, let alone a more sophisticated variant, such as Photoshop. What third-party color management has so far offered, in general, has been applications that are expensive and hard to learn. Considering that they also don't provide separations as good as Photoshop's *most of the time*, it adds up to a pretty severe market handicap.

What does it all mean?

There is no point in avoiding the easy changes that make the separation more accurate, such as increasing Photoshop's dot gain adjustment when separating an image for use in a newspaper. But an accurate separation only goes so far.

The Goya painting is a great example. The bottom one is better, sure, but neither one is very good. Both fail, due to lack of detailing in the black dress. That can be fixed, and it must be, if we are at all interested in quality.

The funny thing is, if we know how to correct the image, it basically makes no difference whether we start with the top or the bottom version. One does a by-the-numbers correction, which is just as easy with one as the other and should yield the same result.

I recently heard from the head of a large in-house production facility, who took askance at some of the things I said a while ago in a newsgroup,

"You described," he wrote me, "how the perfect profile will only produce a perfect CMYK file, maybe 15% of the time. You also made several other comments about people's expectations of color management.

"Well, it's five months later and after testing every profile-generating software I could get my hands on, damn if you're not right. Please understand I had to see this for myself. There are some advantages to using ICC profiles if you happen to use a ColorSync-savvy application such as LinoColor. But even LinoColor has its own print tables that I've modified and gotten better results than those generated from measured targets. Regardless, every image is different and usually requires intelligent intervention. (Hopefully it will be intelligent)

"It was sure interesting to re-read your post five months later and understand first hand, every point you made."

If there really were one best way to convert into CMYK, it would have been discovered a long time ago. Meanwhile, there are many reasonable variations on the market. If you don't particularly care about image quality, it won't much matter which one you use. If you *do* care, and if you know the right way to get there—well, then it won't matter much, either.

Dan Margulis *can be reached by e-mail at 76270.1033@compuserve.com, or by fax at 973/763-2835. For information on Dan's color-correction tutorials in Chicago and Atlanta, contact Judy Starkweather of PrimeSource at 800/992-4897.*