A Closer Look

Writers are notoriously sensitive about others doing violence to their prose, none more so than the irascible poet, swordsman, romantic, and color theorist, Cyrano de Bergerac. A friend suggests that Cyrano's play might have a better chance if he enters into a business partnership with an influential cardinal.

> DE GUICHE. He's not a bad writer himself.
> He'll edit only a line or two.
> CYRANO. Out of the question, Monsieur.
> My blood curdles at the thought
> Of changing a single comma.

Cyrano's belief in the immutability of his own prose is an excellent introduction to a discussion of a recurrent myth, that converting into LAB somehow damages image quality.

To that, there are two replies:

- It doesn't.
- If it did, it wouldn't matter.

The second, being pure logic that even the maniacal Cyrano would approve of, is easiest to justify. We've now seen several ways to improve images using LAB. If you agree that the images actually are improvements, could you have achieved similar results without the use of LAB? If not—if they're better now, and they couldn't have been made that good any other way—then qu'est-ce que c'est this word *damage?*

As for the first, I was given to understand in 1996, from several persons claiming to be authorities on Photoshop, that converting to LAB even once would inflict "catastrophic" damage on the innocent image. It's not obvious what this means, but a likely guess would be, if you convert to LAB, you can never get back to where you started from.

Why anybody would want to do that, I don't know. The whole point of going to LAB is that we *don't* like the current image. But, shrugging my shoulders, I offer Figure 6.8. One copy (it should be as obvious as Cyrano's nose, some say, so I will not bother to identify which is which) is an original professional image, appearing here at 250 pixels per inch. The other has been converted to LAB, and then back to RGB.

And back again to LAB, and back, and back. 25 trips to LAB, 25 trips back to RGB. Because there could be defects or artifacts that aren't apparent at 100% magnification (Figures 6.8A and 6.8B are at the correct size for the image's resolution), I've also provided blowups of three different areas of the two images, including one of the critical green channel in the face area.

Similar demonstrations in the past have been met by commentary to the effect that the disastrous effects of the conversion to LAB only occur during the *first* conversion, and that repeatedly converting the same file over and over proves nothing, but that if it were converted at several points during the correction process, then, morbleu, quelle abomination!

While disposing of that, I want to deal with two other myths. First, it has been asserted that bad things happen if you work in a big-gamut space like LAB without being in 16 bits/Channel mode, which doubles the size of all files and adds several irritating steps. Second, it is alleged that the effect is particularly pronounced when the modifications are extreme, and even more pronounced when there is a series of modifications rather than just one or two.

Everything humanly possible has been done with the next example to incorporate the kinds of things that are supposed to provoke these problems. The original data quality is good. The bottom half is a digital capture, the top half is scanned film of a scene of different character.

Unlike Figure 6.8, where we took a good original and threw it back and forth 25 times, this time we have to make large changes to the file, just as the myth requires. Also, instead of doing it all in one move, I catered to the myth by dividing it into seven smaller moves. Furthermore, I've thrown in conversions to and from LAB after each step. Since the file is in a different state each time, these are seven of the catastrophic first conversions.

Figure 6.9, which continues with enlarged pieces in Figure 6.10, has four versions.

The start point, which is not shown, is a 16-bit RGB file, which had been artificially lightened to enable a series of darkening moves. As you can see by the cast, I wasn't particularly interested in color fidelity, because it has no bearing on the present topic. The claim is that the damage will show up as excessive noise.

The test image consists of two halves. Both images were professionally shot; the top half is scanned film, the bottom half a digital capture.

These four versions came about as follows:

- 1. The seven sets of corrections were applied to the 16-bit original in 16-bit RGB. This is the politically correct approach. The file was only converted to 8-bit for printing, as printers cannot accept 16-bit files.
- 2. Same as above, except the file was converted to 8-bit immediately and all corrections were applied in that mode.
- 3. Same as #1, except after each of the seven steps, the file was converted to 16-bit LAB and then back to 16-bit RGB.
- 4. Same as #2, except after each step, the file went to 8-bit LAB and then back to RGB.

As with Figure 6.8, all variants are in a random order. If you care to guess which is which, the answers are in the "Notes & Credits" section starting on Page 351.

Personally, although I don't know whether it will be visible in print, if the proofs are accurate I think I could take a flyer about three of the five versions of Figure 6.9. But the difference is slight, we're looking at single channels at high magnification, and anybody who has to do corrections this extreme has many other

***Figure 6.8** One of these images is the original. The other has been converted into LAB and back into RGB 25 times. Opposite page: magnifications of key areas at 200%, 300% (green channel only), and 400%. In each case, the placement of the two versions is random. Can you tell which is which?*


C
D
E
F
G
H

things to worry about than bit depth and conversions. As for Figure 6.8, I state categorically that the two versions are identical for any conceivable professional purpose.

And Why Not Look, If You Please?

A boring bystander is unfortunate enough to be confronted by an enraged Cyrano, who imagines the man is staring at his nose.

> CYRANO. What do you think? Is it not a phenomenon?
> THE BORE. But I knew better than to look!
> CYRANO. And why *not* look at it, if you please?
> THE BORE. I was…
> CYRANO. Does it disgust you?
> THE BORE. Monsieur…
> CYRANO. Perhaps you do not like its color?

At the close of the scene, the bore is lucky to escape by being smacked on the side of the head rather than being run through with an épée. A good cuffing might do wonders for his counterparts in the Photoshop world, those who are so certain of their ground that they know better than to look.

And now that we *have* looked, and know the correct answer, it must be conceded that sometimes the theory seems so obviously true as to render any alternative inconceivable. This is a compelling example. How can moving to LAB *not* cause damage? We're throwing away (so they say) a third of the colors!

Translation: the original RGB file consists of three channels, each of which has 256 possible values, or *levels of tonality.* If we consider two channels simultaneously, each of the 256 values in the first channel has 256 more possibilities in the second, for a total of 65,536 possibilities. If we add a third channel, each of these 65,536 has 256 more, for a grand total of 16,777,216 possible combinations.

I don't know how many distinct colors are in the original version of Figure 6.8, but it isn't 16,777,216. For example, there's no bright yellow anywhere. $210^R210^G40^B$, which is a fairly subdued greenish yellow, isn't likely to be found. And neither is anything with higher red, green, or both, coupled with an equal or lower blue.

If you agree, 86,756 possible colors have just fallen on their swords. Pastel blues, brilliant greens, and all cyans are also among the missing. Plus, there may be some luck of the cards. $50^R50^G10^B$ might easily be found in the woman's jacket, but there's no guarantee that even a single pixel will have exactly that value.

Some programs can analyze exactly how many discrete colors such a file contains, but I don't own one. My guess is that in this image it's a lot more like 10 million than 17 million.

But now, let's take it into LAB for the first time. There should be around 256 values in the L, granted. But there won't be anything like that in the A or B. With no really brilliant colors in the image, it would be surprising to see values more than ±50 in the AB channels. So, there are maybe 100 values in each one, tops.

Having just said goodbye to 14,217,216 colors, it only gets worse. As the L gets closer to its endpoints, the AB possibilities are sharply reduced. By the time we're at 5^L or 95^L we may be down to only 20 real possibilities in each AB channel.

To be generous, let's shortcut a lot of arithmetic and estimate that for each L value there are 60 possibilities in the A and B. If that wild guess is exactly correct, there are 921,600 possible colors in the LAB version. Since it isn't, let's call it a million. And we estimate that the RGB picture contains 10 million colors. We are throwing 9 million of them away by converting, no?

This is much worse than the advertised loss of

***Figure 6.9** These two images, one digital, one from film, are joined in one file. Originally they were quite light, but instead of correcting in one pass, this drastic change was done in seven separate steps. In one version, all steps were done in 16-bit RGB; in another, 8-bit RGB; and a third and fourth were done in 16-bit and 8-bit respectively, but after each of the seven steps, the file was converted into LAB and back into RGB. On the opposite and next two pages, the four versions are shown in random order. Can you tell which is which?*


A


B


C


D

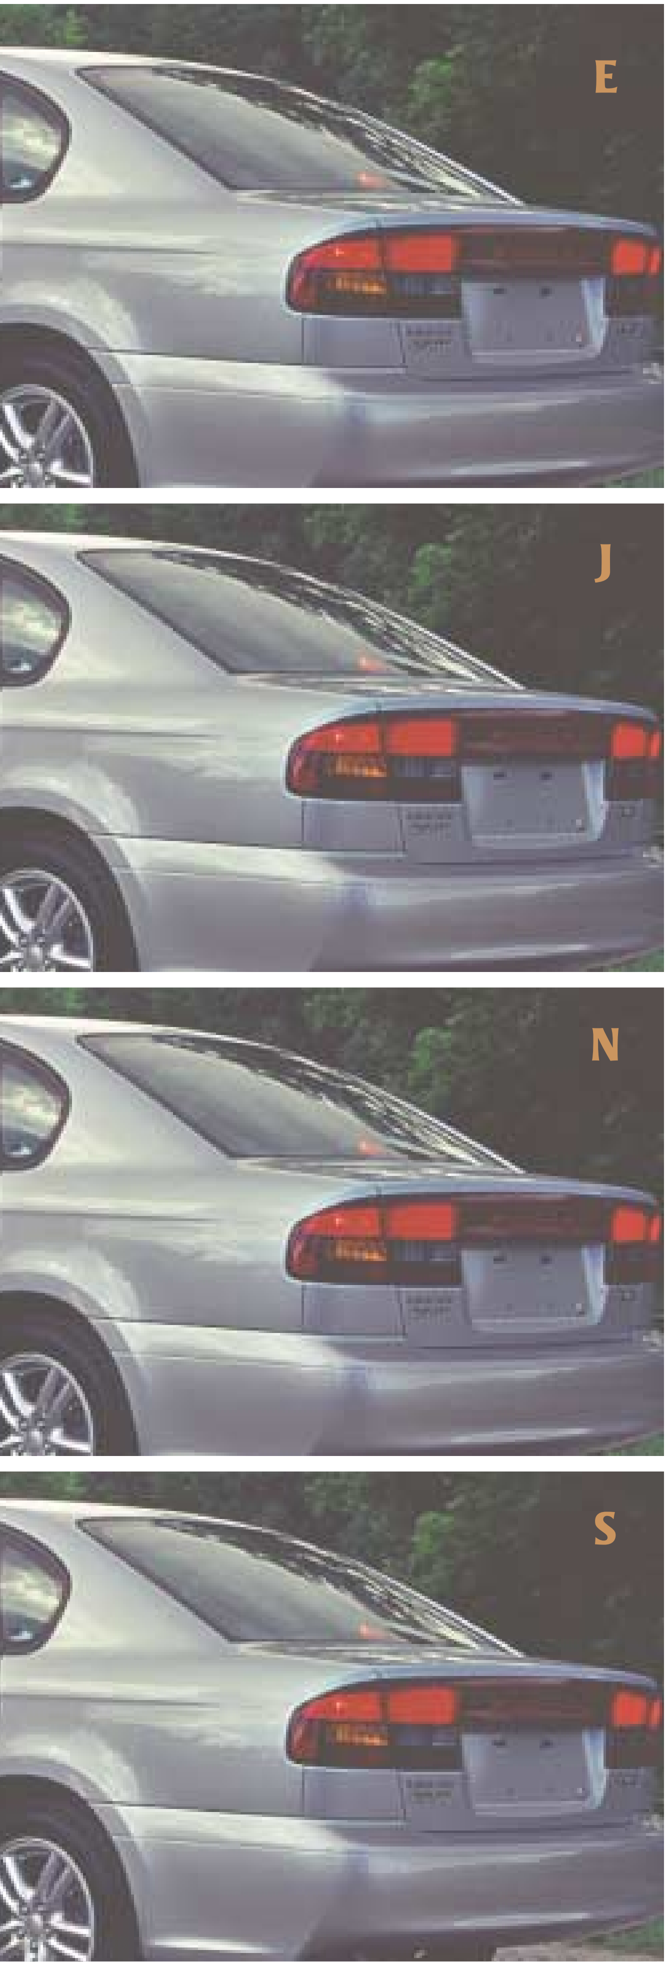

***Figure 6.10** Views of the four versions of Figure 6.9 at various sizes. Left to right, the magnifications are 200%, 250%, 400%, and 500% (showing the green channel only).*


G


abrikation
H


L


abrikation
M


Q


abrikation
R


U


abrikation
V

a third of the colors. We've lost nine-tenths of them! Surely, it is madness to suggest that converting into LAB is safe!

A great theory, seemingly irrefutable. And yet there's Figure 6.8, big as life, laughing at us, demonstrating that there's no loss at all, not even after 25 conversions to and from LAB.

When the Impossible Happens

This book assumes that our RGB is the variant known as sRGB, a choice of convenience, not an endorsement. Many professional photographers believe that sRGB is unduly limiting. Its definitions of the primary colors are relatively dull. Those who subscribe to this criticism generally prefer the definition Adobe RGB, which permits more brilliant colors at the expense of some subtlety. A few feel that Adobe RGB isn't wide-gamut enough and use an even more brilliant definition.

An Adobe RGB user who wishes to work on a file that was prepared for sRGB has to convert it, using Image: Mode>Convert to Profile, just as we LAB users need to convert out of whatever our own RGB is to do our thing.

So, here's the challenge. Suppose Figure 6.8 was prepared not by converting sRGB to LAB to sRGB 25 times in a row, but rather by converting sRGB to Adobe RGB to sRGB 25 times. How much closer to the original version would it be than the image with multiple LAB conversions?

Adobe RGB is certainly a much closer relative to sRGB than LAB is. It does waste a certain amount of real estate on colors that don't exist in sRGB, but still, if there are 10 million distinct colors in the sRGB version of Figure 6.8, I'd have to suspect that there would be 9 million in an Adobe RGB version. So it *has* to be less damaging to convert to Adobe RGB than to LAB—right?

Wrong.

If you do this test—and I have—a most perplexing thing occurs. The multiple-LAB conversion is closer to the original than the multiple-Adobe RGB version is. None of the three versions can be easily told apart, at least I can't, but we can apply statistical measures to verify that the impossible is indeed true.

This is becoming surreal, and we haven't even hit the clincher yet. Create a new RGB file. Choose a couple of unlike colors for Foreground and Background Colors, activate the gradient tool, and create a vignette. Make a copy of the file. Convert it to LAB, and then back to RGB. Hideous! Banding in several areas.

Seems fairly conclusive—but then again, there's Figure 6.9. Tortured almost beyond belief, converted again and again, when it's a real picture and not a computer-generated gradient, all four versions are so close as to be indistinguishable for any practical purpose.

Every logical way of looking at it suggests that the LAB versions have to be much worse than the RGB originals. But they aren't. Therefore, something about the reasoning is incorrect; it only remains to figure out what part.

Faced with things I don't understand, I find it useful to curse at the monitor. If that fails to resolve the problem, Armagnac, or on extremely rare occasions a cigar, may make an appearance to help the thought process along. I forget how much of this was necessary ten years ago, when I first tried to figure out how there could possibly not be a visible loss when going to LAB. Anyhow, there are two basic answers:

1. In mathematics, the symbols $+$ and $\times$ do not mean the same thing.

2. In a photograph, the blood does not curdle at the thought of altering a single comma.

Of Salaries and Pixels

Numbers make excellent servants, poor masters. An overweening and unwarranted belief in the power of their precision has been the hallmark of those who cry data loss every time there's a minor move in the image.

John Jones makes $50,000 per year. How much does he make per week?

A computer programmer would answer, is it a leap year, or not? A statistician would answer, about $1,000.

Someone who thinks that converting to LAB is damaging would answer, $958.9041095890.

We need clarification. Does what we have been told really, literally mean that he makes $50,000.00, not a penny more or less, in the course of one non-leap year? Or is $50,000 merely shorthand for somewhere between $45,000 and $55,000? Or between $49,000 and $51,000?

Knowing as little as we do, the statistician's answer is correct. It really sounds like $50,000 is some kind of rough estimate. Any answer more precise than $1,000 a week makes an unwarranted assumption about the reliability of the data. $958.9041095890 sounds ever so much more authoritative, and so impresses some Photoshop authorities that they call the $1,000 answer "quantization error." In fact, from the statistical point of view, it's far more accurate than making unwarranted assumptions about how many significant digits we start with. Anything other than the first digit after the dollar sign is a random number, for all we know now.

The same analysis applies to digital images. Cameras and scanners do not return perfect data. We should have more confidence in the reliability of midtone captures than those of extreme lights and darks; in less saturated rather than brilliant colors; in the green channel rather than either of the other two. But in any case, the very act of capturing the image has introduced unwanted variation.

Even if the data is very good (and how would you prove that it is?), it can never be fully reliable. Suppose you own the finest camera or scanner in the world. You claim that it's capable of resolving 1,000 different levels of gray, and that a certain pixel measures 437, and that's the correct value, period, amen.

The response is, how can you be so sure? The device is actually trying to juggle a lot more than 1,000 values, and it's doing some rounding. What 437 really tells us is that the pixel measures somewhere between 436.51 and 437.49. But is the device actually that good? Because if it's off by as much as .02, it could conceivably be reporting something as 437 that actually should have been rounded to 436 or 438. And if you say yes, the device is really that good, I'll ask whether it's good enough to know the difference between 436.4999999, which rounds to 436, and 436.5000001, which should be reported as 437; and I'll keep adding decimal places until you give up and admit that it's theoretically possible that 437 is not technically the correct value.

Back in the real world, the results are reported on a scale of 0 to 255, or 256 values in all. We use this scale because 256 happens to be the number of possibilities that can be described with eight bits of computer data. That is, a single bit is either on or off, yes or no, 0 or 1. Two bits give us four possibilities: 00, 01, 10, and 11. Three bits permit eight, since any of the above four two-digit numbers could be followed by either a 0 or a 1. Each time a new bit is added, the possibilities redouble. Four bits allows 16, five 32, six 64, seven 128, and eight 256.

All modern capture devices nominally use more bits. They may think they're getting 1,024 values, or even 4,096. The question is whether the numbers are particularly accurate. Some people are so buffaloed by arithmetic and so in awe of any kind of measurement by machine that they forget to ask it.

No computer program can verify whether a given pixel is correct. We have only our gut feelings as to how accurate the capture is. My own is, I don't think any devices can make accurate real-world captures in more than thousandth-part increments, and that's only under the very best conditions at certain levels of lightness. If it's a digital capture taken in relatively dark conditions, I don't think the camera gets even close to 256 accurate values. Under better conditions, I think most cameras record accurately to within a level of the ideal, particularly in the critical green channel. That is, if the camera records 128^G, I doubt, but don't rule out the possibility, that 126^G or 130^G would have been more accurate. A difference of one level, that's another story.

Of Translations and Transfers

Cyrano never exactly said any of the things quoted so far. He couldn't have—he was speaking French. What you've read is a translation, a restatement of what he said, just as an LAB file is a translation, a restatement, of the RGB one.

Cyrano says, *"Mon sang se coagule."* The first two words can be matched exactly in English: *my blood*. The second two are harder. The cognate *coagulates itself* doesn't carry the proper sense. I vote for *curdles,* but would accept *runs cold* or *congeals.* The three choices are not identical, but equivalent for all practical purposes.

Now, suppose someone without access to the original text retranslates *my blood curdles* back into French. The first two words would be restored to the original *mon sang,* for sure. There are several possibilities for the third—all just as good as the original to everyone except Cyrano, whose blood curdled at the thought of changing a single comma.

If we retranslated the entire play, from French to English and back again, each phrase would compare to the original in one of the following ways:

- Identical.
- Worse.
- Equivalent.
- Better.

The phrase we've been discussing would be partially identical, partially equivalent. The chances are that much of the rest of the play would be worse, because there really is loss in certain translations. (On the other hand, a book of the collected speeches of George W. Bush might well read better if it were translated from English to Russian and back again.)

The point is, identical is not only unlikely, but it isn't even desirable, provided the retranslation is equivalent or better. And so it is with color files. Around two-thirds of the pixels in the version of Figure 6.8 that was translated 25 times in and out of LAB are identical to the original. The remaining third could conceivably be worse than the original—but conceivably some are just as good, and others may even be better. We just don't know. Unless the pixels fall outside of our range of uncertainty, which is always at least one level, to insist that they match the original exactly is to go to the last hundred millionth of a cent when your margin of error is a thousand dollars; to announce that your blood curdles at the thought of changing a single comma.

And that's the fundamental difference between photographs and computer-generated art, one that renders the test of a gradient being converted to LAB pointless. In gradients, the change of any comma would indeed be blood-curdling.

A Photoshop value of, say, 127, is an approximation, if it's a photograph. Maybe if this were a perfect world, with infinitely precise cameras, its real value would be 126.67289, which rounds to 127 but can go to 126 instead without any worries. In our world, the range is considerably wider, so 126 might well be not just equivalent to but better than 127.

But if it's a gradient, then the correct value in a perfect world is 127.00000. Any change is by definition wrong. If the retranslation doesn't come back identical, then it's worse. Better and equivalent are no longer possibilities.

If a whole row of pixels in a gradient jumps by two levels rather than one, it's visible, even though in a normal photograph, a two-level variation can be seen by the naked eye about as frequently as Halley's Comet.

Theorizing that converting to LAB causes damage and attempting to prove it by converting a gradient is circular reasoning. It assumes that a single value is uniquely correct, tests a method that is sure to change it, and then concludes that the method is inaccurate. It is a statement that *my blood curdles* is the one and only correct way to translate Cyrano's phrase and that any other phrase is data loss.

Incidentally, the problem of gradients in conversion is by no means limited to LAB. Many people face needless frustration when they prepare gradients (particularly blue ones) in RGB for files that are eventually going to CMYK. This begs for banding or other evil consequences. Gradients should be created in the same colorspace as the output device—in this case, CMYK.

The Most Useful Statistic

An architect planning to build something in a strange city needs to know what temperatures are likely to be encountered, so that appropriate heating and air conditioning systems can be ordered. The information that the average noontime temperature in my New Jersey home town is around 53 degrees Fahrenheit would not be enough for that purpose. That average temperature is similar to that of Kansas City, Missouri, which, not being close to any ocean, has more extreme heat and cold. Yet summer days where I live are frequently hotter than in San Juan, Puerto Rico, which has a much higher average temperature overall. As a matter of fact, Fairbanks, Alaska, is sometimes as hot as San Juan in the summertime.

The average temperature is not as important as how much it fluctuates. And the architect would need something better than all the temperature records of the last few years. For example, I don't recall noontime temperatures of higher than 95° in the last five years. However, around 15 years ago, it hit a ghastly 106° and stayed there for several days.

The supremely important statistic known as *standard deviation* would have informed the architect that such a heat wave was possible, even if the only records available were for the last two years. The concept applies whenever there are many data points clustered more or less uniformly around a mean value, as the weather is. If the mean is 53°, we're equally likely to find 63° as 43°; less likely but still equally likely to find 73° as 33°, and so on.

I haven't gathered the data or done the arithmetic, but I'm going to estimate that the standard deviation in my home town is around 14°, and the cities mentioned above as follows: Kansas City, 17°; San Juan, 5°; Fairbanks, 24°.

High standard deviations are generally bad things. If you had to choose which of these cities to live in based solely on their climates, you would certainly choose them in the order of lowest standard deviation—even if you don't know precisely what standard deviation means or how it is computed.

In fact, almost everything having to do with process control in the graphic arts amounts to a struggle to reduce the standard deviation, because variation is bad and variation is what the standard deviation measures. For example, the printer of this book, whose presses are run by mortals, sometimes prints jobs lighter or darker than his average. I am hoping very hard that his standard deviation is low and that this book will fall close to the mean when printed.

Once enough data exists for a standard deviation to be computed, it can be used to predict the likelihood of various events. For example, the variation of noontime temperatures over the period of a year is likely to be slightly less than six times the standard deviation, meaning in my case that the hottest day is around 80° hotter than the coolest. Fairbanks, I am given to understand, has the highest standard deviation of any major city—around 140° difference between the coldest and hottest days. I can also learn from the standard deviation that my town does occasionally have days in the 90s; that something on the order of the 106° heat wave is apt to occur every 20 years or so, and that a reading of 115° would indicate that either the thermometer is broken or the weather recording station is on fire.

The Odds Are Against It

As you may have conjectured, standard deviation can also be part of image analysis. Like the histogram, I consider it worthless as an aid to image manipulation. Neither can tell us about the visual quality of an image as accurately as our own eyes do.

Both are, however, sometimes helpful in trying to figure out why something is happening that we don't understand, like, for example, why converting to LAB is safe when logic seems to dictate otherwise.

To learn how close the two halves of Figure 6.8 are, I applied one to the other in Difference mode. This blend, which can be done in several ways, creates a black file, except in pixels where the two images aren't identical.

For an RGB image, Photoshop offers six different sets of statistics to accompany the histogram, in locations that vary with the version of Photoshop. The most important stats are those for the green channel and for luminosity, which is a weighted average of red, green, and blue. Photoshop reports that in the green channel the mean variation between the original of Figure 6.8 and the version that went in and out of LAB 25 times is .15 levels and the standard deviation .36; in luminosity the numbers are .10 and .30.

Let me offer, er, a translation. The numbers indicate that the variation is approximately equivalent in impact to the soft noise or dither that Photoshop by default inserts every time an image is converted from one colorspace to another. If you didn't know that Photoshop does so, you're not alone—it's undetectable, useful, and harmless. (If you're going to be converting files 50 times, though, you should turn it off, as I did for these tests.)

Further, if these numbers are correct, around 80 percent of the pixels in the two green channels are identical, and essentially all others are one level apart. Variation of two or more levels would occur, if at all, less than one time in every 5,000 pixels.

Also, remember that we never see individual pixels except on the monitor. When the image is printed, there's always an averaging process to convert the original pixels into the form that the output device requires. This is true regardless of how the image gets printed. In the case of this book, the press requires halftone dots, tiny blobs of cyan, magenta, yellow, and black ink. Each dot is calculated by averaging, usually, the values of three or four pixels. Take a loupe to either half of Figure 6.8, and if you have a few weeks to spare you'll be able to count some 2.7 million halftone dots, averaged down from around 7.5 million pixels in the CMYK Photoshop file.

What would it take for us to notice roughness, any degradation in quality? I'd say, a dot, not a pixel, that varied from its proper value by at least two percentage points. Although printing dots are usually referred to in terms of percentages, they in fact are constructed on a 256-level scale, just as pixels are. Two percentage points equals five levels.

But let's be ultra-conservative and say that a dot might be detected if it were only two *levels* larger or smaller than it should be. Being that it's camouflaged by three other correct dots of different colors that are intersecting with it to some extent, it would be almost impossible to see, but let's theorize that we are going to edit the file so drastically that the difference might show up later.

The Torture Test: 25 Times Back and Forth

(All variations from original are expressed in Cyrano Units; lower is better)

sRGB to	Red	Green	Blue	Lum
LAB	1.62	1.16	1.71	0.95
ColorMatch RGB	1.68	0.65	0.49	0.09
Adobe RGB	3.96	0.88	1.99	3.18
Wide Gamut RGB	8.45	12.80	3.66	9.12
LAB (w/dither, 1 conversion)	2.62	2.13	2.82	1.38
LAB (w/dither, 25 conversions)	9.50	7.67	10.04	3.37

***Figure 6.11** The original of Figure 6.8, in addition to the 25 conversions to LAB and back, went through five other sets of conversions, in each case but one being converted in and out of the destination space 25 times. Variation from the original is expressed in terms of "Cyrano Units" as defined in the text. All files except the final two lines were converted with dither disabled.*

Now, let's try to figure the odds of this rogue dot ever showing its long nose. I will skip over considerable arithmetic here in favor of approximations. The precise odds can't be calculated because of irritating complications such as the fact that the data isn't truly randomly distributed around the mean, and that the presence of one incorrect pixel sharply increases the odds that one of its neighbors will be also. So, I will use the traditional prepress technique of the fudge factor. I will assume that one in every 300 pixels varies from the original by two levels.

Cutting to the swordfight, the only sure way to get a two-level variation in the dot is to have four pixels forming a square, all being either two or more levels lighter or two or more levels darker than the mean.

The odds against either event occurring are approximately 65 trillion to one.

A much more reasonable scenario would be to take a cube of 9 pixels. If any four of them were rogues, *and* if there weren't any rogues of the opposite persuasion to cancel their effect out, then it's fairly likely that a two-level variation might occur in a certain halftone dot. The odds against this happening are a much more modest billion and a half to one against.

And remember, even if it happened, you almost certainly wouldn't notice, particularly if it happened in the red or blue channel. And it assumes far too many rogue pixels. In fact, it assumes that the standard deviation is twice as high—like it would have been, if these repeated conversions had been into Adobe RGB rather than LAB.

The Tale of the Tape

Given these tiny variations, the two halves of Figure 6.8 are identical for all practical purposes, and to the extent they vary, nobody can prove which one is better. But what level of variation might actually cause a problem?

Disgracefully little research has been done into the vital issue of the handling of images that have been converted 50 times, an omission I propose to remedy here by offering the following formula: three times the standard deviation plus half the mean. As the inventor of this new standard, I get to name it; and in honor of this chapter, it is hereby dubbed the Cyrano Unit.

If the reconverted version varies from the original by less than two Cyranos, the files run a statistical dead heat. Between two and three, there's a case to be made that the original is better, but it won't make any difference. At values of three in the green or luminosity sets, somewhat higher in the red and much higher in the blue, it becomes conceivable that problems may develop later; at four it becomes probable; and at five it's a definite pain. (The stats labeled RGB and Color are not important.)

Not content with converting a file to LAB and back 25 times, I tried the experiment with four other settings. The results are summarized—in Cyranos—in Figure 6.11, which brings us back to the flabbergasting observation that we get closer to the original if we convert to LAB and back than to Adobe RGB and back.

I also tested conversions into ColorMatch RGB, which covers a smaller gamut, and into Wide Gamut RGB, which, as the name suggests, is huge, as big as LAB itself. Unsurprisingly, the smaller the gamut being tested, the closer the post-conversion file was to the original. The ColorMatch RGB version was slightly closer than its LAB counterpart. In its green channel, 19 of 20 pixels were identical to the original.

The Adobe RGB version isn't quite as close, particularly in the red channel. In fact, let me earn the price of the book by offering an important tip: before commencing work on a file, don't convert it to Adobe RGB 25 times and back. With LAB and ColorMatch RGB, go for it, if you've got a lot of time on your hands. But not Adobe RGB.

The Wide Gamut RGB version is the worst of the lot by a large margin. A lot of people would have thought that the LAB file would have had the same kind of grim numbers. It didn't happen. Let's discuss why.

The Plus Sign and the Times Sign

These three RGB definitions are more alike than different. The red channel in one is very similar to that of another, except the narrower-gamut one will show more contrast. It has to, because it needs to have a lot of action at the extremes if it hopes to match the brilliance of colors that the wider-gamut RGB produces routinely. Therefore, any object occupies slightly more space in the narrower-gamut RGB. It may take 11 levels to portray something for which the wider-gamut one only needs 10, which becomes awkward when converting between the two. With only ten steps in the original, we can't go from 1 to 11 in ten steps of 1.1, as we'd like to. We have to take single steps—except that somewhere along the line there will be one dubious double step. And if we go from 11 to 10 steps, we can't spread the damage among all eleven: ten will get their normal variation and one will vanish. That's potentially the birth of a rogue pixel.

This effect, where the very similarity of the file structure hampers the conversion, is absent in LAB. There's a mild correspondence between the L and every RGB channel, but it's heavily disguised by the impact of the AB, which have as much in common with the RGB channels as the poetry of Edmond Rostand has with that of Eminem. The A and B have disturbingly long ranges between steps, but since the steps don't correspond to anything in RGB, the effect is distributed more uniformly.

That the RGB channels are intact also explains the mystery of how LAB appears to dump nine-tenths of the possible colors without destroying the image. When the RGB channels are sound, it doesn't matter how many colors are missing, because they'll show up sooner or later. That the LAB file doesn't have millions of distinct colors merely means that certain combinations of RGB values are impossible—temporarily. If you have 150^R160^G, then perhaps 170^B may not be a possible companion; you'd have to go to 171^B or 169^B. The value 170^B exists in the file, just not in conjunction with the other two.

If that's an unsatisfactory state of affairs, there are lots of ways to restore millions of colors very quickly, such as:

- Gaussian blur at .1 pixel radius.
- Rotate the image 5 degrees and then rotate it back.
- Upsize the image by one pixel and then downsize it again.
- Make a copy of the file, convert it to LAB, reconvert it to RGB, and apply it to the original at 50% opacity.

In fact, just about any move you make to a single channel will create tens of thousands, if not millions, more color possibilities.

Fortunately, you can save all of the above quackery for the next time some nincompoop complains that your histograms look too ugly. A file that merely is missing a lot of color combinations is no cause for worry. There may have been less than a million distinct colors in the LAB version of Figure 6.8, but there are millions and millions now that ink has hit paper. There would have been millions had we output it on a desktop printer, and there are even millions when we open the file and look at it on screen.

No output process uses the red channel as is. Even a desktop printer that appears to be taking RGB data is converting the incoming file to CMYK. And the cyan channel, although similar to the red, has been heavily munged. Its center has been lightened, and to some extent it's been blended with the previous blue channel. The previous limits on combinations with other channels no longer apply, and the millions of colors are shown to be, like Cyrano orchestrating the courtship of Roxane from underneath her balcony, there all the time, temporarily hidden in the shadows.

A Bit About Bits

The question of whether converting colorspaces causes harm is closely related to one mentioned during the discussion of Figure 6.9: whether there might ever be an advantage in correcting files in Mode: 16 Bits/Channel as opposed to the more conventional 8 Bits/Channel. That subject is academic for us, because all techniques in the book work either way. However, the debate does offer some constructive lessons.

16-bit files are twice as large. They contain 65,536 tonal levels per channel rather than 256. It is logical to think that such a file might be more forgiving of major tonal changes than an 8-bit file would be, particularly if there are several such changes in succession. So, a number of Photoshop authorities, some politely and some imperiously, have suggested that at least major editing should be done in 16-bit mode, without ever showing a single example that suggested there was any merit in doing so. In one notorious case, a prominent photographer announced that anyone who wasn't editing in 16-bit mode was a "recreational user" of Photoshop, rather than a professional.

It sounds sensible, just as the theory that converting to LAB is damaging sounds sensible. And the result is just the same: in practice, the theory doesn't work. On images containing computer-generated gradients, yes. But on color photographs, no. Consider Figure 6.9, which has been massively corrected, seven different times. Yet the version done entirely in 8-bit *and* also converted seven times during the process to and from LAB is just as good as the one done in 16-bit all the way with no unnecessary conversions.

In the last three years, around a dozen people, including me, have made serious efforts to find anything to support the proposition that 16-bit editing might be better under any circumstances. By that, I mean any unretouched color photograph that might possibly be used in the real world, and any sequence of attempts to improve the image, however farfetched, where editing in 16-bit creates a better result than 8-bit. Images have been tortured beyond belief. Nobody has found any quality gain at all.

Neither, of course, has anybody shown that there's any loss by doing so. So, if you have the disk space to spare, and feel like wearing belt and suspenders, go ahead.

As for LAB specifically, a number of people whose opinion I respect think that because LAB is so huge, editing there in 16-bit might make more sense than it would in RGB. With the possible exceptions of nearly neutral images that have been heavily altered in the L channel and of once-in-a-lifetime images where you decide to unsharp mask the B channel, it isn't true. Counterintuitive as it sounds, for a lot of the

8- and 16-Bit: An Exception

Many digital cameras offer the option of producing an 8-bit or 16-bit file, although most consumer-level digicams output in 8-bit only. If you have the option, and are planning to work on the image afterward to any extent, you should open in 16-bit and convert to 8-bit in Photoshop at your convenience, whether you use LAB or not.

In response to a standing challenge, many users have sent me files, together with the actions that were taken, seeking to show that 16-bit corrections were better. There have only been two cases where the claim held up. In each, the user had output both types of file using his camera software and then had applied massive, but identical, corrections to both.

The first user supplied images on a gradated gray background, which posterized badly in the 8-bit version. The second had work featuring dark, rich colors: burgundies and greens. After huge corrections, which included attempting to work on a raw 1.0-gamma file in Adobe RGB, the 8-bit version exhibited ugly dark noise in these areas that wasn't found in the 16-bit.

When these corrections were repeated on 8-bit files that had been generated by converting the 16-bit original to 8-bit *in Photoshop*, however, the results were every bit as good as the ones done in 16-bit all the way.

One, with the 8-bit file prepared by Canon's Digital Photo Professional 1.5, arrived while I was drafting this chapter. I compared it to an 8-bit version generated in Photoshop from the 16-bit file prepared by the Canon software. After verifying that the Photoshop version was extremely close to the 16-bit original, I compared the two 8-bit versions before they were corrected. The variation was a ghastly 7.5 Cyrano Units in the green channel—more than enough to cause problems if the image is edited extensively.

reasons discussed earlier, RGB would need the extra bits more than LAB does. Theoretically only, I hasten to say: in real life neither one needs it at all. However, if we were forced to work in 6-bit—only 64 levels per channel—6-bit LAB would have a lot of advantages over 6-bit RGB. And with that, I think we should stop discussing 6-bit Photoshop and files that are converted 25 times back and forth, and how many angels can dance on the head of a pin, and whether Photoshop books should be written in blank verse.

Ton Nom Est Dans Mon Coeur

Lawyers would describe this half of the chapter as being an attempt to *prove a negative,* an exercise in futility. That is, there's no way of proving absolutely that converting to LAB *never* damages a file or that editing in 16-bit *never* gives better results, or that wearing garlic around the neck while at the computer never prevents shadows from plugging.

Fortunately, it's not up to me. Whoever is advocating doing something inconvenient is responsible for demonstrating why it's necessary. The purpose of this book is to suggest that you should learn an alternative way of working with color, which is certainly inconvenient even before we consider having to convert each file to LAB and then back. Therefore, I'm the one who has to make the case that you will get better results that way. So, where possible, I compare the LAB method to the RGB equivalent—if one exists.

CYRANO. I am so in thrall to your hair
That, like one who stares at the sun,
And thereafter imagines shades of vermilion
  on everything,
When I leave your fire, your luminescence,
My whole life develops a blond cast.
ROXANE. [with trembling voice]
  Oui, c'est bien l'amour…

This is not very good color theory. One who stares at a colored object sees the *complementary* color thereafter, so if Roxane's hair was all that compelling, Cyrano would have been looking at the world through B–negative glasses once he turned away from her—a blue cast, not a blond one. Nevertheless, as Roxane remarked, it is truly love, and mere matters of factual inaccuracy have seldom troubled suitors.

Cyrano was not a Photoshop user himself, but had he been, he would have been a devotee of LAB, the space that liberates color and allows the imagination to put blond casts where it will. If someone suggests you should give that up on a theory, that person should be wielding a picture, not statistics; a sword, not a histogram.

THE FIRST CADET. [shrugging his shoulders]
  Always the sharp, the pointed word.
CYRANO. Indeed, the word is the point.
And when I die, I should like it to be
In the evening, under a rosy sky
As I speak sharp words in favor of beauty.
Ah! To be struck down by the noble arm
Of a man worthy to be my enemy,
On the field of glory, and not in a sickbed,
A point in my heart and a point on my lips.

The Bottom Line

Most problems with LAB derive from using it on images for which it is not suited, from not appreciating that it can produce colors that are wildly out of any known gamut, and particularly, from trying too hard to get a perfect result in the L channel.

LAB should be a Photoshop-only tool. Other programs generally don't support it. Sending an LAB file to an output device is a form of Russian roulette.

There is no problem in converting files from RGB to LAB and back, unless the file contains a computer-generated graphic such as a gradient. Such graphics should always be made in the final output space—CMYK if the job is to be printed.

The L channel sometimes serves as a better black and white version of a color file than a direct conversion to grayscale would be. However, the two are close relatives. Blending channels gets superior results.

Notes & Credits

Chapter Six

Beta reader André Dumas offered improved versions of some of the lines in my translations of Rostand's *Cyrano de Bergerac*.

The picture of the horsewoman in Figure 16.8 is by Jim Bean. Versions B, C, G, and H are the ones that went back and forth between LAB and RGB 25 times.

In the original double-picture collage of Figure 16.9, the German marketplace is from Kodak, photo by Alfons Rudolph. The car was photographed for advertising use by Aldas Minkevicius.

As for identifying the various pieces of Figures 16.9 and 16.10:

- The politically correct method, 16-bit RGB all the way, is in versions C, S, F, Q, and R.
- Doing the same thing in 8-bit all the way were versions B, N, P, U, and H.
- The method of converting to LAB and back to RGB after each one of the seven correction steps, while remaining in 16-bit at all times, is shown in versions D, E, K, G, and V.
- The least palatable method, working in 8-bit throughout and converting back and forth from LAB after each correction step, is versions A, J, T, L, and M.

If you're curious as to how well somebody else might do at picking out which was which, I can be the guinea pig. After the text of the book was finalized, I received final contract proofs from the printer. Since I had long forgotten in which order the versions appeared, I took 20 minutes and carefully examined each variant under controlled lighting conditions. I permitted myself reading glasses, but not a loupe.

In the horsewoman shot, I could see no difference at all either at 100% or 400%. In the two intermediate sizes I thought I saw enough to take a guess at which was which, and was correct in both.

I was less successful with the collage. Of my 20 answers, only three were right. Considering that someone who answered the questions on the basis of coin flips would average five correct answers, this was not particularly impressive evidence of a detectable difference between the four variants.

I was most accurate where it could be expected—at the highest magnification, looking at a single channel. In that set, I correctly identified both 8-bit versions. However, I mixed up the 16-bit versions, declaring that the one converted to and from LAB at seven different points in the process was actually the one that had never been converted at all. With respect to the four sets at lower magnifications, my responses appeared to be completely random—no discernible relation to the true order of the images.