

CFS-276

Complete Color Integrity

David Dunthorn
C F Systems
February 24, 2009



Complete Color Integrity

CFS-276

February 24, 2009

This document is a continuation of and, in an important way, a completion of a series of three documents from early 2004 investigating some aspects of digital photography. These companion documents are CFS-242, *Film Gamma Versus Video Gamma*, CFS-243, *Maintaining Color Integrity in Digital Photography*, and CFS-244, *Negative to Positive*. In general this document is much less mathematical, less formal in structure, but much more conclusive.

This document is viewable only. We believe this will be adequate for people who do not intend to study it. Please contact us through our web site if you need a printable version. We are aware that the no-print can be defeated, and of course you can print and assemble much the same information from the web-html version but again we ask that you contact us instead. We would like to know if and how people are finding this document useful, and this seems one of the few ways we have to encourage feedback.

David Dunthorn

www.c-f-systems.com

Much of the material in this document and the analyses are original with us.

This document is

Copyright © 2009 by C F Systems

If you plan to use our original material or analyses in any published form, please contact us at www.c-f-systems.com for terms and conditions. Before-the-fact terms typically require no more than appropriate acknowledgement.

Go To Perspectives and Comments (Table of Contents)

Introduction

Several years ago in these web pages I began to explore the matter of *color integrity* in digital imaging. This came about because like most people, when I started into digital photography I was hypnotized by the marvelous things that are possible – even easy. Then the marvel began to wear off when I started to realize that something was *wrong* with the color in many digital images. Not *all* digital images, but in so many of them; not only in snapshots, but in professional photographs, even in the highest quality magazines. Worst of all, I found this disturbing color effect in some of my own photos.

I did have success in finding the main cause of this lack of *color integrity* but there was a serious problem in that my early work was mathematical and difficult for most people to appreciate. Very recently I have had a real breakthrough that has so greatly expanded my understanding of *color integrity* that now I can explain it in simple terms and without using mathematics. In fact, after seeing this explanation you may find it difficult to believe that it is not already a core element of digital imaging. Yet, clearly it is not. The tools for maintaining color integrity prove to be very basic and very simple indeed, yet you will not readily find them in Photoshop. As this is being written popular tutorials and common practice both emphasize tools that actually *destroy* color integrity instead. That is why the problem is so pervasive even in professional work.

I feel very fortunate to have made this breakthrough. It is a real rarity to have several lines of study of a complex topic converge so beautifully into such a simple explanation. Join us on our *Complete Color Integrity* page to learn the full story.

For the past several months I have told some of you that this document would be published on the web site "soon." I apologize that it has taken this long. The new concepts are so fundamental that they touch upon many areas of digital imaging in significant ways and it has taken a long time to tie the loose ends satisfactorily.

Complete Color Integrity

We begin by explaining *color integrity* in simple terms. Following that there are several routes to take toward a more thorough understanding of *color integrity*, and the reader can follow whichever routes are familiar and comfortable. I need to be clear from the start, however, that I have nothing against the creative, artistic use of color in photography. However, I do believe that it is important to start with an image that has color integrity before getting creative with the color. Otherwise a lot of creative energy is wasted in trying to compensate for what the image lacks.

The concept of *color integrity* deals with the behavior of color as it appears when in bright areas of a scene and when in darker areas of the same scene. Imagine a colored patch placed in a dark area of a scene. If I move this patch to a light area of the same scene the patch becomes lighter but it does not change its apparent color. (Of course this model assumes the light source itself is not different in the light and dark areas of the scene, but cases where the light does change are easily built upon this framework.) If we place this colored patch in an area where it is shaded at one side and the light is unveiled gradually to full brightness the other side, again the color of the gradient does not change as it becomes lighter. This is how the eye expects colors to behave, whether looking at the folds of a drape, the leaves of a tree, or the contours of a face. When the shading of colors in an image behaves in this natural manner we call that *color integrity*. When the colors in an image change unnaturally going from dark to light the eye becomes confused and the lack of *color integrity* can produce an unsettling effect.

In photography there are two distinct steps involved in producing images with this *color integrity*. First we must produce a digital image that corresponds to the colors and variations of color in the original scene with acceptable accuracy. This step of creating the initial accurate image having *color integrity* is called *calibration*. We find it best to use a calibration system based on a stepped grayscale. In digital photography gray is composed of equal parts of Red, Green, and Blue so calibrating with a stepped grayscale exercises all three color channels identically and because the grayscale steps through the range from black to white the calibration covers the entire response of the digital camera sensors (or the film) over the complete range of all three colors, red, green, and blue, that the digital camera or film sees. Using a correct grayscale calibration will produce an image file that starts out with color integrity. We have a web page on grayscale calibration: <http://www.c-f-systems.com/DunthornCalibration.html> and the **ColorNeg** and **ColorPos** plug-ins <http://www.c-f-systems.com/Plug-ins.html> include several methods of grayscale calibration. For those who are curious about the choice of grayscale calibration and whether it is really adequate – and preferable – to using colored patches, in addition to the web links, **Comments About Calibrating Digital Images** in this document has more detail on calibration and why many of the current methods of calibrating digital cameras and film, particularly "profiling," often lose color integrity right from the start.

The second step in producing images with *color integrity* is in making sure that adjustments, especially initial major adjustments, to this initial calibrated image can be made while maintaining its *color integrity*. Here we are particularly interested in adjustments which

lighten dark areas of the image or darken light areas of the image and also in adjustments to produce color balance, both general and in localized areas of the image. This is where my new discoveries come into play.

To explore how we do this, let us take a calibrated image fresh from a digital camera in which the colors and their shadings do behave as described above – that is, a starting image that has *color integrity*. We want to make the image darker, but we want to maintain *color integrity* as we do so. Above we moved an imaginary colored patch into different lighting conditions to see how it behaved. Now it is helpful to think of each tiny bit of color in our image is such a "patch." Whatever happened to the color of our reference patch when we made it darker by moving it into areas with dimmer lighting, that same thing has to happen uniformly to all the colored "patches" in our image. As we prove elsewhere on these pages, doing this is very simple – we just add black. To darken the "patches" of color that make up the image we must add black to each color in the same way that an artist-painter might add black paint to darken a colored paint on the palette. To darken the entire image uniformly we need to mix in the same amount of black to the color of each color patch in the image, resulting in exactly the same effect as if we had dimmed the light source. This idea of "adding black" to darken an image probably seems obvious now that it has been explained and is not at all difficult to do to a digital image. But you won't find anything similar to "add black" in the Photoshop plethora of commands nor is it particularly easy to ferret out exactly how you might achieve that result in Photoshop. Even then the tool may work in some cases and not in others. Apparently this "obvious" approach really is not all that obvious!

Suppose now, instead of darkening the image, we wish to make it lighter. Again, it is very easy to achieve that. We simply reverse the above and *remove* black instead of *adding* it. To uniformly lighten the image we remove the same amount of "black paint" from the color of each patch of color in the image. Again, this is easy to do for a digital image, just the reverse of adding black. A painter would find it difficult to remove black paint from an already mixed color!

"Adding black" also plays a less obvious part in properly dealing with color. The terms "white balance" and "color balance" are commonly used in dealing with the fact that the human eye automatically "accommodates" so that it sees as white – or nearly so – whatever light source predominates in any scene under view. Neither photographic film nor the sensors in digital cameras do this and so the image nearly always must be adjusted – white balanced or color balanced – to achieve the same effect as the eye's accommodation. To understand what is happening here we need to look at the scene in the way the camera does, with three image sensors, one for Red light, one for Green light, and one for Blue light. When this set of sensors produces an image under what it regards as "white" light the resulting calibrated and uncorrected digital image will look correct when properly displayed. However, suppose the light is really reddish (like a regular light bulb, compared to daylight). The human eye will compensate, so the scene still will look normal, but the resulting digital image will appear reddish because the Red component of the digital image is too bright. We can deal with this just as we do when the entire scene is too bright. We found that adding black to the color of a color patch produces the same effect that dimming of the original light source would have done. In this case, however, only the red component of the patch is too bright, so we add

black *only to the red component* rather than to all the colors of the color patch. In this way we dim the red only and exactly compensate for the non-whiteness of the light source in the scene. Of course we can use the same procedure of adding – or removing – different amounts of black to the red component, to the green component and to the blue component in order to color balance the image with a result that is similar to what the eye sees automatically. Again what we describe here is something an artist painter cannot do, but in this case it is no great loss since the artist finds it unnecessary. The eyes serve as sensors for a painter and so the artist has no problem similar to photographic color balance in the first place.

In the above we have learned that by using actions which effectively "add black" or "remove black" from the colored "patches" (pixels) in a digital image we preserve and maintain the *color integrity* of an image. Black is colorless, so this makes sense and in fact seems obvious. But white is also colorless. Is "adding white" the same as "removing black"? The answer is no. When we say "adding" here we do not mean "adding" in the mathematical sense, but in the sense of taking some colored paint and adding white paint to it or adding black paint to it. If you add some black paint to some colored paint and find that you have added too much black, adding white paint will not reverse the effect and get you back where you wanted to be. Seeing this in your mind is easier if you take an extreme case. Suppose you had a spoonful of red paint and you mixed it into a cup full of black paint, so the result was nearly black with a very slight red tinge to it. If you mixed a cup of white paint into that, clearly you would not get back your original red color, but instead a middle gray (black plus white) with only a slight red tint to it.

Nonetheless, as you would expect, white has no color and so "adding white" is another tool that preserves color integrity. The reason that "adding black" works is that it makes color "patches" behave as they would as the image lighting produces different levels of dark and lightness in the scene and so "adding black" is recognized by the eye as being natural. "Adding white" turns out to be exactly what happens in the areas of a scene that are bright enough to produce glare (or specular reflection). As the glare increases in bright objects more and more white is added to the color in the color patches. "Adding white" also turns out to be exactly what happens in a natural scene where there is fog. The individual "color patches" in the scene are the color that they would be without the fog but with white added to them. Again, just as it is possible to "remove black" in digital images it also is quite possible to "remove white" and indeed as is shown in the **Fog Example** section of this document, removing white from an image of a foggy scene can do a quite convincing job of removing the fog. You will not find any tool similar to "add white" in Photoshop. Unlike black, the tool for adding white is not just hidden in an unexpected place, it simply is not there. You can sometimes combine several actions and achieve the action of adding white but even that is not dependable.

In summary we have learned that to have color integrity in a final image the first step is to get the image into the computer using a calibrated device, and that the most problem-free calibration is done using a grayscale. The resulting image will have color integrity as a starting point. Then the major adjustments of color balance and of the tonal scale should be done using tools which preserve the color integrity. We have found that the required tools are

those which can add and remove black and those which can add and remove white to the color in the "color patches" that make up the image. For tonal scale adjustments there is nothing saying that there must be a consistent amount of white or black added to or removed from each of the color patches, so adjustment of the tonal scale can be made quite differently in different parts of the image without affecting color integrity – but only if it is done by adding or removing black or white to or from the color that would be in the patch. In the case of adding black in different amounts to the single color channels of Red, Green, and Blue in order to achieve color balance, there does need to be a consistent amount of black added to groups of patches which have the same light source.

So, the starting point for *color integrity* is a properly calibrated image which accurately represents the image colors as they appear in darker and lighter areas. Then to maintain color integrity as you adjust the image or parts of the image, the rule is to work by adding or removing black or white. Once you have grasped it, this seems such an obvious rule that surely it *must* be a core principle behind digital photography. Yet as this is being written, far from being a core principle, it is so completely unknown that in Photoshop, acknowledged as the premier image editing program, the tools to perform these four basic actions range from difficult to find to nearly impossible to achieve. In general, other image editing programs fare no better. In addition, many actions taken under the guise of "Color Management" actually subvert the well-planned ICC design. As one result, it has become more and more difficult to get an image from a digital camera into a computer without having *color integrity* lost from the start. Our **ColorNeg** and **ColorPos** Photoshop plug-ins are capable of adding and removing black and white, but there is room for improvement as a result of this study.

Perspectives and Comments

Introduction

Complete Color Integrity

Genesis of the Idea – what led to the discovery of these simple facts.

Why We Give Few Illustrative Examples of Color Integrity

Color Integrity from the Viewpoint of Artistic Painting.

Fog Example

Color Integrity from the Viewpoint of Basic Physics and Mathematics.

Trying to Deal With *Color Integrity* in Photoshop

Color Integrity and Color Balance – A Few Examples

Comments on Calibrating Digital Images.

"Acceptable Accuracy" in Calibration

Calibration and Color Profiling

The Pitfalls of Using Profiling as Camera Calibration

Genesis of the Idea – what led to the discovery of these simple facts.

While reviewing my earlier work on *color integrity* I made two key observations where the rules of color integrity as I had defined them did not seem to work as expected. First, with my mathematical approach, managing the colors in order to get and then maintain color integrity in a digital image required considerable care, often in knowing what to avoid doing. Yet it is rare to see a painting with the same sort of unsettling color. My past study of the methods of oil painting did not turn up difficulties of this nature. Oil paintings (that were intended to) typically and obviously showed color integrity. Of course oil painting requires considerable knowledge of color and how colors work together, but why did the color integrity problem not befall oil painters as it has photographers?

Equally puzzling was the phenomenon of "blackpoint," a common adjustment for digital images which completely violated what I then understood to be the rules necessary for maintaining color integrity. Yet the colors in an image remained very natural over wide ranges of blackpoint adjustment. Furthermore, the blackpoint adjustment is most often done incorrectly (as is the case in Photoshop). I found that when using the *incorrect* form of the adjustment, the colors lost their natural appearance. Blackpoint was regarded primarily as a correction for common flaws in the light sensitivity behavior of film and image sensors in general. If blackpoint really was just correcting for a flaw then it would be expected that there would be just one correct setting but instead the colors looked natural over a wide range when the correct form of the blackpoint adjustment was used.

These two puzzles caused me to explore both artistic painting and the mathematics of blackpoint to see if I could find answers. As it finally turned out, it really was necessary to run the two studies in tandem to arrive at a final solution that is both very surprising and yet quite obvious – after the fact. The study of artistic painting led me to understand that my earlier version of *color integrity* was exactly equivalent to the practice of adding various amounts of black to a palette color to produce the darker shades of the same color for a painting and that this in turn was exactly equivalent to lowering the light level on the original scene. Adding black does not change colors and so using it to preserve *color integrity* seemed obvious and that led me to explore adding white, which also does not change colors. I was amazed to find that adding white was exactly equivalent to using the *correct* blackpoint adjustment! This completely explained the blackpoint mystery. Adding white by various means in painting is used to produce glare in highlights and also to produce fog effects. Indeed, it is exactly equivalent to the presence of fog in the original scene. Thus the blackpoint adjustment can be used to increase or decrease the level of fog in a photographic scene, which it can do with marvelous accuracy. However, this requires the *correct* blackpoint adjustment. Photoshop and other image editing programs normally do the blackpoint adjustment incorrectly and so visibly damage color integrity long before significant fog can be added to or removed from an image.

Still, at this point I had discovered the solution to maintaining color integrity in an image – by simply adding or removing black or white – and it was mathematically obvious how this simple rule could be extended to deal with color balance while maintaining color integrity.

Why We Give Few Illustrative Examples of Color Integrity

I have been asked why I do not provide examples of images with and without *color integrity* to make it easier to see exactly what it is that I mean. In the **Color Integrity and Color Balance** section of this document there are two examples showing the general effect, but I do not hold them to be good examples. Creating demonstrations of this effect is more difficult to do than it sounds. If I start with an exhibit of images showing good color integrity then I need to intentionally *misadjust* those same images to get the examples with color integrity problems. This is what I have done in the two examples mentioned above. The resulting images would not likely satisfy any critical photographer. The professional images that you see every day that lack color integrity are created by good photographers trying their best to adjust their images out of the hole that Photoshop (or some other image editing program) has dug for them. Even when the color has the disturbing effects brought about by lack of color integrity the result is often still quite convincing. I cannot do that. I frankly am no longer capable of working naively on images that lack color integrity – because I understand color integrity, when I try to make the images convincing I just end up restoring color integrity.

Alternatively I might take images from a source like the National Geographic and attempt to recover them by reverse engineering the loss of color integrity. Even if that was not a difficult task in itself there would be copyright problems. I believe most readers who have gotten this far already have started to recognize the color integrity problem, at least in the back of their minds. As I have said from the start, the best way to see the problem is to take a magazine with a good, long-standing reputation for quality color images and compare some current issues with issues of the same magazine that are 20 or more years old, thus pre-Photoshop and probably pre-digital. The difference will become obvious – while the images in the new issues may have a more spectacular appearance, the color in the older issues will appear more natural, like you were really there. I suggest *National Geographic* because long runs are readily available in most libraries, but any quality picture magazine – including photography magazines – with a long enough run is also a good candidate.

Color Integrity from the Viewpoint of Artistic Painting.

Here we explore some very basic techniques used by artists painting in colors. Do not worry if you are not familiar with artistic painting as this will be very fundamental and simple, using techniques artists have known for centuries. Although photographers rarely if ever use these techniques directly, knowing them is essential to understanding how photographic color works. If you do artistic painting please bear with me as my terminology may differ from your own and my description of these basic techniques may have a different emphasis.

We will paint in color the following object, a ceramic bell:




The first basic technique that interests us is preparation of a base shade of a color to be used for the bell. To do this, we may start with a quite pure color directly from a paint tube or we may blend colors from paint tubes to get a source color which typically is fairly intense. As examples here we choose a blue and a bronze:




Next we prepare a "base color" for the bell. We mix the source color, with white and perhaps with black until we get the correct tone of the bell for areas that are directly illuminated; that is, the tone for the brightest fully colored area of the bell that does not show glare. This gives us our base color. Here the source colors have just had white added to produce the base colors.



The shape of the bell leads to shadow or shading of the areas of the bell which are not as directly illuminated. To produce the correct tone for the shadowed areas the artist mixes black paint into the base color; progressively more for the areas in deeper shadow:



This adding black is one of the two key steps in producing tones which are visually correct. The shadow areas are produced by adding black to the base color. The artists among you will know that while this is true, in practice a dark color is often used in place of black. That slightly advanced technique works just as well but includes more than we wish to at this point. The important thing to understand here is that *all* tones in the complete shading are produced by the blending of just two elements, the base color and either black or a dark color.



In the above bells, the brightest highlights are colored, which does not look right. They should be white. Rather, the *brightest* areas should be white but the surrounding slightly less bright areas should show some color. The artist handles this by going back to the base color and adding white to it:



All of the affected highlights in the bell are produced by this mixing of the base color and white (or very light gray). The resulting bell looks like this:



Again, all the tones are achieved by mixing just two paints, the base color and white. So, this means that *all* the tones in each bell are completely controlled through the use of just black paint and white paint.

Of course few objects have just one base color:



The same methods work for each color in an image. Mix the base color and black to produce shade and shadows for areas below full illumination on an object, mix the base color with white, sometimes with a little black added, to produce highlights that contain glare (specular reflections).

The fact that this works just the same for any color in an image is of particular interest to photographers. Since the methods work for each color in an image they also work for the Red, Green, and Blue primaries so widely used in photography. This is very important.

Through the centuries artists have come up with many variations on these basic methods, including many different ways of achieving the blending of paints. For instance an artist may observe that the light falling on the shadow areas of a subject has a brown or green or blue cast compared with brightly lit areas and so will use a very dark brown or very dark green or very dark blue instead of black in order to mix the shadow tones. The artist also has several different ways to do the actual mixing. While the use of these basic techniques varies from artist to artist, the basic principles we have covered above have been extensively and very successfully used for centuries and we will leave further discussion of variations to the artist painters.

Fog Example

There is another artistic technique that is of great importance to us, however. When an artist wishes to show a scene in hazy conditions or partially in fog, it is normally done using white paint. (Actually it may be slightly tinted white paint instead for effects like blue haze, but at this point we want to avoid tinting and so will use white.) With glare or specular reflections we added successively *more* white to the color as we went into successively brighter areas of the glare. Within an area covered by fog of a particular density we will add the *same* amount of white to each color regardless of the color or how dark or light it is. There are numerous ways to do this and here we will use a simple, direct technique. The easy way to produce a fog effect in a painting is to first do the painting with normal unfogged colors and allow it to dry. Then paint over the fogged parts of the scene with a light wash of white paint, so that the original scene shows through. The amount of white added varies with the depth of the fog, so we may use more white for more distant areas, or areas closer to the ground. In any one area of the painting where the fog level is constant, however, the same amount of white is applied for any color and whether light or dark:



In this example, some of the foreground objects do not have the white fog effect. But the above images are misleading. The appearance is that a pretty good fog effect has been added to the image on the right by adding white to the fog area. However, the image on the right is really the *original*. What actually was done was to *remove* white from the image on the right, producing the nearly fog-less image on the left! Adding and removing white is equally as effective going in either direction.

Artists usually prefer actually mixing paints to achieve fog effects. Instead of the "wash" technique described above an artist may take some base color and some black and mix the two in varying amounts as required for the painting. Then an even-width thick strip of the

shaded color is put on the palette, varying from dark to light along the length of the strip. A thick strip of white of the required width is laid along side and the two strips blended across. Using the same method and strip sizes for other colors puts the same amount of white into each shade of the blended black and colors, as required to produce the fog effect. Paint may be used freely as required from any point along the blended strips in building the painting. Again, more white can be added to the mix when painting areas in deeper fog.

All of the above "paintings" have been simulated in Photoshop. This was not easy to do with precision but within the confines of jpeg and gif images, the images are generally close to being accurate.

Color Integrity from the Viewpoint of Basic Physics and Mathematics

The physics and mathematics of adding black and adding white

When we speak of "adding black" in this document we mean "adding" in the sense of adding black paint to a sample of colored paint. In terms of a useful physical model, this can be thought of as taking a patch of color – for instance, a pixel – which has an area A and putting black paint over some fraction b of the area. Black does not reflect light, so only the remaining colored patch $(1-b)A$ reflects light. Since the reflecting portion is part of the original patch that reflects light, the color will remain the same but since only $(1-b)$ of the pixel will show, it will be only a fraction $(1-b)$ of its former brightness. If the light source illuminating the original colored patch were to be dimmed to $(1-b)$ of its former brightness the effect would be exactly the same as adding the black paint. Either way the colored patch or pixel remains the same color but reflects only $(1-b)$ of the light it did originally. Note that in both cases the initial color or lightness of the colored patch does not make any difference. No matter what the starting color and lightness is, the color stays the same but it is reduced to $(1-b)$ of its original lightness.

What does this mean mathematically in terms of pixel values? Each pixel in a digital image is described in terms of its Red, Green, and Blue (R,G,B) components, the three primary colors of digital imaging and the pixel represents a small point in the original scene. For a calibrated image ideally for each pixel the R component is proportional to the amount of red light coming from that point, G is proportional to the green light and B is proportional to the blue. R, G, and B are normalized so that when $R = G = B$ the color is neutral gray, ranging all the way from black to white. (This process of normalization is more commonly known as "white balance" or "color balance.") Because R, G, and B are each proportional to their respective components in the light source they individually behave exactly as described above for the entire pixel. When the fraction $(1-b)$ is applied to (R,G,B) the result is $((1-b)R, (1-b)G, (1-b)B)$. If we want to dim the entire image without losing color integrity we simply apply the fraction $(1-b)$ to each component of every pixel in the image. This is the same as "adding black paint" to every pixel in the image. But since we are working with a computer image and not with paint we can also lighten the image using the same method in reverse, effectively "removing black paint" from every pixel. Where a fraction $(1-b) < 1$ is used to darken the image, we simply apply a fraction $a > 1$ to lighten the image: (aR, aG, aB) . Note that this "adding black" and "removing black," that is, (aR, aG, aB) for $a > 0$ preserves the color of the pixel. This means that using this method we can lighten some parts of an image and darken others selectively and in as many places as we choose and still retain color integrity; that is, the colors will not change.

When we try to apply this very simple rule in Photoshop we immediately run into a problem. The R, G, and B in Photoshop are typically (almost always) "gamma-encoded." We have gone into the historical reasons for this elsewhere <http://www.c-f-systems.com/Docs/ColorIntegrityCFS-243.pdf> but Photoshop (and most of digital imaging) stores image pixels as $(R^\gamma, G^\gamma, B^\gamma)$ instead of (R,G,B) and so the image pixel components are no longer proportional to the red, green, and blue components of light represented by that pixel. As luck would have it (and clearly luck it was rather than planning) this form still works, sort of, with our simple "adding black" operation. If we take $(aR^\gamma, aG^\gamma,$

aB^γ) and let $a = c^\gamma$, then we have $(c^\gamma R^\gamma, c^\gamma G^\gamma, c^\gamma B^\gamma) = ((cR)^\gamma, (cG)^\gamma, (cB)^\gamma)$. So, multiplying the R, G, and B components of a gamma-encoded pixel by a we still get lightening or darkening of the image that preserves color integrity although the fractional amount of darkening or lightening is $c = a^{1/\gamma}$. In particular this means that color integrity preserving lightening or darkening of the image can be done in Photoshop by using the highlights control of the Levels tool. (The rest of the functions of the Levels tool damage color integrity, however.)

This bit of serendipity does not always work. For those who have changed to L* encoding in place of gamma encoding in the hope of getting more accurate color (if you have done this, you'll know it) will be thrilled to know that the above trick does not work with L* encoding, nor does it work with certain other profiles such as sRGB which do not adhere to gamma-encoding. For those cases Levels highlights and the few other tools in Photoshop that function by multiplying pixel components by a constant fraction do not preserve color integrity. In fact we are hard-pressed to come up with any means within Photoshop of properly "adding or removing black" without first converting out of these profiles. Our ColorNeg and ColorPos plug-ins do function correctly with L* encoding, however.

We have noted that the "adding or removing black" operation, which mathematically is multiplying each of the Red, Green, and Blue light intensity values of the pixels by the same constant fraction, is exactly equivalent to what would happen if we brightened or dimmed the source lighting for the original scene. Suppose that instead of brightening the source lighting we shifted its color so that the red component of the source was brighter but the Green and Blue components stayed just the same. Our camera or film would then record the image as too red. To compensate for this we need leave the Green and Blue pixel values alone but to decrease the Red pixel values in order to compensate for the increase in red source lighting. If the red light component increased by 5% then the red light reflected by the object in the scene will increase by 5% and the Red pixel values of the image will be 1.05 times as large. To correct this we need to divide the Red pixel values by 1.05 so that $(1/1.05)(1.05R) = R$ and the Red pixel values will be restored to their correct values. Doing this is called white balance or color balance. In order to correctly adjust the color balance of a (calibrated) image all of the Red pixel intensities in an image need to be multiplied by the same fraction and the same type of adjustment may be required using different fractions for Green pixel intensities and for Blue pixel intensities. The highlights adjustment in the Photoshop Levels tool may be used to do this, operating individually on the color channels. Again, this also will work with gamma-encoded pixels, but not with other non-linear encodings such as L*.

We have dealt with adding and removing black. How about adding and removing white? By "adding white" again we mean in the sense of adding white paint to a pixel patch of color. In this case the patch area is again A and we cover a fraction of the d of the patch (pixel) with white paint. With black paint we reflected no light, but with white paint it is necessary to explain in more detail what we mean by "white." In digital imaging, white is taken to mean the brightest pixel value possible. If we call this maximum value W , then white is the condition where $R_w = G_w = B_w = W$. Returning to our patch A covered with fraction d of white paint, the original pixel (R,G,B) will become $((1-d)R+dW, (1-d)G+dW, (1-d)B+dW)$. This is an "addition of black" or dimming of the original color where the white partly covers

the patch combined with the same constant quantity of white added to each color component for the white portion of the color patch.

This model applies to fog, where something – like real fog – is between the main subject matter of a scene and the camera. If we consider typical white fog, droplets which are between the subject and the camera reflect white light themselves while they block light coming from the subject. Light from the subject passes between the fog droplets, however, so taking the fraction of the area that the fog droplets actually obscure as d , the physical situation matches the mathematical model. This model for "adding white" also applies for specular (mirror) reflections in the highlights, where the nature of the surfaces and the angles cause the light source to be partially or completely directly reflected instead of being diffusely reflected as is required to show the color of an object. Where the transition to specular (mirror) reflection is still partial it still diffusely reflects some colored light added to the specular white light. There is a fraction d involved just as for fog, with the value of d changing with the transition to full specular reflection.

Returning to the "adding white" $((1-d)R+dW, (1-d)G+dW, (1-d)B+dW)$ form, we can see that this is really a combination of adding white – the $+dW$ in each of the three color component terms – and adding black to account for the part of the color patch which we have covered with the white paint – the $(1-d)$ multiplier in each of the three terms. For the fog relationship the "black" part of the term and the "white" part of the terms are related through the fractional area d , but in the generalized of the black/white transformation for (R,G,B) we have $(aR+dW, aG+dW, aB+dW)$. Note that $a < 1$ is adding black, $a > 1$ is removing black, $d > 0$ is adding white and $d < 0$ is removing white. Since this transformation involves the addition or removal of black or white the transformed colors have color integrity. The appearance of a calibrated image in which the pixels have been altered in this way will retain the natural appearance of the colors. To use this form, the same a and d constants need to be applied to each of the R, G, and B components of a pixel, but there is no requirement that the a and d be the same for *all* pixels and indeed they can be varied as required to produce different effects for different parts of the image and the image will still retain its color integrity.

For the above we assumed we were adding pure white, $R_w = G_w = B_w = W$. Suppose instead that R_w , G_w , and B_w are a little different so that the white has a tint, $R_w = d_r W$, $G_w = d_g W$, and $B_w = d_b W$. Then $(aR+d_r dW, aG+d_g dW, aB+d_b dW)$ is adding tinted white, which is also a natural situation. For this form, R_w , G_w , and B_w should of course be kept the same over larger areas to represent tinted fog, but a and d are under no more constraint than for the black/white generalized form.

Finally, we can make use of a completely generalized form $(a_r aR+d_r dW, a_g aG+d_g dW, a_b aB+d_b dW)$. Here in addition to tinted fog we also include color balance or white balance. For this form R_w , G_w , and B_w should again be kept the same over larger areas and in addition a_r , a_g and a_b determine color balance and should be kept the same over areas that are subject to the same lighting. But once again, a and d are under no more constraint than for the black/white generalized form.

This final generalized form

$$(R,G,B) \rightarrow (a_r a R + d_r d W, a_g a G + d_g d W, a_b a B + d_b d W)$$

is a complete representation of adjustments that can be made without affecting color integrity. Note that the (R,G,B) used here and throughout this section are quantities that are proportional to the normalized light intensities of red, green, and blue. That is, the (R,G,B) is linear and not gamma-encoded. Furthermore, while the "adding black" form serendipitously worked directly with gamma encoded pixels, the generalized form definitely does not. It is necessary to linearize encoded (R, G, B) values before use, regardless of their encoding.

Investigating the common blackpoint or shadow adjustment led to much of the above. Some form of this adjustment needs to be made to nearly all digital images:

$$(R,G,B) \rightarrow (R-b, G-b, B-b)$$

or

$$(R,G,B) \rightarrow (R-R_0, G-G_0, B-B_0)$$

Various reasons are given for making this adjustment, typically having to do with various defects in the response to light of camera sensors or film or other equipment defects. While this may in fact be true it is also true that the blackpoint adjustment is the same as we found above for removing white, so that it effectively lifts fog from the image, whether white fog as in the first form or tinted fog as in the second form. For all practical purposes, this adjustment is not available within Photoshop itself for normal gamma-encoded images although the "shadows" (earlier versions) or "black" (later versions) adjustment in the ACR (Adobe Camera Raw) plug-in appears to do the first form above. This blackpoint adjustment is very often applied directly to a gamma-encoded image, losing color integrity from the start.

Since the addition of white actually is a mathematical addition, you might think that to add white in Photoshop it would be easy to simply put in a layer which is white – even a tinted white – and then blend a certain percentage of that with the image. After all,

$$(R, G, B) + k (R_w, G_w, B_w) = (R+kR_w, G+kG_w, B+kB_w)$$

Except that again Photoshop works directly with gamma-encoded pixel values:

$$(R^\gamma, G^\gamma, B^\gamma) + k (R_w^\gamma, G_w^\gamma, B_w^\gamma) = (R^\gamma+kR_w^\gamma, G^\gamma+kG_w^\gamma, B^\gamma+kB_w^\gamma)$$

and that result is not equal to the required $((R+kR_w)^\gamma, (G+kG_w)^\gamma, (B+kB_w)^\gamma)$.

Trying to Deal With *Color Integrity* in Photoshop

Adding Black and Adding White in Photoshop

When we try to "add black" or "add white" in Photoshop we immediately run into a problem. The (R,G,B) values are typically (almost always) "gamma-encoded." To really understand the reason for this requires mathematics, but the gist is that traditionally most digital images have been 24-bits per pixel (which is the same as 8 bits/channel). When pixel values are expressed directly in terms of light intensities for each of red, green, and blue and crammed into 24 bits many visible colors get pushed out the edges and are lost. To cram the pixel values into the 24 bits without losing visible colors the pixel values had to be compressed, the same general idea as compressing zip files. This was done using a process called gamma-encoding. Gamma-encoding was convenient because it was already in widespread use in the analog video world for an entirely different reason. Unfortunately, Photoshop and most other image editing programs generally try to work directly on the encoded pixels rather than on the intensity values that they represent. Sometimes this actually works satisfactorily. Often, and to varying degrees, it does not. Unlike the earlier days, Photoshop now offers the capability for images to have more than 24 bits per pixel. Although this makes gamma-encoding unnecessary Photoshop still uses gamma-encoding in almost all cases and worse, still works directly on the encoded pixels.

One place where working on gamma-encoded images does work satisfactorily in many cases is with "adding black." Photoshop has several tools which perform a function equivalent to adding black, but none of them are obvious. The preferred tool is Levels, where the "highlights" portion of the tool does the correct action of adding black or removing black over much of its range. The shadow and mid-range functions of Levels destroy color integrity and are to be avoided.

The fact that Levels highlights control will work to properly add or remove black from a gamma-encoded image is pure serendipity. Although most working profiles are gamma-encoded, some such as sRGB are not and other new working profiles are coming into use which use L* encoding rather than gamma encoding. For these profiles, Levels highlights does not accurately function to add or remove black.

When we try to "add white" the situation is considerably worse. The Levels shadow or blackpoint tool might have served this purpose just as the Levels highlight tool will correctly add and remove black in many cases. But the Levels shadows tool works directly on the gamma-encoded image and in this case it simply does not work, often destroying color integrity quite noticeably. For digital camera images the ACR (Adobe Camera RAW) plug-in "shadows" slider (earlier versions) or "black" slider (later versions) appears to make the correct adjustment for removing white, but it does not provide for adding white or for working with tinted white. You might think that putting in a separate layer of white and then adjusting its percentage transparency should work, but the additions of layers appear to be done to the already gamma-encoded image, so even that does not work (see at the end of ***Color Integrity from the Viewpoint of Basic Physics and Mathematics*** for details on this). Technically it would be possible to convert the gamma-encoded image to a "linear" image, make the "add white" adjustments and later convert back, but that is tedious. It depends upon

Photoshop to correctly convert the image encoding twice as well as making the adjustment for adding white, all without applying hidden "improvements" along the way, never a good bet. To avoid what baggage might be hidden in a "profile conversion" one could choose to do the gamma-encoding and decoding using the Levels middle-gray adjustment which is widely known to be a gamma adjustment. But Levels middle gray deviates from being gamma and the deviations are greatest just where they will do the most harm to this decoding-encoding. (See <http://www.c-f-systems.com/Docs/ColorIntegrityCFS-243.pdf> page 18 for details.) Hidden somewhere among the plethora of Photoshop tools there may be one or two that are capable of adding or removing white accurately, but as I write this I have not found any.

Also as I write this the ColorNeg and ColorPos plug-ins do correctly add and remove black with the "lightness" slider and do correctly add and remove white with the "shadow" slider. Technically when adding white with the shadow slider you also need to add a little black with the lightness slider to account for the colored areas that the white "paint" covers. In the future ColorNeg, ColorPos and perhaps a new plug-in will make more complete use of what has been learned in the study that led to this web page.

Color Integrity and Color balance – A Few Examples

People who are new to the concept of color integrity tend to confuse it with color balance. We have explained and used both of these concepts in several ways on this *Color Integrity Complete* page. Here we will give just a few examples. I make no claim that these are good photographs or that they are good illustrations of what I mean by lack of color integrity. See the section **Why We Give Few Illustrative Examples of Color Integrity** to understand. First,



The image on the left has color integrity. You may not agree with its color balance, but if you don't it is an easy matter to change it to your liking. The image on the right does not have color integrity. You may not agree with its color either, but it will not respond well to changes in color balance. In fact, you would find it difficult to adjust it by any means and really get it right.

Similarly,





Again, the image on the left has reasonably good color integrity. And again you may disagree with its color balance – in fact I often think it is too red myself. And since it has color integrity its color balance can be easily adjusted. The image on the right does not have good color integrity. If you compare the skin tones in the two images you will find that they are very similar while the ocean is very different in color. It would be nearly impossible to get the image on the right to look right – natural. For both the elephant and the beach walkers the differences are in adjustments that have been made to the *entire* image. In neither case was a selection made and part of the image treated differently.

Sometimes getting the color balance right can require very exacting adjustments of the colors of an image. We present here an image in which three different color balances have been applied:



In this image there really is no object where the viewer is really certain of its precise color, yet most people – photographers, at least – viewing these three versions will definitely prefer one to the other two. However, not all people will choose the same one of the three images and some people will think none of the three is really correctly color balanced. These images are really quite close in color balance (plus or minus 5CC), especially compared to what you see below. The principal reason for the sensitivity to color balance is the rock. We know it is a gray rock but that it likely is not precisely gray. Our differing experiences lead to disagreement as to exactly what tint it might take on. Seen by themselves rather than side by side any one of the three images might look properly color balanced.


The next image is at the other end of the color balance sensitivity range. This is a scan of a print which I made years ago using wet darkroom processes – the reddish spot at lower right is actually a defect of aging.



The image on the left is approximately as it would have been white balanced. The image on the right shows a very different color balance which most people would accept or even prefer to the white balance. Even placing these two images side by side does not make one or the other really jump out as wrong. But we can go farther than that:



This shows that the same scene can be color balanced much more to the red and still not be objectionable to most people, especially if not seen in comparison to the others. We do not wish to give the impression that this is all a warm-cool, red-blue effect, so here is one that has been shifted to the green:



This extreme case works as well as it does for two reasons. First, the color of the lighting in a hazy sunset is ambiguous. The eye expects this and furthermore it expects the lighting to be changing rapidly with time at sunset. But this still would not work well if the image did not have reasonably good color integrity to start with. In each case the colors running from dark to light blend as they would in natural lighting.

Comments on Calibrating Digital Images.

Calibration is the process of characterizing the camera or the film scanner so that the images it produces correspond to the colors and variations of color in the original scene with acceptable accuracy. Note that this does not mean the image will be what we want as a final product, only that it accurately represents the original scene.

We find that it is best to calibrate a camera or scanner using a target having a stepped grayscale for which the RGB pixel values are known for each step. The quality of the grayscale is important as the steps should be uniformly gray with no tinting. The calibration results from comparing the camera's image of this grayscale with the known values for each step. This comparison involves aligning three elements for each of the three channels (R,G,B) so the camera image and the known grayscale values best match. The three elements are 1) a white balance adjustment to account for the lighting used to take the camera image, 2) a blackpoint adjustment, and 3) a "curve" which describes how the camera or film responds to increasing light levels. Of these three elements, the first is specific to the lighting and will be different for different lighting conditions. The second typically results from a combination of several factors and will be different for images taken under different conditions. Only the third element depends just on the digital camera or film being tested and so only the third actually is a calibration of the digital camera or film. These third corrections, the "curves," need to be applied to all images from that digital camera or that type of film. Blackpoint adjustments and white balance (color balance) differ from image to image. These cannot be corrected automatically as part of the calibration but must be determined as required for each different photographic situation; sometimes for each image. As we discuss several other places in this document, the first adjustment, establishing color balance, is done by "adding and removing black" from the color channels and so does not affect color integrity. The second adjustment, blackpoint, is done by "adding or removing white" from the color channels and so it also does not affect color integrity. Since neither of these adjustments affect color integrity only the third adjustment actively causes the calibration to establish color integrity in the image. Calibration methods that are currently in use often make the mistake of including the first and/or second element as part of the calibration. The first and second elements need to be accounted for while calculating the calibration, but should not form part of the calibration itself.

We have a web page < <http://www.c-f-systems.com/DunthornCalibration.html> > that goes into detail on grayscale calibration and describes how it can be done in Photoshop. Our **ColorNeg** and **ColorPos** plug-ins <<http://www.c-f-systems.com/Plug-ins.html>> for Photoshop provide several methods of grayscale calibration, using known grayscales as described above, but also using grayscales in which the target pixel values for the steps of the grayscale are not known and even for a "grayscale" selected from the natural grays available in many images. The plug-ins also have a non-grayscale calibration method called "FilmType" which actually uses the fact that most people will select an image with color integrity as more natural from a series of images with varying loss of color integrity. The curves that produce the most natural image are thus a good estimate of the calibration curves. The calibration methods in these plug-ins pre-date the work that led to this *Complete Color*

Integrity page and will likely be improved as the result of this work, but in their current state they still work reasonably well with the new concepts.

All of our calibrations use *gamma* curves for the calibration of films, cameras, and scanners. It is certainly possible to use other more complicated curves for calibration and as *gamma* curves are nearly as old as photography itself, it certainly is valid to question using them in this age of computers. But our experience is that using a more complicated form rarely results in improved performance of the system. With film the actual calibration curves are never known with high precision due to variations in processing and handling. The use of a system of gamma curves is much more forgiving of variations in the film than is the use of more complicated curves, particularly where the complicated curve shape is different for Red, Green, and Blue. As for most digital cameras, if the image data is available in truly raw form, the "curve" does not depart much from a straight line (or a gamma of 1) except possibly for the very brightest part of the image – a small range of 1/3 stop or so – and there is very little that can be done about that brightest somewhat unstable part in any event. In fact, the main problem with digital cameras is getting the raw camera image data into a computer image file without having the program doing the transfer "improve" the image as the transfer is made. We will have a web page dealing with this problem very soon, we hope. Meanwhile, the literature on our **ColorPos** plug-in has some suggestions.

One obvious question about the above is why use just a grayscale? This is a color calibration – don't you need color patches to do an adequate calibration? In the past we have explained why using color patches is not *necessary*. In a digital image gray is composed of equal amounts of Red, Green, and Blue. We calibrate the camera or film so that the responses to Red, Green, and Blue are equal on each step of a grayscale. This fully defines the way the film or sensor responds to Blue light as it goes from dark to light and the same is true for Green light and for Red light. But as you will learn below, it is also true that you cannot do a better calibration than using a grayscale although many – probably most – digital camera makers claim to.

"Acceptable Accuracy" in Calibration

The key to proper calibration is in understanding exactly what "acceptable accuracy" *must* mean for calibration to work properly. Cameras – digital cameras or film cameras – are simply not capable of producing a completely visually accurate record of the colors in a scene, but a camera *can* produce scenes with colors that are completely visually realistic. This is because the sensitivity of the eye to colored light is different than the camera sensitivity to colored light. Generally the two can see colors as nearly the same, often visually identical, but also quite often there is a visible difference. An object you see as a particular shade of blue the camera might record as a slightly more greenish blue. But this difference is not consistent. There may be another, different object that you see as an identical shade of blue to the first blue object, but the camera records that object as a slightly more reddish blue than you see. It is important to realize that these differences cannot be corrected by adjusting the camera image to be slightly less green to agree with the first object because then the second object would have an even more pronounced reddish cast. Likewise, making the camera image slightly less red would correct the second object but then the first

object would appear as even more greenish. The phenomenon that leads to these differences in the way the camera sees and the way you see is called *metamerism*. It is the same phenomenon that sometimes causes an item of clothing to be a different color under store lighting than it is when you get it out in daylight. If you want to more fully understand metamerism, you can do no better than looking in the standard textbook on color, Roy A. Berns' "Billmeyer and Saltzman's Principles of Color Technology."

Just above, metamerism made the camera see as two different shades two color patches which we see as matching. The reverse of this is just as common. The camera can see two color patches as identical in color while to us one patch may be slightly more red than the other. Again, it is impossible to "fix" this situation. If the two patches have an identical color in the camera's image there is no way to tell whether or not that color came from a color patch that we see as more red so there is no way to know whether or not the color should be adjusted to be more red. That is, there is no way in general to know which sort of color "patch" that color came from. If we are taking a picture of a target with colored patches, it is possible to use our special knowledge to tell exactly which color patch is which. And that is where camera calibrators get into trouble. They can "calibrate" the colored patches to be different than what the camera actually sees.

Metamerism is a natural phenomenon. The eye routinely has to deal with metameric effects as the character of the light changes in a scene. Even daylight has large changes in character (technically "in spectrum") as cloud conditions change, as natural reflections alter the light, as time of day alters the light, etc. And so metameric color shifts are usually seen as natural unless the shift is uncommonly large or there is direct comparison between the image and the original scene. Another property of metameric color shifts is that they are consistent whether the patch is brightly or dimly lighted. If the camera sees a particular blue object in bright light as slightly redder than you do, then in dim light from the same light source it will also see the same object as slightly redder than you do. This consistency in the light/dark behavior of metamerism is important. As a result metamERICALLY shifted colors will behave as is required for *color integrity* and will look natural even though they may not match the colors of the original scene. So, for an image to have color of "acceptable accuracy" the key is to have the image give visually the *same* color for a given colored object whether the object is brightly or more dimly lit in the scene. This means that in calibration of a camera or film, metameric effects can be ignored and in fact it is best to do so. As we have shown above, trying to compensate for metameric effects at best just decreases the metameric color shifts of some colors at the expense of increasing the metameric shifts of other colors. Worse, as we show in the section on profiling, some methods of "calibration" that are in common use can actually destroy color integrity in such a way that it is irretrievable.

Calibration and Color Profiling

In the present state of digital imaging nearly all attempts at calibration of digital cameras or film are referred to as *profiling*. This is because most of the methods used are based on the ICC Color Management (<http://www.color.org>) system in which *profiles* are used to characterize various imaging devices. There are very serious problems with this approach. This is not a fault of the ICC Color Management system itself, which is in fact well-designed,

ingenious, and very successfully applied in color managing displays, printers, and commercial printing. The problem is that when it comes to cameras the ICC system is widely misunderstood and has often been grossly misapplied.

The Pitfalls of Using Profiling as Camera Calibration

Given a digital image – that is, a computer file representing an image – the purpose of the ICC Color Management system is to make that file visually appear as similar as possible when rendered on various devices, CRT displays, LCD displays, ink-jet printers, laser printers, commercial printing services, etc. It does this by using *profiles* that describe how each such device responds to image data to produce color images. "Device" as used here is a very rigid term. When you use a different type of paper in a printer it becomes a different "device" and requires a different profile even though you think of it as the same printer. As a CRT ages its behavior changes so that it becomes a different "device" and needs a different profile to make its displayed colors acceptably accurate. ICC profiling also extends to scanners. Scanners have a built-in light source that is controlled by the scanner. A page that is scanned in a scanner today can be expected to produce the same result if it is scanned in that same scanner next week. So, the scanner is a "device" and can be profiled. If we use ICC profiles to scan a page on a scanner, view it on a display, and print it on an ink-jet printer we can expect that the displayed and the printed page will look very much like the original paper page that was scanned.

In trying to extend this to cameras we first find that a camera is really not a "device." The source of illumination is not constant as is the case with scanners. While some photographic studio or photographic laboratory situations can be exceptions, for general photography the light can come from many different sources at many different intensities and can be altered by sky conditions, reflection and in many other ways before it reaches the subject. Properly applied, the ICC system would require a different profile for each one of these myriad light sources. Since that is impossible, the typical approach has been to pick *one* light source and profile for just that one source. This is given an official appearance by choosing an official-looking light source, typically D50 or D65, and going to great pains to be exact about it. But neither D50 nor D65 illumination match any real lighting conditions and so any real image that is "calibrated" by using a D50 camera profile will not be accurate and will be very inaccurate for images taken in, say, daylight. To compensate for this, fudge factors that have nothing to do with ICC profiling are introduced, but the whole package is still called "profiling" so that it appears to be under the wing of ICC. Apparently this makes everyone feel better.

This would be bad enough, but there is another current trend that is even worse. In dealing with cameras, film or digital, or scanners we are (nearly) always dealing with systems of three primary colors, additive Red, Green, Blue or sometimes the subtractive Cyan, Magenta, Yellow set. The ICC has profile formats which are designed to deal with the special considerations required for three-primary systems. But the big triumph of ICC Color Management has been in dealing with printers, both desk and large printing service systems. Color printers are commonly based on multi-color systems rather than three-primary systems. The most common multi-color system is CMYK (Cyan-Magenta-Yellow-black) but printers often use even more colors. The ICC has profile formats which are designed to deal with the

special considerations of these multi-color systems. These multi-color profiles naturally have a lot of control over color matching, and can allow quite precise choice of inks which best match the various colors and tones throughout much of the visible range. This is possible because when there are more than three colors of ink, some visible colors and tones typically can be represented by many different combinations of the inks and in other cases colors and tones can be represented that would not be possible if just three of the ink colors were used. Where multiple choices are possible the choice may be made on economic grounds.

At the start traditional standard color charts were used for checking and calibrating digital cameras. These charts had, in one form or another, numerous patches of color as well as a gray scale and there had been colorimetric measurements made on the various color patches which defined each of their colors and thereby which RGB values each patch ideally should have in a computer representation. When using specified lighting and one of these color charts to calibrate a digital camera to a three primary color ICC profile it was possible to get a reasonably close match to most of the color RGB values but naturally, due to metamerism some patches might closely match while some others might be visibly different. At some time in the development of digital cameras someone had the bright idea of using the ICC multi-color profiles intended for printers with these three primary color RGB systems. By choosing one of the most detailed multi-color profiles they found they could actually "calibrate" the camera so that every one of the color patches and gray scales was an exact match! Thus the color chart would be duplicated exactly and of course the "superior" color ability of the system became a major selling point. The problem is that this approach is completely fallacious. Given correct lighting conditions the resulting profile will indeed make images of *the color chart* that are perfect matches to the original. At the same time, the profile will make the color response of the camera worse for nearly everything other than the color chart itself. With a camera profiled in this patchwork quilt manner colors can shift several times going from darker to brighter illumination – color integrity is not only lost but becomes nearly impossible to regain.

The serious problem described above really should be obvious, at least to persons with a mathematical background as presumably would be the case for persons setting up profiling calibrations. We can be generous and assume it really must not be that obvious rather than conclude that we have a lot of disingenuous people who have programmed calibration profiles for digital cameras. Being generous, let me try to explain the problem. Nearly all digital cameras have a single sensor chip with an array of identical light sensitive cells. To make the chip color sensitive, a tiny red, green, or blue filter is placed over each sensor cell. Since all the sensors are identical that means that the response to light as it goes from darker to lighter is identical for all cells regardless of color. In addition, the chip is normally designed so that the output it produces is proportional to the light intensity it sees (that is, it is linear) over all or nearly all of the range of interest. Therefore, any calibration should deviate very little from linear for all three color channels. Furthermore, there is no crosstalk between the color channels that is of a predictable nature. The red channel simply measures the amount of red light and that measurement is not influenced in any predictable way by the amount of light that the blue and green sensors see. For this reason any "calibration" which tries to adjust the red value of a pixel differently for different readings in the corresponding blue and green pixels is a false calibration. This in effect changes the near-linear "curve" of the red (or other

color) channel to what might be described as lumpy, and with a lumpiness that varies according to the color levels sensed in the other channels. This plays havoc with color integrity, which demands a smooth transition from dark to light.

This can be confusing because the BAYER interpolation usually used on digital camera images does adjust the red value of a pixel differently according to the readings of nearby blue and green pixels. But this is done according to the geometric relationship of several surrounding pixels and the values each pixel is sensing. Geometry within the image is the key. A specific pair of blue and green pixels will influence the red pixel value in *different* ways depending upon the geometry of placement of the pixel values within the image. The calibration form we treat above requires that the red pixel be changed the *same* way any time a particular (R,G,B) color is found. In fact BAYER interpolation generally does damage to color integrity, but BAYER interpolation has its largest influence at edges within the image and does little to the larger areas of similar color to which the eye is most sensitive. BAYER interpolation is necessary and the effect it has is not very harmful visually.

Finally, and primarily because I get asked about this, a comment on the transfer matrices which play a part in ICC profiles. These are linear transformations which are intended to accurately convert colors from their expression in one set of three primaries, r, g, b , to another set of three primaries r', g', b' . We will assume for the moment that the primaries r', g', b' are known for the target system. In order for the transformation to be meaningful, the three primaries r, g, b of our camera or film also need to be known. They aren't. So, the use of transfer matrices with film or digital cameras is basically meaningless and can be harmful.

To explain, a CRT display uses phosphors which radiate the red, green, and blue colors each at a very specific wavelength of light. This direct connection between the primaries and specific wavelengths gives the CRT very specific values of r, g, b . With film and digital cameras there is no direct connection between the primaries and a specific wavelength of light. For example, the red filters used in a digital camera pass a wide band of wavelengths of light throughout the red region of the spectrum. When single wavelengths "primaries" are given for a digital camera they represent averages (integral averages) over the band of wavelengths actually passed by the filters and moreover the average is weighted according to a standard light source (D50 for example). If the many different light sources which the camera will actually experience were used in the calculation instead of D50, there would result a whole collection of different r, g, b "primary" values each of which is as valid a characterization of the camera primaries as any other. All of these result in are very similar rgb systems and it makes much more sense to just use the target r', g', b' as the primaries for the camera rather than converting to r', g', b' from some artificially averaged r, g, b for the camera that is bound to be incorrect most of the time. The situation is basically the same for film cameras, see <http://www.c-f-systems.com/Docs/ColorNegativeFAQ.html#CM> for more detail.

By this I do not mean transforms via transfer matrices are not ever useful. In situations where the difference in primaries is larger and/or the source has a fixed lighting system, transfer matrices are very useful. But for a typical general use camera or film the transform is wishful thinking at best and may well serve to degrade the image rather than improving its accuracy.