

ICC Profiling for Digital Cameras

Tak Auyeung, Ph.D.

June 21, 2005

1 Problems to be Solved by Profiling

So you bit the bullet and purchased an expensive digital camera. You also bit the bullet to purchase the best lenses. You take some pictures, you enjoy the sharpness and resolution. *But*, you find the colors, tone and contrast of the produced JPEG files not to be exactly what you want.

No big deal, just PhotoShop it. However, you find that with only 256 shades of red, green and blue, any major adjustment creates artifacts that should not be associated with your expensive digital camera. You get furious, but it is too late to return the digital camera.

Wait! The problem is not with your digital camera, but with the way your camera internally processes the captured image. Although JPEG only allows 256 values for red, green and blue (for each pixel), most expensive digital cameras (US\$500 and up) captures 1024, 4096 or even more values for each color channel. More color resolution means more room for manipulation before artifacts are visible.

The best news is that you can get that 10-bit, 12-bit or 16-bit (per color channel) image! Many cameras can record in “raw” mode. This means the captured file is simply what a camera’s sensor sees, before any additional processing is performed. No contrast enhancement, no tone modification, no white balancing, nothing. You get *all* the information that the camera captures, not a bit less.

However, no processing also means there is no color management. You are simply looking at the ADC (analog to digital converter) values of each red, green and blue sensor on the camera. There is a little catch here. The sensitivity of a color sensor is only centered at a particular frequency, but it is also sensitive to neighboring frequencies. Different sensors have different characteristics. As a result, if you look at a raw file, the colors are “off”.

Also, the raw file records exactly what a sensor sees, without any adjustment of “gamma”. We’ll talk about this later, but a raw file would appear too dark because of this.

And that’s why we need color management for raw files!

2 Geeky Background Stuff

If you don’t want to get all technical and *understand* some of the science, skip this section altogether.

2.1 Gamma

A camera pixel sensor records “linearly”. This means that if a sensor is recording a light level that is twice of another one, the former has a value that is twice that of the latter. If this is the case, how come raw images appear too dark?

When an image is displayed on screen, the displaying software *assumes* the image is the result of a device with $\gamma = 0.45$. This convention has a root that traces back to the NTSC video standard. This constant is used in the equation $o = i^\gamma$, in which o is the output, i is the input. Both o and i are scaled from 0 to 1.




Figure 1: This is what a linear raw file looks like when displayed as a “normal” image.




Figure 2: This is what figure 1 *should* look, after applying $o = i^{0.45}$.

Given a constant input (light level, or luminance) of 0.5, the pixel value in an image file assuming $\gamma = 0.45$ is $0.5^{0.45} = 0.732$. However, a raw file assumes $\gamma = 1$, which means the same luminance has a pixel value of just 0.5.

When both images are display on screen side-by-side, the raw image appears dark because of this. If one applies the function $o = i^{0.45}$ to a raw file, then display the output file, the brightness and contrast should be approximately “normal”.

2.2 Color Space

It is all about what we can see, and how to encode that in a file. Most of us are familiar with the RGB (red, green and blue) standard. Afterall, that’s how we specify colors in HTML documents. As it turns out, RGB is not enough to represent all colors that are visible to the human eye.

For a technical treatment, you should read the original article at http://www.normankoren.com/color_management.html. In a nutshell, our retinas have “RGB” cone sensors, much like a camera has RGB sensors. However, the response curves of our retina sensors have very wide bases. Indeed, the red and green cones are sensitive to the entire visible spectrum (with different sensitivity at different frequencies, obviously).

At the same time, the RGB pixels on a screen (regardless of technology) are very “spiky”. This means each color pixel has a very narrow emittance bandwidth. As a result, the produced image can never fill the color space of what a human eye can perceive. In order for a screen to produce all perceivable colors, the RGB emitting pixels must have band responses that match those of a human eye.

If RGB cannot color the entire space that a human eye can perceive, how do we encode all possible colors?

The answer is the CIE XYZ specification. The specification is based on a human eye color response model (by experimentation). X, Y and Z are basically the sensed level of red, green and blue cones, respectively. From this specification, a diagram called “chromaticity diagram” is drawn. Follow <http://www.efg2.com/Lab/Graphics/Colors/Chromaticity.htm> to view it.

In this diagram, there are only two axes, x and y. This is because it is normalized so that only the ratios among X, Y and Z are important. $z = 1 - x - y$. The edge of the horseshoe defines the pure colors that we can perceive, and one can associate frequencies to each point of the edge of the horseshoe.

But how about all those “muddy” colors that we cannot find in the chromaticity diagram? Recall that the chromaticity diagram is only representing the ratios of X, Y and Z. We know that $z = 1 - x - y$, so we do not need to represent z . To identify a single color, all we need to use include one of X, Y or Z so we know the “scale”. Because Y (green) is centered in the visible spectrum, we usually use that for scaling purposes. As a result, any perceivable color (and then some) can be represented as xyY (x and y for the color, Y for scaling). This is also sometimes called xyL, and L stands for luminance.

Now we finally have a (not so intuitive) method to encode all perceivable colors!

2.3 sRGB

sRGB is a color encoding standard based on the ratios of red, green and blue. To be more specific, the three colors correspond to the emittance of “standard” red, green and blue phosphorous pixels. Not very surprisingly, this standard is made to encode all colors that can be displayed by a standard CRT color monitor.

The color space of sRGB is a triangle in the horseshoe (in the chromaticity diagram). Refer to http://www.normankoren.com/EFG_chromaticiey_SMPTE.jpg. This means there are colors that we can see, but a monitor cannot display. From the positioning of the sRGB triangle, a lot of “pure” green colors, for example, cannot be displayed. Some “deep red” and “deep violet” colors are also missing.

sRGB is the most popular and basic standard for most images for online viewing. Nonetheless, one should remember that once an image is expressed in sRGB space, some of the original colors may be lost.

2.4 Camera Color Space

RGB sensors in a camera have response curves that are wider than the emittance curve of RGB phosphorous on a monitor. This is good news because it means a camera can record colors *out of the sRGB space*. It is absolutely necessary for a camera to record out of the sRGB space because it is recording what is reflected from a surface, and that can be changed by factors other than the actual color of an object. A bigger color space gives us more room for color balancing.

If you view the raw file of a camera, like figure 2, you kind of feel the saturation is way too low. In other words, the picture looks dull. This is because of the wide response bandwidth of camera sensors (when compared to phosphorous emittance bandwidth). You should not be concerned at all, because this means your camera is recording more colors than your monitor can display, and that's a good thing.

3 Profiles

A profile (ICC profile) is a file that specifies how a device responds to colors. An input profile takes an input image, and transforms it to an “internal color space” that is capable of representing all colors (such as xyL). An output profile transforms an image represented in the internal color space back into an image file to be rendered on an output device.

In order to get the most out of your expensive digital camera, you will need to find an input profile that specifies how the sensor responds to colors. There are several ways you can do this:

- Beg your camera manufacturer for one. I don't know how these big manufacturer think, but no profiles are supplied with even expensive digital cameras in the order of a few thousand dollars!
- Purchase workflow software. Most workflow software (from capture to print) includes ICC profiles for most expensive cameras. You may be looking at spending from US\$100 to US\$500.
- Install evaluation versions of commercial workflow software, and extract the ICC profiles. Besides being amoral, there are other reasons you may not want to do this.
- Purchase profiling software and make your own profiles. Profile Prism is one. It costs about US\$60. Most commercial (non-free) profiling software includes a nice GUI. If you are command-line-phobic, this may be the best option. Keep in mind, however, that the EULA of most commercial profile programs specifically says you cannot share the generated profiles!
- Use free software to do it all by yourself. You need to know more about using command-line tools, but you get to control what you want to do with the generated profiles.

4 Show Me the Goods (Creating Profiles)

You are reading this to know *how* to get your pictures to look right. Here is the meat!

4.1 Getting the Tools

Get the tools and documentation from <http://www.argyllcms.com/>. These tools are available for major operating systems (Linux, Mac OS X and Windows). Best of all, these tools are free!

4.2 Getting an IT8.7/2 Target

I suggest you visit <http://www.targets.coloraid.de> and order one. For digital cameras, you should use either R1 or C1. I use an R1 (glossy) without any problem. Please read on for tips to use a glossy target.

Kodak, Agfa and other film companies also make IT8.7/2 targets. I found that my Kodak target not to be made as well as the ones mentioned above.

When you get a target, be sure to get the data files as well. The data file describes the color of each patch on the target as measured by a precision colorimeter. If you order from Coloraid, the data files are included.

4.3 Photograph the Target

I recommend at least taking a picture of the target under a noon sun. The sun is a full-spectrum light source, and most pictures are taken under the sun. You can take pictures of the target under other lighting conditions, too. However, at least have one taken under noon sunlight.

If your target (for example, an R1 target) does not have any firm backing, you should attach it to cardboard, foamboard or some other backing. I found that masking tape works well. However, if you worry about damaging the target, you can put “buffering” paper between the adhesive side of tape and the surface of the target.

Next, select your longest lens that can give you a full-frame image of the target. For point-and-shoot cameras, you may need to use the macro mode. The reason to choose the longest lens is to give you more working room, and it also helps to make sure you do not cast shadow on the target. I use a 100mm lens, which gives me about three feet of distance.

Use a tripod whenever possible. Although we don’t need absolute sharpness for profiling, it is nice to free up our hands. A tripod also makes framing easier. Instead of adjusting the tripod, adjust the target instead. The target should be completely rectangular with no perspective tilt. Align the edges of the target with the edges of the viewfinder.

This next step is the key to use a glossy target. The main reason *not* to use a glossy target is glare and reflection. To overcome this problem, maintain an incident angle of 45 degrees between the target surface and the light source. Also, use a large black object on top to block any reflection. This means it is much more convenient to use a tripod and use the timer on the camera, because you need to place the black object directly behind the camera. A black foamboard purchased at an office supply store works great. Be sure that the reflection block does not cast shadow on the target.

Take a few shots with slightly different exposure settings (bracket). Most raw files are not affected by other post-processing options (such as color saturation, tone, sharpness, etc.). To be sure, however, read the camera’s manual and make sure the camera does not post process the raw image.

4.4 Conversion to TIFF

Most cameras do not record raw images in common formats. You’ll need a conversion tool that does not perform any processing to generate the necessary 16-bit TIFF file. I use `ufraw` (<http://ufraw.sourceforge.net>). Follow these steps:

- Open the image, like `ufraw b3fv9854.tif` (the Canon EOS 1Ds uses `.tif` extension for its raw files).
- Set both the input and output profiles to `sRGB`. This ensures no color space conversion takes place.
- Use a linear curve (no curve correction). ICC profiles work best on linear files ($\gamma = 1.0$).
- Reset to default settings.
- Put the white balance selection spot on GS0, then select spot white balance. I suggest you use a relatively large sampling square for accuracy.
- Adjust exposure so that there is no clipping (no underexposed or overexposed pixels).
- Click “Save As”, select 16-bit TIFF without LZW compression.




Figure 3: An IT8.7/2 target converted from raw file, using a profile created by Argyll CMS programs.

- Name the file accordingly, then click “Save”.

4.5 Commands (This is It!)

Let’s assume the TIFF file is called `it8.tif`. Now you need to extract the colors from the file:

```
scanin -v it8.tif /usr/local/cms/ref/it8.cht R050301.txt
```

Obviously, this command only works because of the following:

- The programs are installed where `PATH` includes. I actually uncompress the files into `/usr/local/cms/`, then use soft links to link all programs to `/usr/local/bin`. You can choose other ways to do it (especially if you use Windows!).
- Both `it8.tif` and `R050301.txt` are in the current directory. `R050301.txt` is included with the target from Coloraid.

This generates an intermediate file. The next command uses it:

```
profile -v -E"1Ds sunny" -qm -as it8
```

This creates the profile called `it8.icc`. You probably want to rename it and move it to another folder where you keep all the other profiles.

Be sure to read the full document from <http://www.argyllcms.com/doc/ArgyllDoc.html>. Both commands have many options that can help you diagnose problems or achieve better results.

5 The Finished Profile

Figure 3 is the result of applying the generated profile to the raw image. The output is in sRGB space, which means it is only suitable for online viewing on a monitor. Nonetheless, it is very close to the original target when viewed under 5000K lighting.

The ICC profile is located at `/plays/1ds-day-linear.icc`. Since this is generated by a GPLed program, you can freely download it, use it, and/or redistribute it.




Figure 4: Fluorescent lit picture processed by a day (sun) profile.

6 Additional Profiles

If you take pictures under different illumination conditions, you may want to consider taking more pictures of the target and generating different profiles. The “sun” profile is good enough for most naturally lit situations. However, it assumes a relatively broad and wide illumination spectrum.

Some light sources are not based on broad and wide spectrum. Fluorescent bulbs have narrow bands spikes. Such spikes mean how a camera sensor sees in fluorescent bulb lit conditions can be quite different from how we do it. Worse yet, the spiky characteristics of fluorescent bulbs can be different from one model to the next. It may be necessary to generate one profile for each type of fluorescent bulbs. I have found that the sensor of the EOS 1Ds has a reception bandwidth that is wide enough to handle most fluorescent bulbs.

Figure 4 is a picture of the IT8.7/2 target shot under fluorescent light, but “fixed” with a daylight profile. Figure 5 is the same picture, but fixed with a profile generated for that fluorescent light. The differences are subtle. Most of the differences are visible on column 19 and column 15. The “sun profile” turns yellow greenish, and pushes blue a little to the violet side.

http://www.palagems.com/gem_lighting2.htm has some really scientific analysis of various artificial sources of light. You can visually see the spiky nature of fluorescent bulbs.




Figure 5: Fluorescent lit picture processed by a matching profile specifically made for the fluorescent light.

7 White Balance

Assuming we have the proper ICC profile for each unique lighting condition, colors can still be off. This is because recorded “white” is absolute, and it depends on the color of the object as well as lighting.

The best way to get accurate white in the image is to include a neutral gray card in one of the images for calibration. You don’t need a gray card in each picture, just one in each unique lighting condition. Then, you can use the gray card to establish color balance (as well as luminance level).

8 Fine Tuning

`profile` supports a few ways to express the *transformation* of color spaces (from the input space into the standard space). Here is a quick summary of how each one fits a camera profile application.

8.1 Gamma Curve with Matrix

There are two options. `-ag` (lowercase) uses a gamma constant for each color channel, while `-aG` (uppercase) uses a single constant for all color channels. In either case, a matrix is used to transform from one color space to the standard color space.

This option is the quickest, and `profile` finds a “best fit” profile within seconds. Depending on the device, this option may very well be sufficient. If a digital camera has a sensor that is truly linear or logarithmic, then a gamma curve is the best fit.

The difference between `-ag` and `-aG` can be significant if the RGB sensors have different responses. You can usually tell. Generate a profile based on `-aG`, and see if the gray tones appear to have different color tints. It’s alright if they all have the same tint. However, if they have different tints, then you may want to consider using `-ag`.

If a sensor is not logarithmic (which includes linear), then you should consider the following sections. The Canon EOS 1Ds sensors are fairly logarithmic, and I get fairly good results with the `-ag` option.

8.2 Shaper Curve with Matrix

There are actually two individual options here. `-as` (lowercase) specifies a different shaper curve for each color channel, whereas `-aS` (uppercase) specifies a single shaper curve for all color channels. In either case, a matrix is used to transform from the original color space to the standard color space.

A shaper curve is a more arbitrary curve when compared to a gamma curve. In other words, a shaper curve can include more “twists and turns” to model the response function of RGB sensors. Consequently, if a camera has RGB sensors that are not truly logarithmic, shaper curves can produce better profiles.

That said, however, I found that a shaper-curve based profile performs worse than a gamma-curve based profile for an CanonEOS 1Ds.

You can experiment with this option if `-ag` or `-aG` does not produce a satisfactory profile. Sometimes, if `-ag` or `-aG` fails to produce a quality profile, you should consider shooting the target again. Inconsistent lighting, glare and etc. can pollute the IT8.7/2 image to the point that the captured bits are not longer representing how the camera sees

8.3 Lookup Tables

Lookup Tables (LUTs) are approximation of curves by many segments of straight lines. As a result, they are far more flexible than even shaper curves.

However, LUTs can easily overfit a sample of the IT8.7/2 target so that the profile may not be very useful for other images. Unless you can get the lighting to be absolutely perfect, this means LUTs may generate profiles that are not very useful in general.

One problem that I find with LUT-based profiles is that when pixels go beyond saturation, artifacts can appear. Normally, when you increase the brightness of a picture, everything will wash out and become closer to white. With LUTs, the wash out transition actually brings a pixel through various colors.

This is not a problem with the LUT approach, but rather bad LUTs generated to over fit flawed samples of the IT8.7/2 target. If you see a LUT curve that has dips, you know that there is a problem.

LUTs work best when there is no need to adjust brightness. If you can guarantee that the exposure of every frame is dead on, then LUT may be an option you can consider. However, if you need to adjust the brightness of a picture during raw file processing, I suggest you stay away from LUT based profiles.

LUT based profiles work best for devices that do not have brightness adjustments, such as printers. LUT based profiles also work great for scanners, assuming no brightness adjustment needs to be done.