2/27/08

# Introduction to Color Management

We have all walked into a video store before and seen a bank of televisions sets broadcasting the same exact signal, and yet the colors, contrast and brightness look different from one TV to the next. Or, you have viewed two printed photos or brochures generated from the same exact file, yet they don't come close to matching. Perhaps most frustrating of all, what you see on the monitor looks different from the print coming off your inkjet printer. Color Management strives to solve these inconsistencies.

Color management in the digital world is a methodology for delivering consistency, repeatability and accuracy throughout every stage of the digital imaging process ~ from image capture (digital camera or scanning), monitor display, editing and file preparation to final output, whether sent to a printing press, inkjet, laser, digital press, photo processor, monitor or other device.

Unfortunately, the Internet and most monitors are *not* color managed (with few exceptions), so websites can look very different from one computer to the next. Images that look perfect on my color managed work station may appear washed out, dark, too blue, too yellow or overly saturated to another viewer. The tools are either not yet available or not widely implemented on the worldwide web. Thankfully, this is not the case in the world of print, or even on individual work stations.

## The Complexity of Human Vision

Human vision is complex. We do not really see an object, but the light reflected from that object, which then passes through our lens, strikes our retina, where light sensitive rods and cones are stimulated. The cones are color receptors and need adequate light intensity to distinguish color, and the rods are more black & white, but able to function in much lower light. Once stimulated, the rods and cones pass this information through the optic nerve to the brain, where the information is processed and “seen” as an image. Pretty miraculous. But that is just the tip of the iceberg. It gets a whole lot more complex.

Eyesight was an adaptation to survival. In fact, scientists say we see differently from many other species, and our acute color vision is one thing that sets us apart from many animals. When we look at a scene closely, our visual system does its best to differentiate and separate colors. For example, we might hear a sound in the bushes, so we focus closely on the greenery, and our visual processing system actually increases the contrast, saturation and differentiation between close shades of green in an attempt to isolate the source of the sound. Not only that, the aperture of our eye changes rapidly to accommodate different brightness levels, and our focus changes rapidly. All this happens instantaneously without our effort or awareness. This partially explains why a photograph of a particular scene may look flat and lifeless compared to the ‘way we remembered it’. The camera records the scene objectively just once, but our visual system is ‘ scanning and processing’, and it varies from one scene to the next. It may tone down colors and contrast on a brilliant day, but it may ‘pump them up’ on a flat, overcast day.

As if that weren’t enough, the human visual system also has its own built-in color adaptation system ~ a sort of ‘automatic white balance’. We have all seen or taken pictures in mixed lighting situations, for example, daylight and incandescent light. A picture taken from inside your house using incandescent lights will look fine, but daylight through the window will look very blue and harsh. Conversely, a picture taken outdoors will look perfectly balanced to our eyes, but home windows lighted from inside with regular bulbs will appear yellow. These mixed lighting systems are confusing to the mind, since artificial lighting is a recent invention, and our visual system did not evolve under ‘mixed light’ situations. It did, however, evolve to compensate for varying light qualities at different times of the day. This brings us to the subject of color temperature.

## Color Temperature

The quality and ‘color temperature’ of daylight varies throughout the day. A sunset bathes us in ‘warm’ colors as the light passes through more of our atmosphere and longer wavelengths of light predominate. At noon, light is overhead and passes through less atmosphere and we have a ‘normal’ balance of long and short wavelengths. If the day is heavily overcast, many of the longer, warmer wavelengths are filtered out, so that shorter wavelengths now predominate, resulting in a colder color of light.

The color and quality of light also varies in the artificial light world. Candlelight is one of the warmest sources of light with very long wavelengths. Incandescent is also very warm, but not as warm as candlelight. Fluorescent light is a bit cooler still, and tends to have a greenish tinge, though this varies with the type of fixture. Electronic flash is cooler still.

Color scientists, in an attempt to simplify, assign a ‘color temperature number’ to a given light source. This single number generally tells us whether we have warm, ‘normal’, or cool light, but it leaves a lot of information out. Still, it is a helpful guide. Color temperature is plotted on the Kelvin scale, and is based on the color of light irradiated from a ‘black body radiator’ when heated to various high temperatures. This color temperature is abbreviated as °K (degrees Kelvin).

When the black body is heated to a temperature of around 1800-2000 °K, you will see a very warm, reddish yellow light similar to candlelight. When heated to 2500-2800 °K, it will be close to incandescent light. As it is heated to higher temperatures, the color of the light becomes whiter and then bluer. The below chart provides *very rough estimates* of some color temperatures for different light sources.

| ***Approx. Temp °K*** | ***Light Source*** |
|---|---|
| **1500-2200 °K** | Candles or fireplace |
| **2500-2800 °K** | Incandescent bulbs |
| **3200 °K** | Photographic ‘hot’ lights |
| **4000 °K** | Fluorescent lights |
| **5500 °K** | Daylight on a clear day at noon |
| **6000 °K** | Electronic flash |
| **6000-7000 °K** | Overcast day |
| **7500-9000 °K** | Shade or heavy overcast (bluish storm clouds) |

## Spectral Distribution Curves

As useful as these single numbers are, they over simplify the nature of light and can be misleading. Light is more accurately described using *spectral data*. Human vision is only responsive to a small portion of the electromagnetic spectrum (the same spectrum that gives us infrared, ultra violet, microwaves, etc). We generally are able to see only those wavelengths that fall between about 380-730 nm (nanometers) in the electromagnetic spectrum. Hummingbirds and bees, for example, can see further into the UV spectrum and use this to locate food sources.

By ‘slicing’ the visible spectrum into very small bands, it is possible to plot how much of each wavelength is generated by a specific light source. For example, we can say that a warm incandescent light source has a color temperature of 2550 °K, but what exactly does that mean? If that light is accurately measured using a good spectrophotometer, (as I have done below) the light can be plotted on a curve. To the right, is what we call the ‘spectral distribution curve’, which plots exactly how much of each wavelength goes into making up this light source. In this case, the light source is very deficient in the violet, blue and green portions of the spectrum, but is heavily weighted in the yellows, oranges and reds. The curve is smooth and does not have a lot of spikes (more on this later). The center plot represents this as a single point (ie, color temperature) on a visual plot called an L u’v’ chart. You can see that the plotted point falls in the yellow area. The “Light

Quality' portion to the left shows that this light source has a Color Rendering Index (CRI) of 99 out of 100, meaning that the distribution of the spectral curve is very close to the associated 'standard' established by the color scientists.

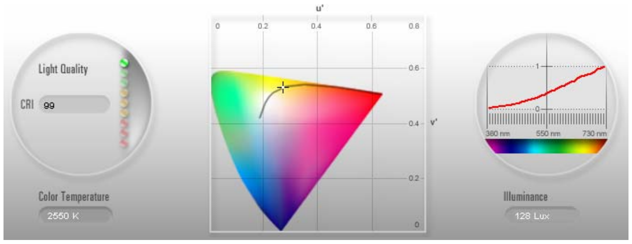

Spectral data for an incandescent light source

You could measure a different light source, and it may also be reported as 2550 °K and yet the spectral distribution curve may look very different, with blips, spikes and a completely different shape. Why would we care? Because, an object viewed under these two light different sources will appear different to our eyes, even though they are 'supposedly' both 2550 °K.

You recall that we don't really see an object; we see the light reflected *from* the object. If we look at a multicolored object in 'normal daylight' we will see what we have learned to call a *normal* distribution of colors. Reds, yellows, greens, blues, violets, and whites will look 'normal'. But, view that same object in the above light source and your blues, violets and to some extent, greens, will be subdued. Why? Well, the light source doesn't contain much of these shorter wavelengths, so if it isn't there to begin with, it can't reflect back to the eye. In fact, those cool colors will look darker and less saturated than they would when viewed under normal daylight. The longer, warm wavelengths, however, are available in this light source in abundance, so the reds, yellows and oranges will reflect back strongly. They will be light and bright since they reflect their energy back to your eye strongly.

Now view the same object under fluorescent lights (I use the term loosely, since there are so many different flavors of fluorescents with radically different spectral distribution curves). Fluorescent lights then to have a lot more 'spikes' in the spectral curve. The plot below shows one of the better fluorescent curves. The average color temperature is about 4503 °K, and the CRI (color rendering index, which measures it against a standard) is 92 out of 100. Very good, especially compared to most fluorescents. But, look at the spectral distribution curve on the right. It certainly is not smooth. It is somewhat weak in the reds, oranges, and violets, but has a lot of spectral power and spikes in certain wavelengths of yellow-green, green, and blue. But, overall, it is flatter than the incandescent curve we displayed earlier. Also, the central plot shows the 'average' color temperature moving closer to the white portion of the visual color spectrum, compared to the incandescent light source measured previously. So, the color temperature in °K does tell us honesty that this is a cooler light source and closer to daylight color temperature, but if fails to tell us which colors this light source will accentuate. The spectral curve, however, says it will over emphasize the greens, blues and yellow-greens, but short change the violets and reds.

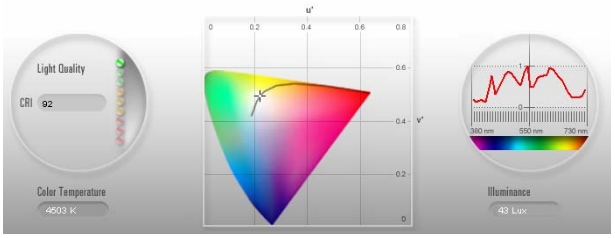

Spectral data for a fluorescent light source

## The Color of White

Notably absent in the above is a discussion of the color white. What we call 'white' is the sum total of all wavelengths reflecting back to the eye under 'normal' lighting conditions, which most color scientists have defined as midday light on a clear summer day (typically 5500-6500 °K). A bright white sheet of paper absorbs very little light energy and most of the wavelengths in the visible spectrum are reflected back to the eye and perceived as white. (It is also why people wear white in a hot climate, since light is reflected instead of being absorbed). Black is the opposite; most of the wavelengths in the visible spectrum are absorbed by the black object, so little or no light makes it back to the eye. (It is also why a black car with black interior gets warmer in the sun).

If you hold that sheet of white paper in the noon day sun, we see it as white. Take that same piece of paper into a dark house and view it under your incandescent lights and it still looks white to our eyes. But, we showed earlier that the color temperature of our incandescent light is a warm, yellow orange in color. It is deficient in violets, blues and greens, so if those wavelengths are unavailable in the light source, and the yellows, oranges and reds are over represented, how is it that the paper still appears white?

It is called Chromatic Adaptation. Remember, human vision involves a brain which processes information and represents it as an image only after processing. The human brain adapts to the light source and 'knows' the paper is white, so it processes the data differently. How does our brain know the paper is white? It looks at the entire scene and 'calibrates' itself using the brightest, lightest object in the field of vision as being white (an over simplification, but useful for our discussion), then applies a global adjustment to other colors. The brain does this automatically and instantaneously, without any effort or awareness on our part.

This works fine under conditions where we have only one light source. But, let's take a walk outside as it is getting dark. The light outside is relatively cool daylight. The street lights come on, and some have a pink-purple cast (mercury vapor lights), and when we look in someone's house window, it is distinctly yellow. Our eyes calibrate to this dominant daylight source of light, so the other light colors use daylight as a reference. When we are *inside* our house in the early evening, our eyes calibrate to our incandescent lighting, since it is the predominant light source.

Human vision adapts quickly and automatically. It has a powerful processor (a brain) that has evolved over eons and knows how to process data quickly and accurately. Mixed light sources are a recent invention, so the brain can become a bit muddled in these situations, but it still does a creditable job.

## How Cameras See Light

Unfortunately, cameras don't see as we do. Whether digital or film, they record wavelengths and photons objectively. They cannot pump up contrast selectively as we do, tone things down when they are too intense, nor can cameras selectively ignore shiny oily skin and 'hot spots' the way human vision does. Human vision compensates. Cameras do not. Cameras are becoming more sophisticated, especially in the digital age. Most digital cameras have an 'automatic white balance' feature that reads the light and attempts to do what our brain does, ie, calibrate itself and adjust colors based on its assessment of the current color temperature. But, camera AWB can be easily fooled and give incorrect readings, especially in unusual lighting situations, or when a large block of color predominates. Many cameras allow you to *manually* select White Balance based on flash, tungsten, fluorescent, shade, daylight and other lighting conditions, which is very helpful, but these selections are approximate. And many cameras allow you to take a photo of a white or gray card to set a 'custom white balance', which is usually the most reliable, at least in static lighting conditions. The film equivalent is selecting daylight or tungsten film, or using color correction filters to 'normalize' the lighting to match the film's white balance expectations.

So, why doesn't the eye just adapt to the print like it does in a real life lighting environment? Let's say you print an un-retouched image taken under greenish fluorescent lighting, but left your camera's white balance set to daylight. Your family will look like a bunch of Martians (no, I won't go there!). Why? The camera faithfully records the wavelengths of light reaching the sensor or film and your printer dutifully prints this on paper. When you look at the print, your brain calibrates itself to the predominant light source, usually trying to find something neutral, or preferably bright white, then makes it's global adjustments. If your print has a white border, the brain sees it as white, but the faces in the print are green. If you trim the white borders from the print, things improve somewhat (unless you lay the print on something white, like a white counter top), since the 'white reference' isn't butting right up to the print, but your brain still knows what color light predominates. But, even if you view the print in isolation, (for example against a black background), you may still see the people as green, especially if there are 'known' neutral or white elements in your photograph. Even with a greenish light source, pure white in your photograph will still print as pure white on paper, and your eye will latch onto that for self calibration.

When viewing your photo, if there are other white/light objects in your field of vision, such as light sources, a white wall or door, the eye will tend to use the most prominent, whitest and brightest objects as a "white reference' to calibrate and adjust our perception of the overall scene, including the photograph. It's similar to setting a custom white balance in a higher end digital camera, but much more sophisticated.

## Color Management's Job

It is the incredibly difficult and complex job of Color Management to reconcile all these issues and deliver an image that the brain sees as accurate, lifelike, color balanced and with a full tonal range. Admittedly, it doesn't do it perfectly, but it does an amazingly good job if implemented properly.

There are many reasons why perfection is not attainable. Much can be attributed to the incredible sophistication and subjectivity of our amazing brains. No computer or processor can hope to attain that level or to sense and think like we do. So, our models and processors are limited and less than 100% accurate, though very good. The rest can be explained by the limitations of monitors, prints, transparencies, dyes, pigments, reflected vs. transmitted light, and laws of physics.

## Monitors, Printers and Computers ~ Oh my!

Not only is human vision complex, subjective, adaptive and sophisticated, it has incredible range. We can see a tremendously wide array of colors and shades. We can see from deep shadows to ultra bright highlights. Human vision adjusts to light, darkness, color temperature and other factors almost instantaneously. When we survey a scene, our eye moves extremely rapidly from one object to the next, adjust aperture and focus, processing the data, adjusting saturation and differentiation, etc. We take a scene apart piece by piece, and the brain scans and processes the picture accordingly.

The camera takes a single exposure in a single moment of time, with a single aperture with single shutter speed. All digital and film cameras have limited dynamic range and color gamut, far more limited than human vision. Yes, some cameras and films can see beyond human vision, but unfortunately, that doesn't do us much good unless we can see them (or unless we somehow convert invisible light into something visible, as is done with infrared or x-ray photography).

Even if they could capture our entire range of vision, we still have to deal with monitors, transparencies and printers, since we cannot see the actual photons on the sensor.

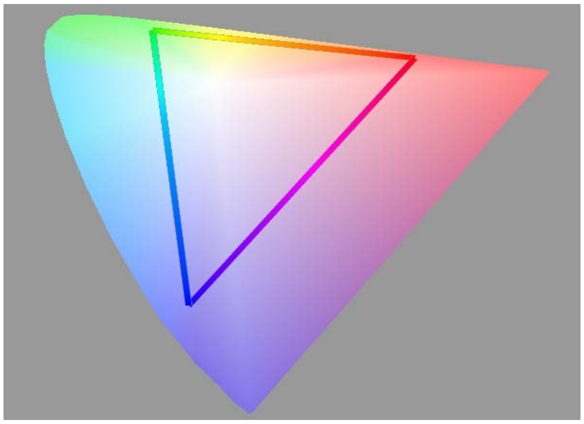

*2D representation of human vision (pastel) vs. displayable colors of an 'average' monitor (triangle)*

Monitors use red, green and blue 'transmitted light' and blend them together to create the illusion of a continuous tone image. If all three colors are firing simultaneously at full intensity, we see white on the monitor. (Red Green and Blue are called the *Additive Primaries*, since when *added* together, they form white). If all are turned off, we see black. But, the white and black we see on screen are less intense and cover a smaller dynamic range than what human vision is capable of seeing in the real world. Red, green and blue components make for very good display of those particular primaries, but don't do as good a job displaying other colors, such as yellow, cyan, magenta and other intermediate colors. The monitor has 'limited color gamut', which in essence means that many colors human vision is capable of seeing cannot be reproduced on a monitor.

Prints, on the other hand, rely on 'reflected light' using cyan, magenta and yellow dyes or pigments (opposite colors of red, green and blue) as the base colors for creating the illusion of continuous tone. White light reflects off the printed image and each color in the image absorbs (or subtracts light, hence the term *subtractive primaries*) some of the light. The light that reflects back to our eye is the color we see. Due to the impurity of the cyan, magenta and yellow dyes and pigments used in printing, it is difficult to get good blacks, so black is usually added to the mix. Black is also necessary for reproducing clean, crisp text and lines. This is the classic CMYK used on a printing press and many printers for printing full color, 'continuous tone' images.

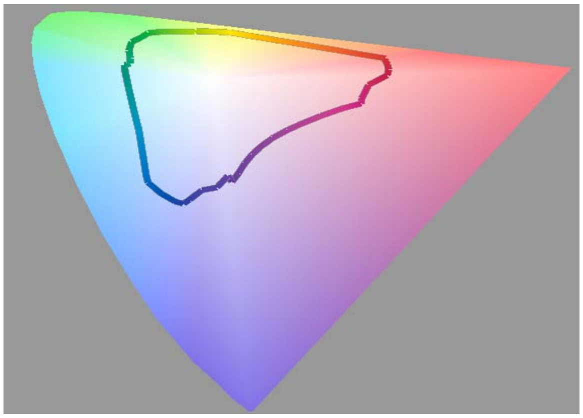

*2D representation of human vision in pastel vs. typical printable colors on matte papers with ink*

Unfortunately, printed images are even more limited than monitors in terms of color gamut (reproducible colors) and dynamic range. Where RGB monitors do better with reds, greens and blues, CMYK printers generally do better with cyan, magenta and yellow. In other words, the gamuts are different, and in fact, overlap. Printed media has another disadvantage – limited dynamic range. The brightest white in the print is dictated by the brightness and whiteness of the paper base, which is duller than what one can get from a monitor. Blacks are also limited by the paper's 'ink limit' which dictates just how deep a black you can achieve. This varies depending on the combination of paper and the ink used. Paper type and quality can make a huge difference on the final print.

It is becoming more common for printers to use additional colors in an attempt to increase the color gamut of the print. Light colors (light cyan, light magenta are examples) help to improve the creation of light pastels in print, since they can provide greater coverage and density, without resulting in too dark a color. Reds, greens, blues and other colors are sometimes added to provide more saturated primaries and further extend the gamut of the print. To improve neutrality and smoothness of neutral gray tones, additional shades of gray are being added to some printers. This is important, since our visual system is extremely sensitive to color casts in light, neutral tones. While these all definitely are excellent advances, we still have to deal with a very limited media….paper. The dynamic range is short, so we need to use every bit of it that is available.

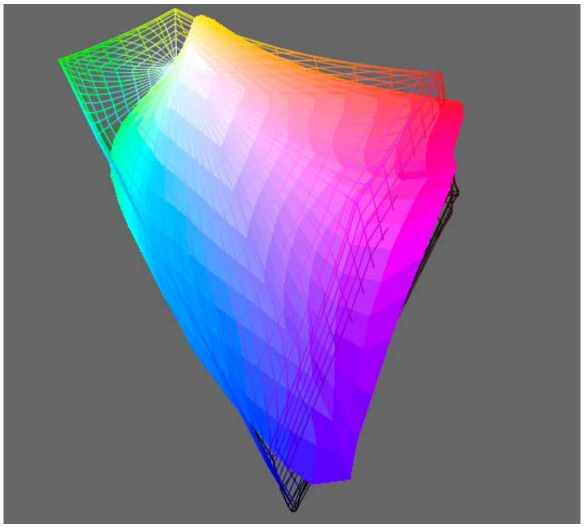

*3D plot showing Adobe RGB gamut in wireframe vs the gamut of ink on glossy media(solid colors)*

As shown above, although, RGB monitors tend to have larger *overall* color gamuts than prints, especially in reds, greens and blues, paper and ink can peek outside a monitor's gamut, notably in cyans, magentas and yellows. Shapes differ.

Printing on transparencies can improve dynamic range, since we now have the ability to transmit a bright light from behind and provide a more brilliant white. Slide shows tend to be more 'colorful' for this reason, and also because we have darkened the surrounding room to boost contrast.

So, how does color management handle all these differences?

## CIE L*a*b* Color Space

In 1913, color scientists formed the International Commission on Illumination, which is usually know as the **CIE** for its French name, commission internationale de l'eclairage. They decided to tackle this incredibly difficult and complex subject we call human color vision.

Essentially, they created a definition of the visible spectrum through intensive testing. Of course, not all people see the exact same color the same way, so they averaged their test data, obtained from a very large sample of people, and created an 'independent color space' called CIE L*a*b* (henceforth, I will use the abbreviation ***Lab*** for short, though the correct notation is always supposed to be CIE L*a*b*).

The Lab 'color space' was designed to represent the entire range of human vision, from the pure white to the total black, bright saturated colors to pastels to neutral grays, and all colors in between. You can think of it as a three dimensional plot of every color we can see. It was also designed in such a way as to be described mathematically and unambiguously. CIE also attempted to make it perceptually uniform, so a difference between two color numbers would be perceived by a person as being the same difference throughout the model. For example, if a difference of 1 ΔE unit (pronounced *Delta E*, which is one mathematical model used to calculate and describe color differences) represented the smallest visible difference between two shades of gray, that numerical difference of 1ΔE would apply whether you were comparing black to the next perceivable shade of dark gray, or a white to a light gray. It isn't perfect, but it is currently one of the best widely used measures available. As mentioned earlier, our vision is much more sensitive to small changes in light, unsaturated colors, such as neutral grays, skin tones, near whites, etc. But, despite its imperfections, the Lab color model is still very helpful and effective in helping us to achieve good control and reproduction of color.

Lab was designed to emulate the way human vision works. Lab has three components, and is therefore called a 'tri-stimulus' color model. (Red/Green/Blue and Hue/Saturation/Brightness are two other popular tri-stimulus color models.) The "L" value establishes the lightness of a color, with 0 being solid black and 100 being bright white. This is pretty straight forward. The a* and b* values are less intuitive.

Human vision sees color in 'opposing terms'. In other words, a solid color can be yellow or it can be blue, but it cannot possibly be both at the same time. So, if we locate a brilliant, saturated blue on one end of a see-saw, and intense, saturated yellow on the other end, and neutral white (or zero saturation) in the center, we have just described the b* scale of the Lab color space.

The a* scale puts fully saturated green (closer to teal green) on one end, saturated red/magenta on the other, and again, neutral in the center.

For our purposes, Photoshop uses the following scale.

L* 100 is pure white
L* 50 is middle gray
L* 0 is black

a* maximum green (teal) value is -128
a* maximum red/magenta value is +127
a* neutral value is 0 (neither red/magenta or green, but right in the middle)

b* maximum blue value is -128
b* maximum yellow value is +127
b* neutral value is 0 (neither blue nor yellow, but right in the middle)

In theory, any color in the human visible spectrum can be defined using 3 numbers. For example, middle gray would be defined as 50L, 0a, 0b, since the lightness component is right in the middle, i.e., neither white nor black, and both the a* and b* numbers are zero (ie, neutral). 80L, 0a, 0b would be a light gray.

A specific color of red could be described as 60L, 80a, 70b. This color is a little lighter than middle gray since it is higher than 50, the a* value is definitely more red/magenta than green, and the b* value is more yellow than blue. Mixing red/magenta and yellow gives you a red color.

The closer the a* and b* values are to the center, the less saturated they are. If a* is zero, the color being described has no red/magenta or green bias and is neutral on that scale, which is very convenient. If the b* value is 0, then the color being described has no yellow or blue bias. If both a* and b* are 0, your color is dead neutral, since all color information is contained on the a* and b* scales (also called channels). You know that color is totally desaturated so it is either black, white or a shade of gray, as determined by the L* value.

The further the a* and b* values move toward either extreme, (ie toward -128 or toward +127) the more saturated they become. Positive values are warm (yellow or red/magenta), and negative values are cool (blue or green).

Of course, some colors that can be mathematically described in Lab lie outside of normal human vision, and some colors are impossible. For example, the color 0L, 0a, 127b would be a color that is totally black (since L* is zero), but is at the same time intensely yellow. That is an impossible color. Sorry, but black is black.

Nonetheless, Lab is the basic framework upon which much of color management rests. CIE created Lab as a "device independent color space" to define color. By device independent, they mean that the color is not tied to a specific device, such as a printer or monitor. Lab is a great reference space for translations of colors, since it purports to accurately describe human vision and a given color is always unique and specific. A Lab color of 50L, -25a, 80b (which happens to be a medium density yellow green) is always that color, regardless of the device it is printed or displayed on (thus it is device independent). At the start of this article, we mentioned different TVs that display colors differently, though the originating signal is identical. All devices, therefore have their own inherent characteristics and properties and cannot be relied upon to produce color accurately, at least not without some help.

## What Color is that Sweater?

No doubt, you've all had the experience of saying to someone, "I love that dark blue sweater." They answer that the sweater isn't blue, it's black, or deep purple or whatever. Part of this is attributable to differences in individual human perception, or our expectations, or the lighting, etc. The same thing happens all the time between photographers, graphic artists, art directors, commercial printers, etc. Someone says they want the dress to be more red. What does that mean? What color red, and what is the saturation and brightness?

These terms are misleading and ambiguous, and subject to our preferences, conditioning, and expectations. But, there is no ambiguity if the art director says, "I want that color to be 50L, 60a, 40b red. I don't care what device you use to display or print that dress, I want it to be that specific color." And color management can help you do that….*as long as that color falls within the color gamut of that monitor, paper or device used to reproduce it*. And if the color does fall outside of the gamut of the paper or monitor, color management can help us attain the closest color that output device is capable of delivering.

## Device Dependent Color

We described CIE L*a*b* as a *device independent* color space, since it is an absolute color space modeled on normal human vision and the boundaries inherent therein. In contrast, printers, monitors and all other output devices are **not** device independent. Remember our walk into the color TV store? All those TVs were receiving identical signals, and yet their colors varied all over the board. The color of red being displayed was dependent on the phosphors or LCD crystals, the design of the TV, the circuit board, the color of the backlight, software, etc. So, you could send the same RGB signal to two different monitors and end up seeing a different color. The color is *dependent* on the device, in this case a monitor.

## How Color Management Works

Color Management is like traveling in a foreign country, where you need a skilled translator to help you understand and convey meaning clearly and accurately. When traveling, your translator needs to be able to speak both languages fluently. A poor translator results in a poor translation and misunderstandings. A great translator conveys meaning precisely and handles the nuances of language and meaning effectively. But even the best translator cannot always convey the meaning exactly. For example, the English word 'home' has no exact equivalent in Spanish. The closest Spanish word is 'casa', or house. But, that misses the nuance and meaning of the English word home.

In color management, the universal translator is CIE L*a*b* (or Lab) color space. It is the model of human vision, imperfect though it might be. Still, it is an excellent and effective translator. But, some colors (like some words) cannot be translated to the final destination. If a color is out of gamut of a device, there is no way to print that color.

The International Color Consortium (ICC) created a methodology whereby different devices could communicate with one another accurately and predictably, within the limitations of those devices. This methodology uses ***ICC profiles***. An ICC profile is nothing more than a map or description of a device's color space, its boundaries, limitations, maximum black, white, etc. It also contains information which allows precise definition of those colors. It can be "table based" or "matrix based". A table based profile has a 'lookup table', akin to a dictionary that translates words between languages. The larger the table, the larger the profile, and generally, the more accurate the profile. A matrix based profile contains curves to describe the color space and translate colors. Most printer profiles are table based, and monitor profiles can be either matrix or table based profiles. There are also 'working space' profiles for editing files, almost all of which are matrix based.

Let's start with a file on your computer's hard disc. At this point, it is simply a digital file full of 'numbers'. If we open a file in an image editing application, such as Photoshop, we will see the image displayed on the monitor. How do we know we have accurate colors?

Using the ICC model, we need two languages and a translator in the middle so the different devices are able to communicate and understand one another. Digital camera files start life as RGB files, and we will assume our file remains in RGB file for this example. The file in Photoshop contains three channels corresponding to Red, Green and Blue, as if you shot three separate, but identical images on B&W film through individual red, green and blue filters. We can edit this file to our heart's content, but we really need to decide what language it is speaking (otherwise our poor translator will get confused). This is where the 'working space' (also called the editing space) comes into play. If you have a color that is 200R, 50G, 50B, you can be sure this is a red, but precisely which color of red? How bright or dark is it, what is its exact hue and saturation? The working space defines boundaries and colors so we can pinpoint exactly what color this is supposed to be.

Working spaces also have two other indispensable features; they should be *gray balanced* and *perceptually uniform*. Gray balanced means that when Red, Green and Blue are all equal the color is totally neutral (completely desaturated). This is very handy during editing and helps us to make sure neutrals, blacks and whites are indeed neutral. Perceptually uniform refers to retaining smooth, incremental changes as the numbers change. For example, if we divided a smooth black to white gradient (256 steps in an 8 bit file) into 32 equally spaced steps, they would be 8 units higher in each step (256÷32). In RGB terms, the pure black step would be defined as 0,0,0, the next lighter step would be 8,8,8, the next 16,16,16 and so on, until you reached pure white at 256,256,256. If everything were perceptually uniform, our eye would see each step clearly and distinctly, and each step would be perceived as being equally spaced. The same would be true when it came to changes in color. This is perceptually uniform means.

Common RGB working or editing spaces include sRGB, Adobe RGB, ColorMatch RGB, Apple RGB, ProPhoto RGB, etc. (These are all ICC profiles, by the way and are widely used standards in the graphic arts field). The color 200R, 50G, 50B is a different shade of red in each of the above color spaces, so defining a file with an editing space is important. We'll get into different editing spaces later and discuss the advantages and disadvantages of each. But, just like when speaking a language, we need to know that the word 'casa' is a Spanish word, just as the words 'home' and 'house' are English words.

Now, we want to display that file accurately on your monitor. How do we translate from our file to the monitor? Remember, as with languages, we need two languages and a translator in the middle. The monitor needs its own ICC profile (ie, language) that defines its color gamut, outer boundaries, and is then able to describe those colors accurately. Monitors come with 'generic' profiles supplied by the monitor vendor. These profiles are better than nothing at all, but they describe an average monitor in new condition, so they will not be very accurate for a specific monitor. Also, monitors drift over time, and age, supply voltage, and other factors can cause further changes. So, it is best to create a custom monitor profile that accurately describes *your* monitor on *your* system. For the moment, we will assume you have such a profile.

Now we have two languages and a translator. The file's *language* is the working space, as described by the working space ICC profile (sRGB, Adobe RGB, etc). We'll assume for this example that it is Adobe RGB. The *monitor's language* is described by the custom monitor profile, which is also an ICC profile. The translator's native language is a device independent color space, such as Lab, or a variant thereof (sometimes xyz color space, but the function is the same). So, when sending our 200R, 50G, 50B red color to the monitor, what happens? The file sends 200R, 50G, 50B to the translator (Lab or xyz). It also informs the translator that its language is Adobe RGB. The translator now knows the precise definition of each color it receives from the file, including this red color. Now, the translator opens a dialog with the monitor, asking for its language and translation table. The monitor provides its custom profile, which contains all the necessary information laid out according to the requirements in the ICC specifications. Then, your translator gets to work. It translates (ie, it *converts*) the file's 200R, 50G, 50B red from Adobe RGB language into its own native language (Lab or xyz). It then retranslates (*converts*) from its native language into the monitor profile's native language using the tables or curves in the custom monitor profile. The monitor then receives instructions from the translator that it will totally understand, since it is now in its native language, asking it to display this precise shade of red.

This process involves *converting* the numbers in the file for monitor display (just like the English word 'home' is converted to the Spanish word 'casa'). The actual numbers in your physical file remain unchanged, but in order to display accurately on the monitor, adjustments need to be made. This is done in the background in the blink of an eye. To correct for your monitor's inaccuracies or bias, your original red may be converted from 200R, 50G, 50B to something like 215R, 42G, 57B. So, the numbers in your file are 'converted' to the new numbers *on the fly* and displayed accordingly. This happens with every single color in your file, and it happens very fast. If our ICC profiles are accurate, the color displayed on the monitor will be the exact shade of red we specified in our file.

So, whenever one device speaks to another, some translation (conversion) is required *unless* they are already speaking the exact same language *and* dialect. In the previous example case, both languages were RGB, but the dialects are different ~ Adobe RGB (the file's working space profile) and Joe's_NEC_Monitor.ICC. The translator is device independent and just translates colors using the independent model of human vision created by CIE.

Some colors in the original file may fall outside the gamut, or capability, of your monitor. The color management engine in the software must handle these anomalies, which we will discuss later.

What about sending a file to a printer? The overall process is exactly the same. We need two languages (ie, ICC profiles) and a translator (Lab). The file's language is still Adobe RGB. We need a profile for our printer/ink/paper combination that accurately describes what the printer can do with that paper and inkset. Then, the translator again does a conversion so that our red color, as specified in the file, can be reproduced as accurately as possible on output. If the printer profile is accurate we will get accurate output *within the limitations of the printer/ink/paper combination*. (Remember, there is no word for home in Spanish, so the translator gets as close as possible). So, if our red color is within the printer's color gamut and *can* be accurately produced, it will be if our profiles are accurate. As before, a physical translation occurs. 200R, 50G, 50B may become 191R, 38G, 61B, but the translator knows what needs to be done to get the right color red in the print, by consulting the lookup table in the printer profile. (The actual conversion into dots of ink is handled by the printer driver or RIP, but it is part of the overall equation.)

In fact, you could email your Adobe RGB file to me, and even though I have a different monitor and printer, I can will see the same image on my screen and be able to output the colors accurately to my printer, as long as my monitor and printer profiles are accurate. The numbers sent to *my* monitor and *my* printer will not be the same as those sent to yours, but we care about accurate colors and densities, not numbers. The translator (the same, independent Lab or xyz color space), knows we speak slightly different languages, but he understands all dialects and makes an accurate translation.

Of course, translations are only as good as the profiles themselves. Well built, highly accurate profiles provide great results. Poorly built, inadequate profiles yield sub par results. You can have a great printer profile, but if you are editing a file by eye using your monitor as a guide, and if your monitor profile is way off, you will end up changing the numbers in your file so they don't jive with what you see. We're back to the TV store with each TV having a different appearance. Which is correct? When you print, the color is unlikely to match. In this case, the problem is the monitor profile, not the printer profile.

And of course, we always have to deal with the fact that different devices have different color gamuts, not only in size, but in shape, as illustrated earlier. So, color management has to wrestle with these issues and offer a suitable solution. More importantly, the solution has to be transportable, from one computer system, monitor and printer to another. Standardized ICC profiles and Lab translators are currently the most widely used solution today.

This article's intent was to discuss some of the more relevant elements of human vision and color management, and how they fit together in broad brushstrokes. Details on how to create good profiles and how to use them will be discussed in the next article of this series.

Even though this paper is lengthier than I had originally anticipated, it remains only a cursory introduction to the concepts of color and color management. Hopefully, it has provided some clarification to a very deep, complex and confusing subject.

Lou Dina