- OVERVIEW OF COLOR MANAGEMENT -

"Color management" is a process where the color characteristics for every device in the imaging chain is known precisely and utilized to better predict and control color reproduction. For digital photography, this imaging chain usually starts with the camera and concludes with the final print, and may include a display device in between.

Many other imaging chains exist, but in general, any device which attempts to reproduce color from another device can benefit from color management. As a photographer, it is often critical that others see your work how it is intended to be seen. Color management cannot guarantee identical color reproduction, as this is rarely possible, but it can at least give you more control over any changes which may occur.

CONCEPT: THE NEED FOR REFERENCE COLORS

Color reproduction has a fundamental problem: different color numbers do not necessarily produce the same color in all devices. We use an example of spiciness to convey both why this creates a problem, and how it is overcome.

Let's say that you're at a restaurant with a friend and are about to order a spicy dish. Although you enjoy spiciness, your threshold for it is limited, and so you also wish to specify a pleasurable amount. The dilemma is this: a "mild" degree of spiciness may represent one level of spice in Thailand, and a completely different level in England. Restaurants could standardize this by establishing that one pepper equals "mild," two equals "medium," and so on, however this would not be universal. Spice varies not just with the number of peppers included in the dish, but also depends on how sensitive the taster is to each pepper. "Mild" would have a different meaning for you and your friend, in addition to meaning something different at other restaurants.

Peppers per Portion		You	Your Friend
	→	Hot	Hot
	→	Medium	Hot
	→	Medium	Hot
	→	Mild	Medium
	→	Mild	Mild
	→	Mild	Mild

To solve your spiciness dilemma, you could undergo a one-time taste test where you eat a series of dishes, each containing slightly more peppers (shown above). You could then create a personalized table to carry with you at restaurants which specifies that 3 equals "mild," 5 equals "medium," and so on (assuming that all peppers are the same).

Computers color manage using a similar principle. Color management requires a personalized table, or "color profile," for every device which associates each number with a measured color. This way, when a computer tries to communicate colors with another device, it does not merely send numbers, but also specifies how those numbers are intended to appear. Color-managed software can then take this profile into account and adjust the numbers sent to the device accordingly. The table below is an example similar to the personalized spiciness table you and your friend created, which compares the input number with an output color.

Input Number (Green)		Output Color	
		Device 1	Device 2
200	—>		
150	—>		
100	—>		
50	—>		

Real-world color profiles include all three colors, more data, and are often more sophisticated than in the above table. In order for these profiles to be useful, they have to be presented in a standardized way which can be read by all programs.

COLOR MANAGEMENT OVERVIEW

The International Color Consortium (ICC) was established in 1993 to create an open, standardized color management system which is now used in most computers. This system involves three key concepts: color profiles, color spaces, and translation between color spaces. A color space relates numbers to actual colors and contains all realizable color combinations. When trying to reproduce color on another device, color spaces can show whether you will be able to retain shadow/highlight detail, color saturation, and by how much either will be compromised. The following diagram shows these concepts for conversion between two typical devices: a monitor and printer.

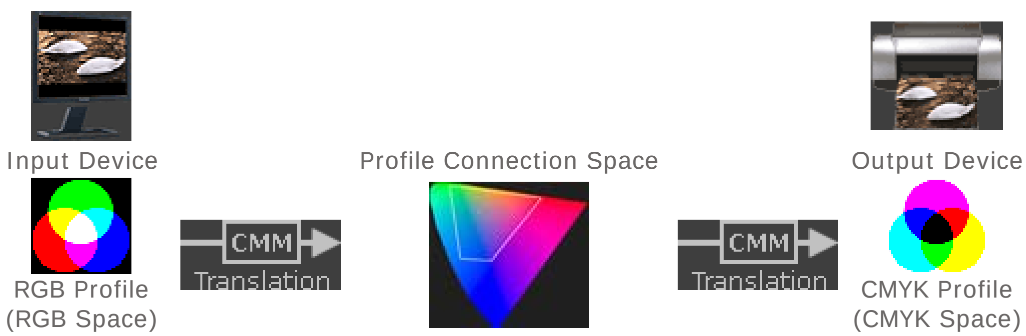

The color profile keeps track of what colors are produced for a particular device's RGB or CMYK numbers, and maps these colors as a subset of the "profile connection space" (PCS). The PCS is a color space which is independent of any device's particular color reproduction methods, and so it serves as a universal translator. The PCS is usually the set of all visible colors defined by the Commission International de l'éclairage (CIE) and used by the ICC. The thin trapezoidal region drawn within the PCS is what is called a "working space." The working space is used in image editing programs (such as Adobe Photoshop) and defines the set of colors available to work with when performing any image editing.

Each step in the above chain specifies the available colors, and thereby defines a color space. If one device has a larger gamut of colors than another device can produce, some of that device's colors will be outside the other's color space. These "out-of-gamut colors" occur with nearly every conversion and are called a "gamut mismatch." A color management module (CMM) performs all calculations needed to translate from one space into another, and is the workhorse of color management. A gamut mismatch requires the CMM to make key approximations that are specified by a "rendering intent." The rendering intent is often specified manually and includes several options for how to deal with out-of-gamut colors.

This all may seem a bit confusing at first, so for a more in-depth explanation of color spaces, profiles, and rendering intent, please visit:

- COLOR MANAGEMENT: COLOR SPACES -

A color space relates numbers to actual colors, and is a three-dimensional object which contains all realizable color combinations. When trying to reproduce color on another device, color spaces can show whether you will be able to retain shadow/highlight detail, color saturation, and by how much either will be compromised.

TYPES

Color spaces can be either dependent to or independent of a given device. Device-dependent spaces express color relative to some other color space, while device-independent color spaces express color in absolute terms. Device-dependent color spaces can tell you valuable information by describing the subset of colors which can be shown with a monitor or printer, or can be captured with a camera or scanner. Devices with a large color space, or "wide gamut," can realize more extreme colors, whereas the opposite is true for a device with a narrow gamut color space.

VISUALIZING COLOR SPACES

Each dimension in "color space" represents some aspect of color, such as lightness, saturation or hue, depending on the type of space. The two diagrams below show the outer surface of a sample color space from two different viewing angles; its surface includes the most extreme colors of the space. The vertical dimension represents luminosity, whereas the two horizontal dimensions represent the red-green and yellow-blue shift. These dimensions could also be described using other color properties.

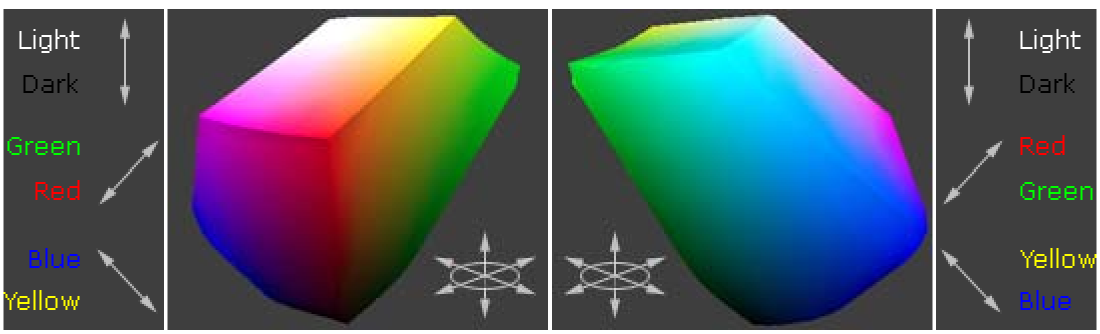

Sample Color Space (Same Space Rotated 180°)

The above color space is intended to help you qualitatively understand and visualize a color space, however it would not be very useful for real-world color management. This is because a color space almost always needs to be compared to another space. In order to visualize this, color spaces are often represented by two-dimensional regions. These are more useful for everyday purposes since they allow you to quickly see the entire boundary of a given cross-section. Unless specified otherwise, two-dimensional diagrams usually show the cross-section containing all colors which are at 50% luminance (a horizontal slice at the vertical midpoint for the color space shown above). The following diagram shows three example color spaces: sRGB, Wide Gamut RGB, and a device-independent reference space. sRGB and Wide Gamut RGB are two working spaces sometimes used for image editing.

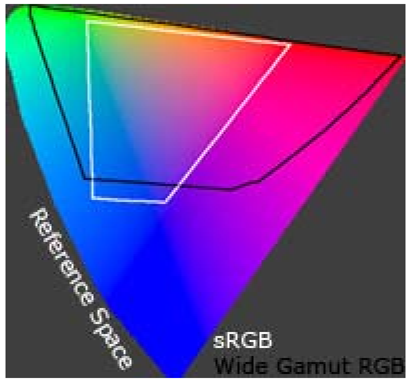

2D Color Space Comparison
(Colors at 50% Luminance)

What can we infer from a 2D color space comparison? Both the black and white outlines show the subset of colors which are reproducible by each color space, as a fraction of some device-independent reference space. Colors shown in the reference color space are only for qualitative visualization, as these depend on how your display device renders color. In addition, the reference space almost always contains more colors than can be shown on a computer display.

For this particular diagram, we see that the "Wide Gamut RGB" color space contains more extreme reds, purples, and greens, whereas the "sRGB" color space contains slightly more blues. Keep in mind that this analysis only applies for colors at 50% luminance, which is what occupies the midtones of an image histogram. If we were interested in the color gamut for the shadows or highlights, we could look at a similar 2D cross-section of the color space at roughly 25% and 75% luminance, respectively.

REFERENCE SPACES

What is the device-independent reference space shown above? Nearly all color management software today uses a device-independent space defined by the Commission International de l' éclairage (CIE) in 1931. This space aims to describe all colors visible to the human eye based upon the average response from a set of people with no vision problems (termed a "standard colorimetric observer"). Nearly all devices are subsets of the visible colors specified by the CIE (including your display device), and so any representation of this space on a monitor should be taken as qualitative and highly inaccurate.

The CIE space of visible color is expressed in several common forms: CIE xyz (1931), CIE L*a*b*, and CIE L u'v' (1976). Each contains the same colors, however they differ in how they distribute color onto a two-dimensional space:

CIE xy CIE a*b* CIE u'v'
(All color spaces shown are 2D cross-sections at 50% Luminance)

CIE xyz is based on a direct graph of the original X, Y and Z tristimulus functions created in 1931. The problem with this representation is that it allocates too much area to the greens. CIE L u'v' was created to correct for this distortion by distributing colors roughly proportional to their perceived color difference. Finally, CIE L*a*b* transforms the CIE colors so that they extend equally on two axes-- conveniently filling a square. Furthermore, each axis in L*a*b* color space represents an easily recognizable property of color, such as the red-green and blue-yellow shifts used in the 3D visualization above.

WORKING SPACES

A working space is used in image editing programs (such as Adobe Photoshop), and defines the set of colors available to work with when performing any image editing. Two of the most commonly used working spaces in digital photography are Adobe RGB 1998 and sRGB IEC61966-2.1. For an in-depth comparison for each of these color spaces, please see sRGB vs. Adobe RGB 1998.

Why not use a working space with the widest gamut possible? It is generally best to use a color space which contains all colors which your final output device can render (usually the printer), but no more. Using a color space with an excessively wide gamut can increase the susceptibility of your image to posterization. This is because the bit depth is stretched over a greater area of colors, and so fewer bits are available to encode a given color gradation.

For further reading, please visit:
Color Management, Part 1
Color Management: Color Space Conversion (Part 3)

-Back to Tutorials-

- COLOR MANAGEMENT: COLOR SPACE CONVERSION -

Color space conversion is what happens when the color management module (CMM) translates color from one device's space to another. Conversion may require approximations in order to preserve the image's most important color qualities. Knowing how these approximations work can help you control how the photo may change-- hopefully maintaining the intended look or mood.

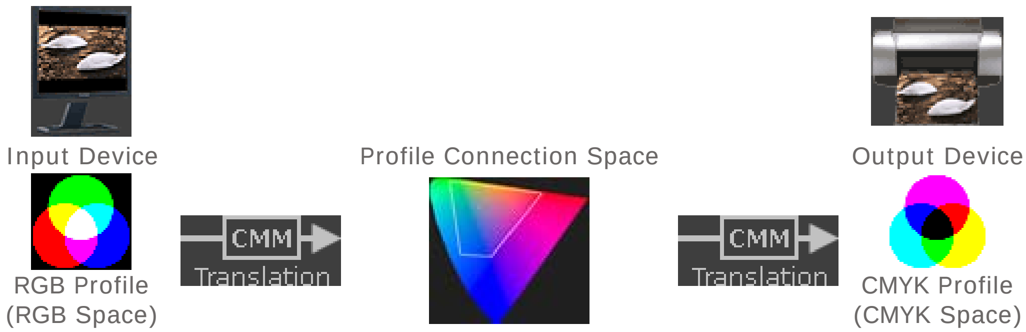

BACKGROUND: GAMUT MISMATCH & RENDERING INTENT

The translation stage attempts to create a best match between devices-- even when seemingly incompatible. If the original device has a larger color gamut than the final device, some of the those colors will be outside the final device's color space. These "out-of-gamut colors" occur with nearly every conversion and are called a *gamut mismatch*.

RGB Color Space

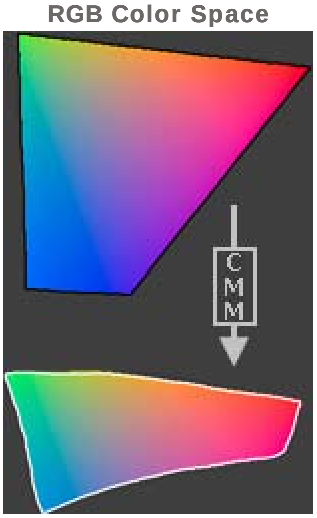

CMYK Color Space
(Destination Space)

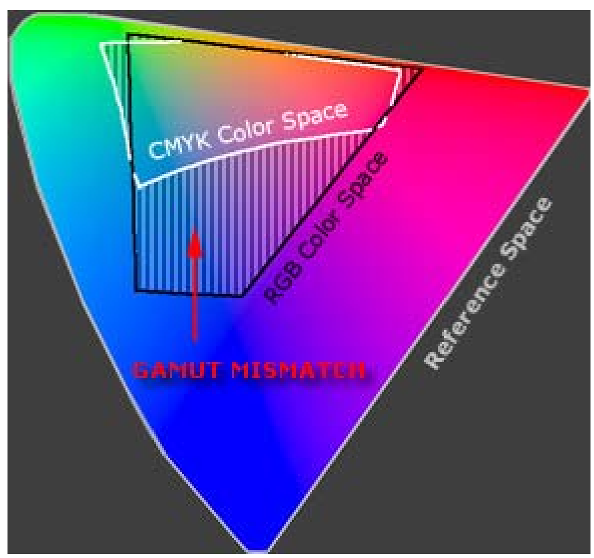

Each time a gamut mismatch occurs, the CMM uses the *rendering intent* to decide what qualities of the image it should prioritize. Common rendering intents include: absolute and relative colorimetric, perceptual, and saturation. Each of these types maintains one property of color at the expense of others (described below).

PERCEPTUAL & RELATIVE COLORIMETRIC INTENT

Perceptual and relative colorimetric rendering are probably the most useful conversion types for digital photography. Each places a different priority on how they render colors within the gamut mismatch region. Relative colorimetric maintains a near exact relationship between in gamut colors, even if this clips out of gamut colors. In contrast, perceptual rendering tries to also preserve some relationship between out of gamut colors, even if this results in inaccuracies for in gamut colors. The following example demonstrates an extreme case for an image within a 1-D black-magenta color space:

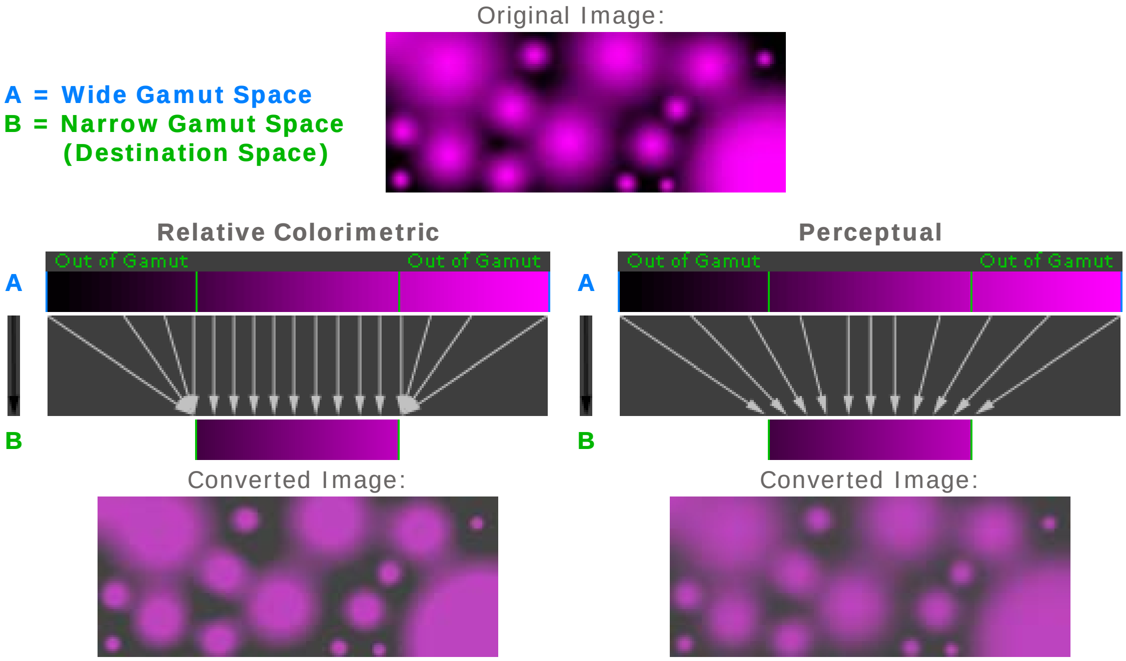

Note how perceptual maintains smooth color gradations throughout by compressing the entire tonal range, whereas relative colorimetric clips out of gamut colors (at center of magenta globules and in the darkness between them). For 2D and 3D color spaces, relative colorimetric maps these to the closest reproducible hue in the destination space.

Even though perceptual rendering compresses the entire gamut, note how it remaps the central tones more precisely than those at the edges of the gamut. The exact conversion depends on what CMM is used for the conversion; Adobe ACE, Microsoft ICM and Apple ColorSynch are some of the most common.

Another distinction is that perceptual does not destroy any color information-- it just redistributes it. Relative colorimetric, on the other hand, does destroy color information. This means that **conversion using relative colorimetric intent is irreversible, while perceptual can be reversed**. This is not to say that converting from space A to B and then back to A again using perceptual will reproduce the original; this would require careful use of tone curves to reverse the color compression caused by the conversion.

ABSOLUTE COLORIMETRIC INTENT

Absolute is similar to relative colorimetric in that it preserves in gamut colors and clips those out of gamut, but they differ in how each handles the white point. The white point is the location of the purest and lightest white in a color space (also see discussion of color temperature). If one were to draw a line between the white and black points, this would pass through the most neutral colors.

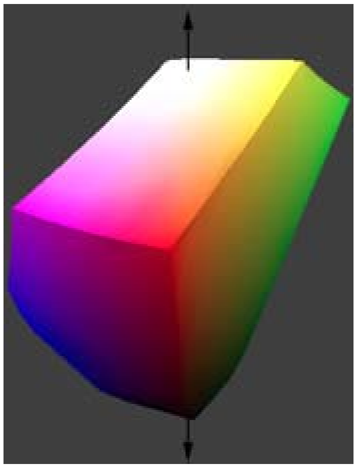

3D Color Space

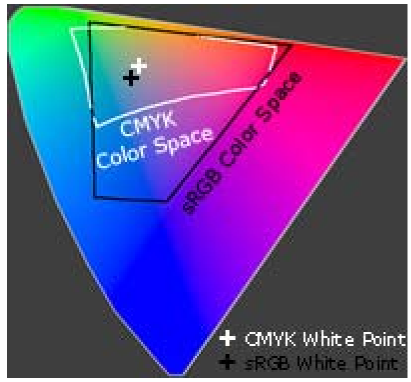

2D Cross-Section
(Two Spaces at 50% Luminance)

The location of this line often changes between color spaces, as shown by the "+" on the top right. Relative colorimetric skews the colors within gamut so that the white point of one space aligns with that of the other, while absolute colorimetric preserves colors exactly (without regard to changing white point). To illustrate this, the example below shows two theoretical spaces that have identical gamuts, but different white points:

X = White Point

Color Space #1 | Color Space #2 | Absolute Colorimetric | Relative Colorimetric

Absolute colorimetric preserves the white point, while relative colorimetric actually displaces the colors so that the old white point aligns with the new one (while still retaining the colors' relative positions). The exact preservation of colors may sound appealing, however relative colorimetric adjusts the white point for a reason. Without this adjustment, **absolute colorimetric results in unsightly image color shifts, and is thus rarely of interest to photographers**.

This color shift results because the white point of the color space usually needs to align with that of the light source or paper tint used. If one were printing to a color space for paper with a bluish tint, absolute colorimetric would ignore this tint change. Relative colorimetric would compensate colors to account for the fact that the whitest and lightest point has a tint of blue.

SATURATION INTENT

Saturation rendering intent tries to preserve saturated colors, and is most useful when trying to retain color purity in computer graphics when converting into a larger color space. If the original RGB device contained pure (fully saturated) colors, then saturation intent ensures that those colors will remain saturated in the new color space-- even if this causes the colors to become relatively more extreme.

Pie chart with fully saturated cyan, blue, magenta and red:

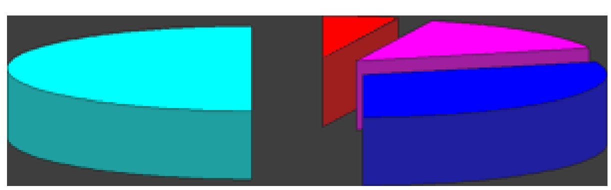

Saturation intent is not desirable for photos because it does not attempt to

maintain color realism. Maintaining color saturation may come at the expense of changes in hue and lightness, which is usually an unacceptable trade-off for photo reproduction. On the other hand, this is often acceptable for computer graphics such as pie charts.

Another use for saturation intent is to avoid visible dithering when printing computer graphics on inkjet printers. Some dithering may be unavoidable as inkjet printers never have an ink to match every color, however saturation intent can minimize those cases where dithering is sparse because the color is very close to being pure.

Visible dithering due to lack of fully saturated colors:

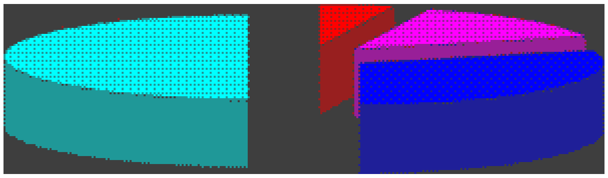

PAY ATTENTION TO IMAGE CONTENT

One must take the range of image colors present into account; just because an image is defined by a large color space does not mean that it actually utilizes all of those extreme colors. If the destination color space fully encompasses the image's colors (despite being smaller than the original space), then relative colorimetric will yield a more accurate result.

Example Image:

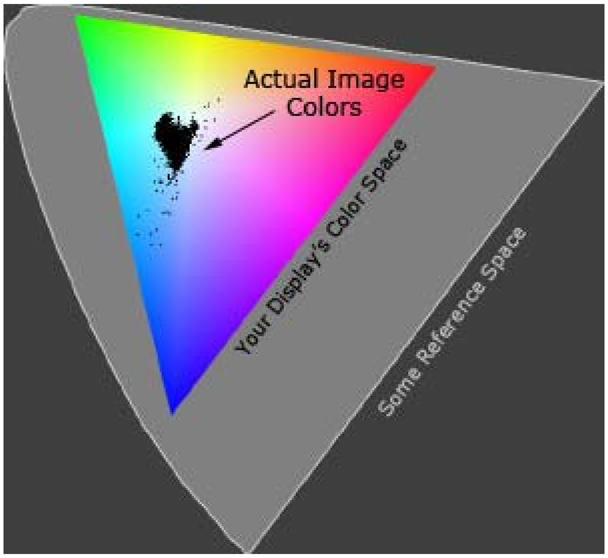

The above image barely utilizes the gamut of your computer display device, which is actually typical of many photographic images. If one were to convert the above image into a destination space which had less saturated reds and greens, this would not place any image colors outside the destination space. For such cases, relative colorimetric would yield more accurate results. This is because perceptual intent compresses the entire color gamut--regardless of whether these colors are actually utilized.

SHADOW & HIGHLIGHT DETAIL IN 3D COLOR SPACES

Real-world photographs utilize three-dimensional color spaces, even though up until now we have been primarily analyzing spaces in one and two dimensions. The most important consequence of rendering intent on 3D color spaces is how it affects shadow and highlight detail.

If the destination space can no longer reproduce subtle dark tones and highlights, this detail may be clipped when using relative/absolute colorimetric intent. Perceptual intent compresses these dark and light tones to fit within the new space, however it does this at the cost of reducing overall contrast (relative to what would have been produced with colorimetric intent).

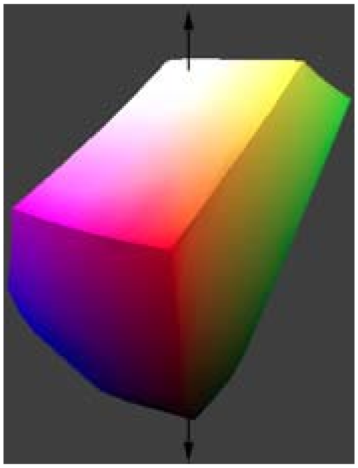

The conversion difference between perceptual and relative colorimetric is similar to what was demonstrated earlier with the magenta image. The main difference is that now the compression or clipping occurs in the vertical dimension-- for shadows and highlight colors. Most prints cannot produce the range of light to dark that we may see on our computer display, so this aspect is of particular importance when making a print of a digital photograph.

Using the "black point compensation" setting can help avoid shadow clipping-- even with absolute and relative colorimetric intents. This is available in the conversion properties of nearly all software which supports color management (such as Adobe Photoshop).

RECOMMENDATIONS

So which is the best rendering intent for digital photography? In general, perceptual and relative colorimetric are best suited for photography because they aim to preserve the same visual appearance as the original.

The decision about when to use each of these depends on image content and the intended purpose. Images with intense colors (such as bright sunsets or well-lit floral arrangements) will preserve more of their color gradation in extreme colors using perceptual intent. On the other hand, this may come at the expense of compressing or dulling more moderate colors. Images with more subtle tones (such as some portraits) often stand to benefit more from the increased accuracy of relative colorimetric (assuming no colors are placed within the gamut mismatch region). Perceptual intent is overall the safest bet for general and batch use, unless you know specifics about each image.

For related reading, please visit:
Color Management, Part 1
Color Management: Color Spaces (Part 2)

-Back to Tutorials-