Understanding Rendering Intents

By Mike Chaney

Used with Permission

Background

Let's take a break from print driver settings this month and talk about rendering intents. Rendering intents are an often misunderstood concept of color management that can affect how you use your ICC profiles. As usual, I'll try to keep things as simple as possible in order to assist in the understanding of how rendering intents affect results when you select perceptual, relative colorimetric, saturation, and absolute colorimetric intents when using ICC profiles. That said, it is impossible to discuss this subject accurately without getting into some technical detail so if you feel like you are getting lost, stick with it and we'll summarize it in the end to hopefully put some meaning into the tech talk. If you've read information on other sites regarding rendering intents and you think you know all there is to know about them, please read on as there is an abundance of inaccurate information on the web regarding rendering intents and how they actually work.

Why are rendering intents needed in the first place?

Color management (the use of ICC profiles) is a game of give and take and involves many compromises. The main reason that we make compromises in color management is because the gamut (range of colors available) on one device is often very different from another. A bright saturated blue color displayed on your monitor for example may not be reproducible on your printer because your printer simply cannot make that color. In this case, we would say that the example blue color is in the monitor's color gamut but outside the printer's color gamut. Trying to reproduce a reasonable visual representation for colors that are out of range on a device is what rendering intents are all about. For example, we may use tricks like reducing the saturation of the entire print so that a color that is out of range still appears a bit more vibrant than ones that are in range. Rendering intents simply use different methods to "trick" the eye into believing that the print can reproduce irreproducible colors.

Visual representation of color gamuts

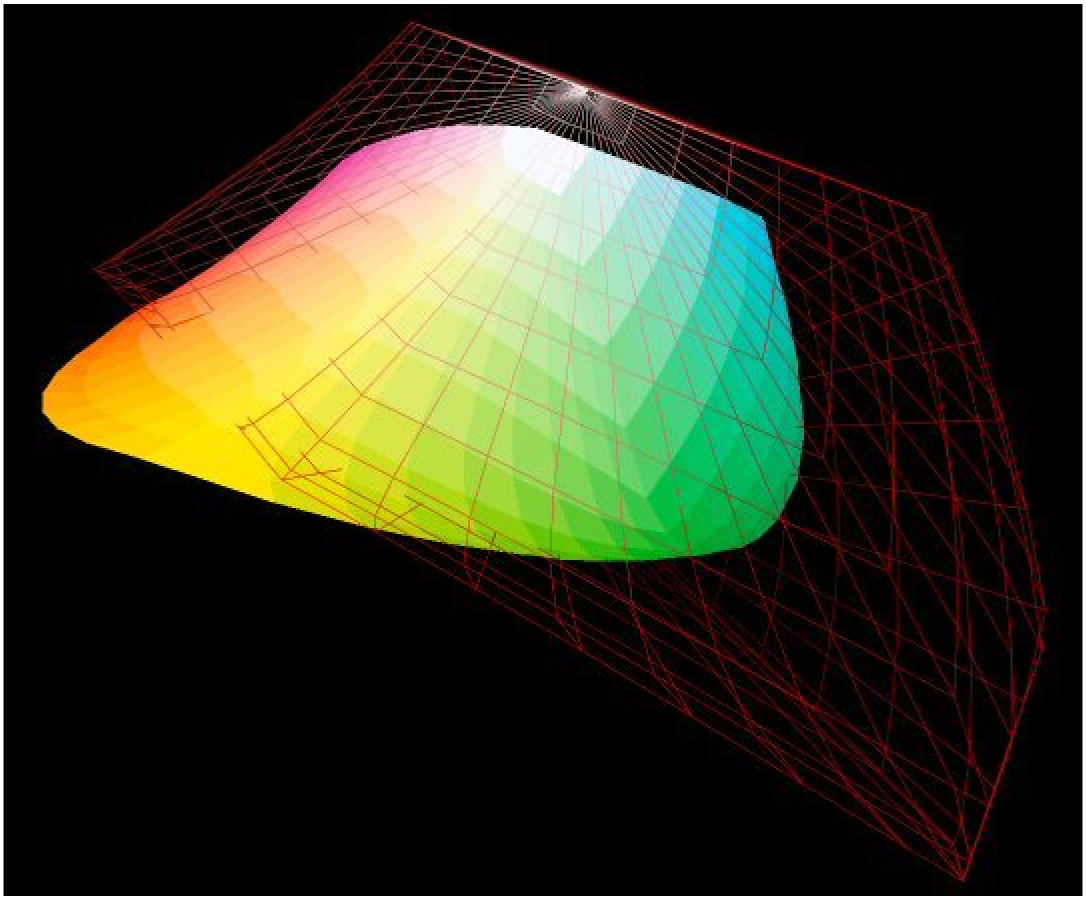

The above represents the gamut of colors available in the Adobe RGB color space (red wire frame) and the gamut of the Canon i9900 printer using Canon Photo Paper Pro (solid color gradient). The hue and saturation of the color (red, green, blue, orange, etc.) is represented in the 2 dimensional X/Y axes while the luminance (brightness) is on the Z axis. The red wire frame shows the range of all possible colors available in your image if your image is stored in Adobe RGB color space. The solid gradient indicates the range of all possible colors reproducible by your printer. As you can see, most of the printer's color gamut is contained within the Adobe RGB wire frame but small sections like the bright yellow peak that protrudes through the Adobe RGB gamut (middle-left) indicate that there are some colors that the printer can reproduce that cannot be captured in the image. There are also large sections of the Adobe RGB color space like the empty space in the wire frame on the bottom-right that indicate colors that might be present in your original image that cannot be reproduced by your printer. It is the handling of these "out of range" colors that we'll refer to in discussing the rendering intents below.

Relative Colorimetric Intent

If we look at the 3D gamut representation above, we can see that our Adobe RGB images might have colors (lower right of the red wire frame) that cannot be reproduced by our printer and that same image might also have colors that are inside both the solid gradient and the wire frame, indicating that those colors are in our image and they are reproducible on the printer. One way to handle the mismatch of gamuts is to: (a) render all colors that are present in the image and ***are*** directly reproducible by the printer to the proper color and (b) render all colors that are ***outside*** the printer gamut to the nearest color on the edge of its gamut (called the "gamut hull"). The gamut hull, visible in the above graphic, simply represents the extremes of what the color space or device can reproduce and usually represents colors that are bright and saturated. Other colors inside the hull (not on the surface) are simply colors that are less vibrant so that they are not as close to the extremes of what the color space or device can reproduce.

Looking at the gamuts above, all colors that are inside both the solid gradient and also the red wire frame are not a problem: we can render those directly from the image color to the same color on the print. Colors that lie inside the wire frame but are outside the solid gradient are colors that are in your image but are not reproducible on the printer. For those, we look at the gamut hull (solid gradient above) and we pick the point on its surface that is closest in distance to the color we are trying to make in the wire frame. This method obviously has some drawbacks. Since there are many colors inside the image gamut that can map to the same point on the printer gamut hull, it means that we may see banding in our prints as a particular gradient (like blue sky for example) might all map to the same spot on the hull of the printer's gamut. On the positive side, at least all of the colors that are reproducible by the printer are reproduced accurately.

Note that the small section of yellow that protrudes through the wire frame indicates colors that the printer ***can*** reproduce but cannot be contained within our original image. Since the image cannot contain those colors, we need not worry about trying to reproduce them in the context of rendering intents because they cannot be present in the image.

Perceptual Intent

As mentioned above, relative colorimetric intent has the benefit of being able to reproduce all colors that are reproducible. That is, if the printer can reproduce the color, it will... accurately. The down side to relative colorimetric intent is that we often have images that exceed the printer's gamut and when this happens, we may see banding (sometimes called posterization) of color in the prints.

Perceptual intent is a rendering method that tries to get around the fact that out-of-range colors might result in banding. With perceptual intent, we compress or "squash" the gamut of the image down a bit so that not as much of the wire frame sticks out beyond the solid gradient of our printer's gamut. If we squash the gamut of the image so that the wire frame is smaller, there won't be as many colors in the image that can't be reproduced on the printer. This will eliminate or at least reduce the amount of banding in the prints because more colors (in our image) are now in range of the printer.

When we artificially squash the image gamut down to try to fit more of it inside the printer's smaller gamut, we generally end up with a print that has reduced saturation. If we reduce saturation by only a little, our eyes may not notice the difference on the print other than the colors not being quite as vibrant as those under the non-squashed relative colorimetric intent. All things being relative, the print can look better because although it is a bit less vibrant in color, the banding present in the relative colorimetric rendering is gone or at least reduced substantially.

Other than reduced color vibrancy, there is another down side to perceptual intent. The current ICC CMM (color management model) is not a "smart" model meaning that it cannot and does not examine the gamut of the actual ***image*** before trying to compress it to fit in the printer's gamut. While the gamut available to your original image (the red wire frame) is large, your actual image may only contain a few dull colors like some green foliage and people wearing light pastel clothing. All colors in your image in this case may be reproducible by the printer. A "smart CMM" might be able to look at the original image and determine that it doesn't ***need*** any squashing to be printed. Unfortunately, the CMM does not have any knowledge about the image being rendered and must perform a sort of "blind rendering" that assumes that all possible colors must be taken into account whether or not they actually exist in your image!

Some misinformation on the web would lead you to believe that because the CMM cannot account for ***image*** gamut, that it simply compresses the entire gamut of the color space used by the image so that it fits inside the printer's gamut. This is, however, also not true. Squashing the entire color space where the image resides into the printer's color space would amount to taking the entire wire frame above and shrinking it in size so that no corners protrude through the printer's solid gamut in the diagram. As you may be able to see by the graphic, that would be an extreme amount of compression that would result in noticeable color desaturation. In addition, it would mean that the same image encoded in two different color spaces of different sizes (say sRGB and Adobe RGB) would result in two different prints with different amounts of desaturation even though the original images should appear the same as they both have all the same colors: they are just encoded differently.

What all this amounts to is the fact that perceptual intent basically uses an arbitrary amount of gamut compression (squashing) in order to reduce the banding effects that might be present in relative colorimetric intent. The amount of compression, which will show up in the printer's ICC profile, is really up to the creator of the printer profile. What is normally done is that when creating a printer profile, the available gamut of the printer is taken into account and a level of compression is chosen so that most colors that are likely to be seen in a photograph will be "pulled back" into the printer's gamut. If it sounds "wishy washy", that's because it is! Many web sites point out the original concept of rendering intents and point out that relative colorimetric intent clips the gamut while perceptual compresses it, but these "ideal" concepts are not what is ultimately going on behind the scenes in the CMM.

What does all this mean?

By now you are probably either kicking yourself for even reading this article because it seemed so simple before, or you're getting close to deciding to just use perceptual intent and not worry about the whole subject of rendering intents. :-) My purpose, however, is not to confuse but to inform. I want people to understand the limitations of color management as it exists today and to understand what is really going on not just the high level concepts. If I were to try to put all of the above as simply as possible, I'd say:

Perceptual Intent: Use this method for most of your work especially if you intend to just set it and leave it alone. Perceptual intent will produce prints with accurate hue and while overall saturation levels may be a bit less, you are unlikely no notice this by just examining the print by itself. In addition, this method reduces artifacts like banding in blue skies.

Relative Colorimetric Intent: Use this rendering method in certain cases where reproducing accurate colors is paramount. This rendering intent is often used when your original image contains only a narrow range of colors. As an example, if you are reproducing an image of the Grand Canyon and there are only rust colored Earth tones in the scene, perceptual intent may take some of the clarity out of the photo because of the compressed color gamut and because there aren't a variety of other colors present for our brains to get a relative "lock" on the entire scene. Using relative colorimetric intent in this case should make the texture of the rock look more realistic and more defined because all of the colors in the photo are likely to be within the gamut of the printer due to the fact that they are not bright, saturated colors.

What about the other intents?

We have two intents left, but I won't spend much time on those since they are of little/no value when reproducing ***photographs***.

Absolute Colorimetric Intent: Absolute colorimetric tries to reproduce the exact colors recorded in the original scene. Sounds even better than relative colorimetric until you realize that absolute colorimetric intent reproduces these colors with no regard (no adaption) for the illuminant or light source. Simply put, using absolute colorimetric intent will usually result in awful color shifts because our eyes will try to adapt to the illuminant (white of the paper, color temperature of the monitor, etc.) and the same color may look different under different lighting. As such, absolute colorimetric is used mainly for reproduction of specific colors like reproductions of fabric or logo colors.

Saturation Intent: With perceptual rendering intent, we may sometimes notice that colors have been a bit desaturated to fit bigger gamuts into smaller ones. To overcome this sense of desaturation, the saturation rendering intent tries to keep accurate saturation while shifting other factors like the hue of the color. This intent can be useful for things like screen captures, bar graphs, and other images where the hue of the color is less important than the overall "pop" of the image. Simply put, in photographs, people are more likely to notice that a stop sign looks too magenta than the stop sign not being vibrant enough. The converse is true when displaying a pie chart where people are much less likely to care that the red in the pie chart looks a little shifted toward magenta but the presentation may have less "impact" if the entire pie chart looks dull!

Summary

These are the games we play when trying to fool our eyes into believing that a print or an image on the monitor looks the same as it did when it was recorded/captured and what is really going on behind the scenes when we make the decision about which rendering intent to use. Hopefully the information in this article has given you a bit more solid a foundation to stand on when dropping down that often misunderstood selection called "rendering intent".