The sYCC Color Space

Douglas A. Kerr, P.E.

Issue 2
July 24, 2015

ABSTRACT

The sYCC color space is an alternative representation of the "sRGB" color space, with a special wrinkle though which it can represent a larger color gamut than the sRGB space proper. In this article we review the definition, principles, and implications of the sYCC color space. We also discuss the matter of characterizing the "chrominance gamut" of a color space.

INTRODUCTION

A widely used color space in which the colors of the pixels in an image are represented in the field of digital photography is the sRGB ("standard RGB") color space. It is the default color space for many modern digital cameras, and is the default color space for images to be presented on Web pages. It is a precisely-specified color space within the broader "RGB" color space family.

The most widely-used file format for digital camera images, the Exif format (JPEG form), prescribes that the standard basic color space in which the images are represented be the sRGB space. But in fact the image is actually transformed to a form of the YCbCr space before encoding under the JPEG compression system.

A particular form of the YCbCr space, called "sYCC", used for this purpose, actually has a larger color gamut than the sRGB space from which it is derived—that is, it can represent colors that cannot be represented in the sRGB space itself (in its normal form). Thus we have the potential of conveying, in an Exif JPEG file, this larger gamut of colors. This could of course be beneficial in cases where the image is to be rendered by a device capable of a larger color gamut than that which can be represented in the sRGB color space itself (true of many photo printers today).

THE COLOR SPACES

The sRGB color space

The sRGB color space, like all RGB-family color spaces, describes a color in terms of a "recipe" for composing the color with designated amounts of light from three "primary" sources, identified by the broad names "red", "green", and "blue (identified as R, G, and B), and having certain specified chromaticities. We will use the symbols **r**, g, and **b** to represent the quantities of these three primaries that make up the color of interest.

However, it is not the quantities **r**, **g**, and **b** that are "sent" to indicate the color of a particular image pixel but rather versions of the same quantities on a non-linear scale. These are identified by the symbols R, G, and B. This nonlinear transformation is often called “gamma precompensation”, since its original motive was to allow the values R, G, and B to directly control the three “guns” of a tricolor cathode ray tube (CRT) display device in a television receiver. In such a device, the luminance generated by a gun is not proportional to the voltage signal applied to the control electrode of the gun, but typically to some power of that voltage (perhaps 2.2). That exponent is called “gamma” by parallel with a similar parameter that describes the response of black ad white negative film to exposure.

The use of “gamma precompensation” in the “sent” representations of the amounts of the three primaries avoids the need, in the display chain, to “linearize” the response of the guns. Although this was the initial motivation for the adoption of gamma precompensation, there are other advantages, which are beyond the scope of this article. The principle of gamma precompensation has been carried forward from television practice to application to RGB color spaces used in digital photography.

Different color spaces in the RGB family use different gamma precompensation functions. In the sRGB color space, it works this way. Here, c represents any of the “linear” variables, **r**, **g**, or **b** (on a scale of 0-1), while C represents the corresponding gamma-precompensated variable, R, G or B (also on a scale of 0-1).

$$C = 12.92c \qquad \text{for } 0 < c \leq 0.0031308$$

$$C = 1.055c^{1/2.4} - 0.055 \qquad \text{for } 0.0031308 < c$$

The YCbCr color space

In analog television broadcasting, although the color representation is initially developed in a certain defined RGB color space, the format in which that information is transmitted uses a transformation of that space called the YUV color space. Similarly, when a still image in RGB form is being prepared for encoding under the JPEG image encoding and compression system, it is transformed to another representation very similar to the YUV space, known as the YCbCr space. Let’s look at its definition.

We begin with the gamma-precompensated R, G, and B values, having a range of 0-1, and derive from them the value Y, thus:

$$Y = 0.299R + 0.587G + 0.114B$$

Note that despite its symbol, Y does not represent the luminance of the color represented by R, G, and B, although it is something like it. It is sometimes called *luma* (a term drawn from television transmission technology). Y can have any value from 0-1.

We then develop two "color difference" values, Cb and Cr, thus:

$$Cb = 0.564(B - Y)$$

$$Cr = 0.713(R - Y)$$

The coefficients in these expressions ensure that (if R, G, and B are within the range of 0-1) Cb and Cr will lie in the range -0.5 to +0.5.

The values of Y, Cb, and Cr, converted to 8-bit digital form, are then used to describe the color of the individual pixel. (Cb and Cr are not represented as "signed integers" but rather in an "offset" form, in which the digital value 128 represents a value of zero.)

Note that in the traditional use of the YCbCr representation of an RGB color, the only YCbCr values that are valid are those that come from valid RGB combinations.

THE sYCC COLOR SPACE

The international standard for the sRGB color space (IEC 61966-2-1) provides, in Annex G, a definition of a special YCbCr color space for use in the sRGB environment. It is designated the "sYCC" color space (meaning it is specifically the YCbCr transformation of the sRGB space). It basically follows the familiar definition of a YCbCr color space, as we discussed above.

But in fact the sYCC color space as defined carries a further, and important, property. All combinations of "legitimate" values of Y, Cb, and Cr that imply visible colors (whether or not those colors exist within the sRGB color space) are deemed to represent valid colors under sYCC. Thus, by this simple stroke of the specification-writer's pen, the sYCC color space is invested with a larger gamut than that of the sRGB space from which it is derived.

Thus, we have in effect defined a special variant of the sRGB space (not having a designation![1]) whose gamut is the same as that of sYCC. To permit this, this special sRGB space must allow a greater range in the values R, G, and B than for the "normal" sRGB space.

For "visible" colors lying in the sYCC gamut, but not in the "normal" sRGB gamut, there still can be a representation in terms of the variables **r**, **g**, and **b** [2] (based on the traditional primaries, R, G, and B), but at least one of the three variables must have a negative value or a value greater than 1 (not allowed by the sRGB specification itself). In this special "parent of sYCC" form of sRGB, such excursions

[1] It is not, for example, the "extended sRGB" (e-sRGB) color space, although it has some things in common with it.

[2] These are the non-gamma-precompensated ("linear light") amounts of the three primaries, R, G, and B. The values represented by R, G, and B in the normal representation of the RGB space are gamma-precompensated.

are allowed. While this cannot have physical meaning (in terms of an "amount" of a certain primary light), it is perfectly acceptable mathematically, and after all, we are not actually "generating" any light at this point—just describing its color.

This of course also implies excursions of the underlying "linear" variables, **r**, **g**, and **b**, into negative or greater-than-1 territory.

To make this possible, the gamma precompensation function needs to be redefined so as to work with negative input (and output) values. This expanded definition is:

$$C = -1.055(-c)^{1/2.4} + 0.055 \qquad \text{for } c < -0.0031308$$

$$C = 12.92c \qquad \text{for } -0.0031308 \le c \le 0.0031308$$

$$C = 1.055c^{1/2.4} - 0.055 \qquad \text{for } 0.0031308 < c$$

Where c represents any of **r**, **g**, or **b** (the non-gamma-precompensated values) and C represents any of R, g, or B (the gamma-precompensated values).

This un-named expanded version of sRGB does not have a standard way of representing it in binary form—that is only prescribed for its sYCC transformation. It is a "virtual" color space, an ethereal step on the way to representing an image in sYCC form. In general, representing a color under this color space in binary form (as will happen in the processing of the image information) requires at least 9 bits per "channel" if we are to retain the customary precision in the values R, G, and B.

Note however that the sYCC representation uses just 8 bits per channel.

GAMUTS OF THE TWO SPACES

Figure 1 shows the gamuts of the two color spaces, presented as an oblique projection in the sYCC coordinate system.

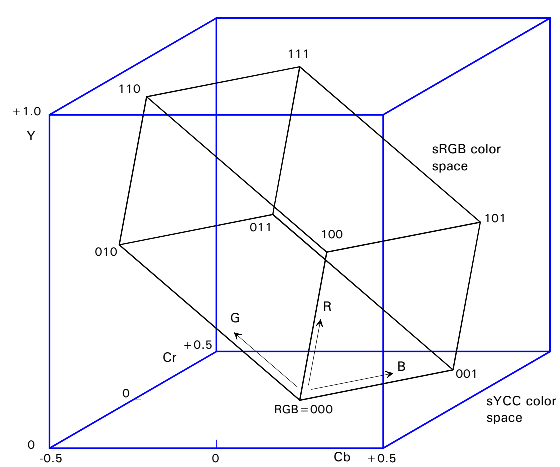

Figure 1. sYCC and sRGB color spaces

You can readily see the much greater gamut of the sYCC space. Note however (and this cannot be seen on this figure) that some of the volume of the YCbCr space that lies outside the RGB space also lies outside the realm of realizable (visible) chromaticities.

The chromaticity gamut

When we speak of the larger gamut or a particular color space as a criterion making it perhaps desirable for use in some photographic image context, we are usually really interested in the chromaticity gamut—the range of “colors” it can represent, as the lay person understands “color”. After all, the luminance property is essentially on an arbitrary scale.

We often see the chromaticity gamut of a color space such as sRGB or Adobe RGB illustrated on the CIE chromaticity diagram (coordinates x-y). But that plot is very misleading.

It turns out that (even for the RGB color space, as we’ll see shortly), the chromaticity gamut depends on where the color falls in the luminance range of the color space. Thus, we must really examine the three-dimensional ”plot” of the gamut to grasp what chromaticity gamut we will have for different luminances.

There are available many interesting “three-dimensional” portrayals of different color spaces available. But almost always, the coordinate system in which these are plotted is the CIE Lab color model. But that’s really not helpful in trying to visualize the chromaticity gamut.

The reason is that the a-b plane of that space does not represent chromaticity—it represents chrominance. Thus, two colors that have the same a-b coordinates, but different Y coordinates (luminance), do not have the same chromaticity—in particular, they have the same hue, but not the same saturation.

Looking at it from the other direction, two colors having the same chromaticity (for example, the same hue and saturation) but different luminance will have different a-b values.

We can see that in this example of two different colors with the same chromaticity (fully-saturated primary red, as a matter of fact) but different luminance (the second, in fact, has half the luminance of the first). We’ll use the CIE xyY model to express luminance and chromaticity. Note that these two colors have different a-b coordinates.

x	y	Y	L	a	b
0.65	0.33	0.20	51.8	78.0	67.4
0.65	0.33	0.10	37.8	61.9	53.4

So if we had plots of an "a-b plane section" of the three-dimensional plot of the gamut of a color space for different values of luminance, we cannot compare them in terms of chrominance gamut.

But the x-y plane of the CIE chromaticity diagram does indicate chromaticity (it certainly should, given its name!). So if we want to compare the chromaticity gamuts of two different color spaces, at different luminance, we would be best to do it on the CIE x-y plane.

Chromaticity gamut of the sRGB color space

I mentioned above that the chromaticity gamut of an RGB model (whether sRGB, Adobe RGB, or some other one) is not the same for the whole range of luminance. How can this be? Isn't the RGB (sRGB) chromaticity gamut bounded by the triangle connecting the locations of the three primaries on the CIE chromaticity diagram, as we see in figure 2? We often speak as if it were.

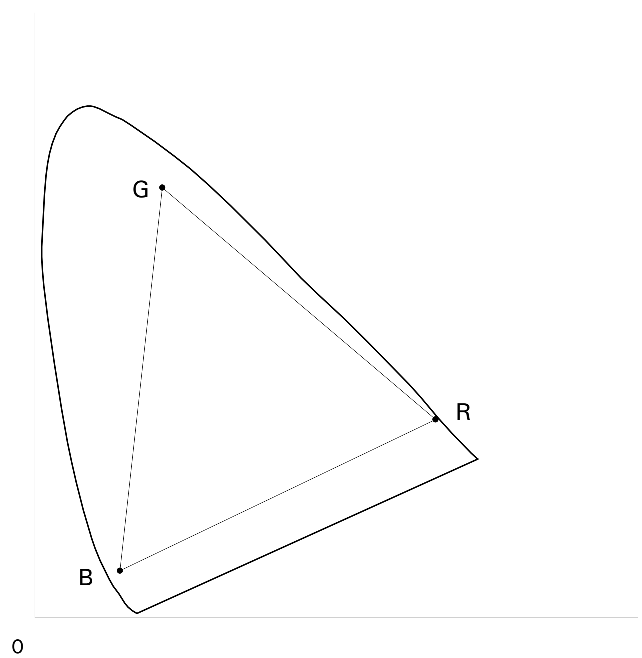

Figure 2. The "sRGB chromaticity gamut"—but not really

But it isn't, and a little thought about an extreme case will show why. Suppose we are interested in the chrominance we can describe as "fully-saturated primary red". To get that, we must use the RGB recipe *n*,0,0, where *n* can be from 1 to 255. The highest luminance fully saturated primary red we can make has the RGB recipe 255,0,0. Its luminance (on the normalized scale of 1-0) is 0.212. Suppose we would like that same chromaticity at a luminance of 0.6? We can't make it!

But, for fully-saturated primary green, we can make a luminance up to 0.715. And for fully-saturated primary blue, we can make a luminance only up to 0.072.[3]

Thus, once we get above a luminance of 0.072, the gamut no longer occupies the entire triangle. And in fact to reach a luminance of 1.000, the only chrominance we can get is the reference white for the color space—the gamut has shrunk to an infinitesimal area.

We can see this in figure 3, which represents a three-dimensional plot of the gamut of the sRGB color space in the CIE xyY coordinate system. The x and y axes are the familiar ones of the CIE chromaticity diagram, and the Y axis is of course relative luminance.

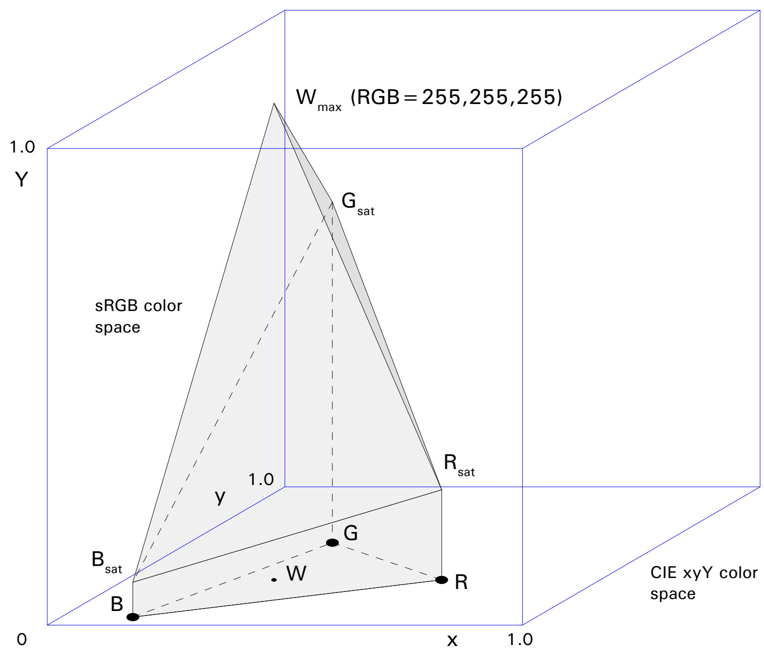

Figure 3. The sRGB gamut seen in CIE xyY space

We have plotted the chromaticities of our three primaries, R, G, and B, on the floor of the xyY space (where Y=0). Actually those points in the three-dimensional xyY chart don't represent any color at all, since the luminance there is zero. The points marked R_{sat}, G_{sat}, and B_{sat} are the highest luminance colors whose chromaticity is fully-saturated primary red, green, and blue, respectively. The point marked W_{max} is the highest luminance color we can make at all, which by definition (since it comes from RGB=255,255,255) is the reference white for this color space in chromaticity. (W shows the chromaticity of the reference white.)

[3] This all comes from the differing sensitivity of the eye to the different primaries.

The classical single portrayal of the sRGB chromaticity gamut on the CIE chromaticity diagram would imply the “triangular prism” we would get if we extruded, along the luminance axis, the triangle joining the points R, G, and B. Here, we see this prism in fact exists up to the level of B_{sat}, but beyond that the cross-section of the figure decreases, its area approaching zero as we approach W_{max}.

Figure 4 shows the often-shown chromaticity gamut (which in reality is only available up to a luminance of Y=0.072), and the actual chromaticity gamut at a luminance of Y=0.715. At intermediate values of luminance, the gamut is of intermediate size (and is in fact not triangular in shape but rather is a polygon).

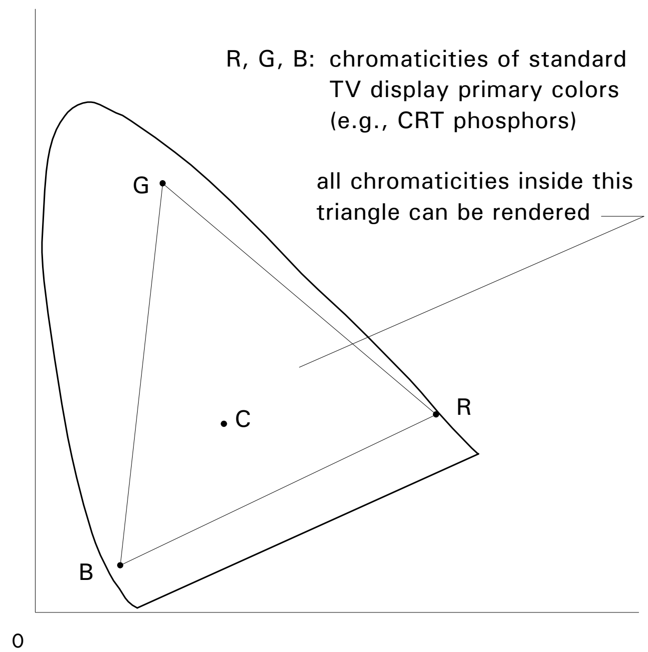

Figure 4. The shrinking chromaticity gamut

So much for the famous triangular chromaticity gamut of RGB color spaces.

Chromaticity gamut of the sYCC color space

Now that we’ve learned how to look at, and think about, the chromaticity gamut of a color space, let’s return to the real subject of this article, the sYCC color space.

Let’s first review what makes this color space different in gamut from its cousin, the sRGB space.

The gamut of the sRGB space is constrained by the limit that the values R, G, and B can only take on values in the range of 0-1 (0-255, in the encoded form). It is this fact, combined with the differing sensitivities of the eye to the three primaries, that gives us the three-dimensional gamut we see on figure 3. And in turn it is responsible for the “shrinking” chromaticity gamut of sRGB we see on figure 4.

The sYCC color space is not bound by that limit—only the limit that Cb and Cr are limited to the range ±0.5 (and of course Y to the range 0-1).

The nature of the transforms involved means that the chrominance gamut of the sYCC space does not shrink much with increasing luminance (although in fact it somewhat shifts its location). This is the major actual advantage of the sYCC space—it gives us a substantially larger gamut then sRGB for the higher luminance values. We can see this on figure 5.

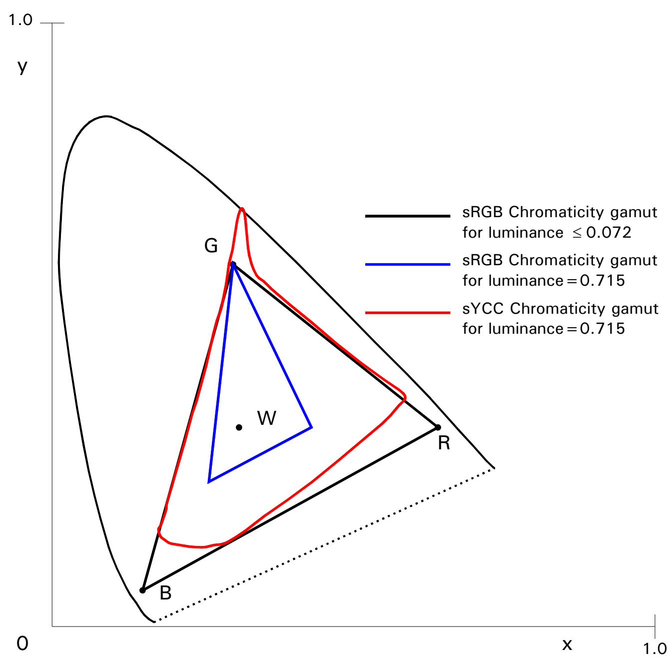

Figure 5. Chrominance gamut of sYCC

Here, we show the chromaticity gamut of the sYCC color space for a luminance of 0.715, the same luminance for which we show the sRGB chromaticity (the "shrunken triangle".

What about the infamous "invisible colors" within the sYCC gamut (those outside the CIE "horseshoe")? These occur only in the lower-luminance portion of the sYCC space.

IMPLICATIONS ON DIGITAL PHOTOGRAPHY

The "JPG" files delivered by most modern digital cameras follow the JPEG option of what is known as the Exif file specification. Under that specification, the image is carried in a JPEG-encoded form of the YCbCr color space—to be precise, the sYCC space.

The file format specification is quite clear that this is to be interpreted as meaning that the representation of colors lying outside the sRGB gamut but within the sYCC gamut is legitimate in an Exif JPEG file.

But it does not seem that, at present, typical applications that decode a JPEG file and deliver an output to computer memory for editing, or to a display or a printer, are prepared to deal with such “out-of-sRGB-gamut” colors that might be found in the JPEG-encoded sYCC data.

Thus, I doubt that any cameras today actually exploit the expanded gamut of sYCC in their JPEG output files. It seems likely that, when delivering images in JPEG files (and assuming the “sRGB” color space has been selected) the output gamut is in fact restricted to that specified for sRGB (just as the selection of the sRGB color space would imply).

However, it certainly appears that both modern Canon and Nikon cameras (perhaps others as well), when operating in a “direct-to-printer” mode, may well exploit the greater potential gamut allowed by the sYCC color space (to the extent that the connected printer is capable of handling that).

Definitive information on this is hard to come by.

Perhaps in the future we may have applications that will accept a JPEG image file in which “out-of-sRGB-gamut” colors appear, and cameras that can be set to use the “sYCC” color space when delivering a JPEG output file.

#