72nd
birthday special

The HSV and HSL Color Models and the Infamous Hexcones

Douglas A. Kerr

Issue 3
May 12, 2008

ABSTRACT

The color models known as HSV and HSL (and specific color spaces based on them) are intended to provide ways of describing color that have a broad relationship to the easily-grasped color attributes hue, saturation, and relative luminance. Almost invariably when an HSV color space is described, it is mentioned that "the color space can be described by a hexcone (a homey synonym for *hexagonal pyramid*)". Similarly, we hear that the HSL color space can be "described by a bi-hexcone" (meaning two hexagonal pyramids joined at their bases).

But these descriptions are paradoxical—the gamuts of these color spaces, "plotted" in their inherent cylindrical coordinate systems, are in fact full circular cylinders. So what might be meant by the allusion to these other, tapered solid figures?

In this article, we describe these two color models and reveal the rationales by which these unusual geometric figures are invoked as representing the associated color spaces.

The article begins with a review of the principles of color, color models, color spaces, and gamuts.

An appendix describes a "different" HSL color model used by Canon in some of its image manipulation software.

SUMMARY

The HSV (sometimes called HSB) and HSL color models (they become "color spaces" when certain details are specified[1]) describe a color in terms of three coordinates that represent hue, "something like saturation", and "something like luminance". Thus, these models are attractive at the human interface (for "picking" colors for drawing, etc.)

---

[1] And we will generally speak of them here as "color spaces" for convenience, assuming that a specific peculiarization is meant.

But unlike a color space that works in terms of true hue, saturation, and luminance, the definitions of the coordinates here admit of simplified calculations in translating color descriptions between these spaces and the RGB space in which the colors are usually actually held in computer memory, a substantial advantage at the time when these spaces were introduced (early in the era of computer graphics).

We often read that "the HSV color space can be represented as a cone (or a hexcone, a figure whose proper name is *hexagonal pyramid*)".

Similarly, we often read that "the HSL color space can be represented as a bi-cone (or a bi-hexcone)".

But what does that mean? The “gamut” of the HSV color space (the range of colors that it can represent), which is identical to the gamut of the RGB gamut color space, plots in the HSV color space as a full (circular) cylinder. The HSL gamut (which is also identical to the RGB gamut) plots in the HSL color space as a full (circular) cylinder. This (single) gamut plots in the RGB color space as a cube. In no color space we encounter in theoretical or practical colorimetry does this gamut plot as a hexcone or bi-hexcone. So what is going on here?

It turns out that this gamut does plot as a hexcone or a bi-hexcone in two different specialized color spaces we never encounter in theoretical or practical colorimetry. What is the significance of these two spaces? Well, they make the RGB/HSV/HSL gamut plot as a hexcone and a bi-hexcone, respectively, that’s what.

These color spaces have one coordinate that is the "lightness" coordinate of the color space of interest (V for HSV, L for HSL), and one that is in effect the “hue” coordinate, but the other coordinate is not one of those of the color space. Rather, the special color space works in a plane defined by "rotation of the axes" from the coordinate system of the RGB color space. So they are in a sense "hybrid (or “bastard”) color spaces".

What does that gambit do for us? Well, authors often draw upon these portrayals (without even stating the color space in which they emerge) to illustrate certain facts about this gamut, or maybe about the HSV or HSL color spaces themselves. These facts are rarely clearly articulated, and if we can see through the fog and discern what the author seems to be trying to say, we often find that it is not exactly true, or not true at all, or is in fact not illustrated exactly, or not illustrated at all, by the hexcone or bi-hexcone portrayal.

BACKGROUND

Color, color models, and color spaces

Color is the property of visible light by which one "flavor" of light may be distinguished from another. It is a property that is wholly defined in terms of the perception of the light by an animal (and we are usually most interested in the human as the observer).

In the case of human response to light, it is known that color is a "three-dimensional" property (in the mathematical, not geometric, sense). By that we mean that to describe a color we must state three numbers.[2] (Note that the concept of "color" in our work includes the aspect of the "brightness" of the light as well as the properties that lay people usually think of as being "color".)

Recall that there are any number of "coordinate systems" in which we can describe the location of a point in geometric space. These include Cartesian, cylindrical, and spherical coordinate systems.

Similarly, there are any number of "coordinate systems" through which we can develop three numbers to describe a particular color. These are often called "color models".

Originally, the term *color space* (borrowed from general mathematical theory) referred to the generic concept of the "three-dimensional space" within which any point represents a color. (It is the three-dimensional equivalent of the "number line" on which we visualize numbers as existing.)

Today, we most often find "color space" used to refer to a particular fully-specified scheme for numerically describing a color, starting with the use of a particular color model but also including a precise definition of how the values of the three coordinates that describe a color are to be determined.

The luminance-hue-saturation color model

Human perception of color can usually be most directly related to three attributes, luminance, hue, and saturation, simplistically described as follows:

---

[2] Note that in the case of animals other than *H. sapiens*, the number of dimensions in their color perception may be less or greater than three—sometimes up to eight!

- Luminance. This tells us the "brightness" of the light whose color is being described.[3]

- Hue. This is the attribute that distinguishes red from blue.

- Saturation. This is the attribute that distinguishes red from pink.

Note that hue and saturation together form the property of *chromaticity* (which is therefore a two-dimensional property in the sense we mentioned before). Chromaticity is in fact the property that most lay people think is meant by "color" (not recognizing that luminance is an attribute of color).

The RGB color model

An important color model (and the basis for an important family of specific color spaces) is the so-called "RGB model". This describes a color by stating the amounts of three light "sources" of different chromaticities (called "primaries") that are to be "added together" to produce the color of interest. Generally, these three primaries have chromaticities whose hues are identified by the familiar but imprecise terms "red", "green", and "blue", hence the designation of the model: RGB.

These quantities are stated by the coordinates R, G, and B. These do not, however, proportionally represent the "quantity" of light of each of the sources to be used in the color recipe; a nonlinear relationship is used for various reasons that are beyond the scope of this paper. (The definition of the specific nonlinear function is a part of the specification of a particular RGB color space).

Why might this color model have been developed, since it doesn't relate well to the human grasp of color perception? Because, in the early days of computer graphics, it related most directly to how we actually produced colors: in a tricolor cathode ray tube (CRT) display.

ABOUT GAMUTS

The gamut of a color space

In this area of work, the *color gamut* of a color space[4] refers to the entire range of colors that can be represented by the valid range of the coordinates of the space. Of course, to describe a color gamut we then have the problem of how we actually describe those colors.

---

[3] Note that in formal work "brightness" has a meaning slightly different from luminance.

[4] Often called just the *gamut*, but there is a risk there which we will see shortly.

For example, with no thought at all we can say that the gamut of any RGB color space of the conventional form is the colors represented by all possible values of R, G, and B in the range of 0-255. Sure—of course, but just what colors does this include?

So often when we try to describe, or even graphically present, the gamut of a color space (especially a newly-introduced one), we must do so in some well-defined color space, preferably one whose coordinates can be related to reconstructable perceptual principles. An example would be the CIE xyY color space, in which Y represents luminance (in terms of the standard scientific definition of luminance) and x and y are the coordinates of a plane on which chromaticity can be plotted (again in a way that can be traced to fundamental perceptual principles).

Color and chromaticity gamuts

When we speak of the gamut of a color space without further qualification, we should mean the color gamut, which has already been defined. If we want to "plot" this graphically, since the color space in which we plot it has three coordinates (not necessarily the same three coordinates as the color space itself—a pivotal matter in what will follow), a "three dimensional" plot is needed. The color gamut is represented by a three dimensional figure (sometimes called the *color gamut solid*), whose surface encloses the points representing all the colors that can be represented by the color space.

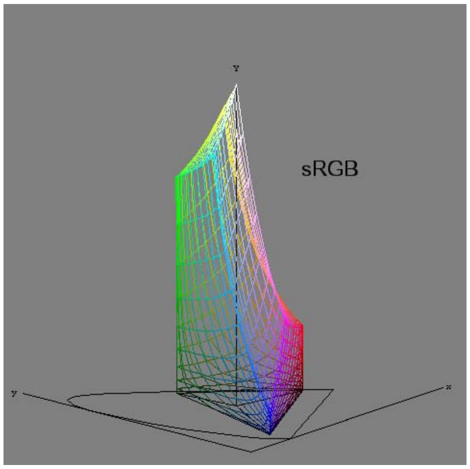

Figure 1. sRGB color gamut in the CIE xyY coordinate system

Of course, unless we want to work with modeling clay, we can't really do this physically. So we typically use a "two-dimensional view of a three dimensional object". Often we are able to do this with the aid of three-dimensional visualization software, so we can rotate the "color solid" in our view as we wish, in order to be able to comprehend its shape.

In figure 1, we see (in a two-dimensional view) the three-dimensional solid that represents the color gamut of the sRGB color space, as “plotted” in the CIE xyY coordinate system.

In figure 2, we see a two-dimensional view of the three-dimensional solid that represents the color gamut of any RGB-family color space, “plotted” in its own coordinate system. (Here, and from this point on, we will treat the range of R, G, and B to be 0-1 rather than 0-255.)

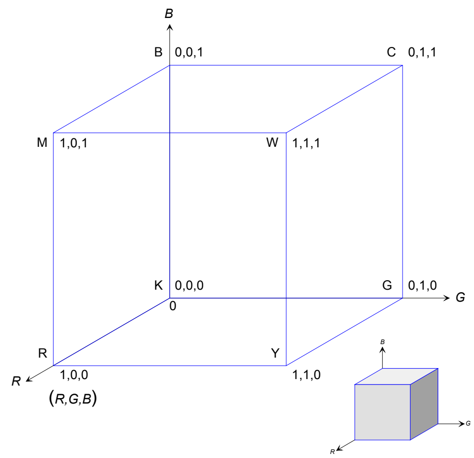

Figure 2. RGB color gamut in the RGB coordinate system

Chromaticity gamuts

Often, someone may say, “I realize that the color space here can only represent colors up to some arbitrary luminance, but all I want to know is what range of chromaticities it can represent. Can you show me that plot?”

Well, not really (although it seems that we see them all the time.) In most of the color spaces we use, as the luminance increases, the range of chromaticity that can be represented decreases (usually the full possible range of hue remains, but the available saturation decreases, perhaps differently for each hue).

We can easily see why by considering a typical RGB color space. Note the following:

- The luminance of a color represented in that color space is a weighted summation of the values of r, g, and b (R,, G, and B after the nonlinearity of the representation is backed out).
- 100% saturation (as saturation is defined in an RGB color space) only occurs if at least one of R, G, and B is zero and at least one of R, G, and B is non zero.

Now suppose we have an RGB representation in which only R is non-zero. This will give us a “100% saturated” color of a hue we call "red". The luminance will depend on the specific value of R. Now raise R to the maximum allowed by the color space (1.0 in our 0-1 scale). That will give a certain luminance (not the maximum allowed by the color space). How can we get a higher luminance 100% saturated red? Well, to raise the luminance (since we can't raise R any further) we would need to raise G or B. If we raise one, the hue will change from "red". If we raise both, the hue may change (depending on the proportions). In any case, the saturation will certainly decrease.

We can of course see this graphically by reviewing figure 1. We note that at the "bottom" of the color gamut solid (very low values of luminance), the cross-section of the solid occupies a certain area on the chromaticity plane. But if we consider some higher luminance, and imagine a plane parallel to the chromaticity plane through that luminance point on the Y axis, the cross-section (which represents the range of chromaticity available at that luminance) will be smaller.

In fact, we can see this here, on which the chromaticity gamut is shown both for small luminance (less than 0.072 of the maximum defined by the color space) and for a luminance of 0.500 of the maximum defined by the color space:

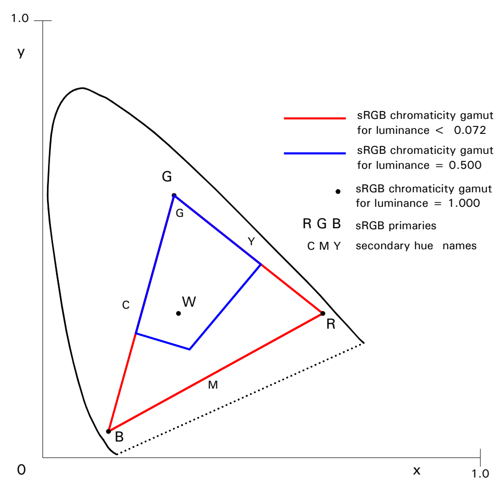

Figure 3. sRGB chromaticity gamut at different luminances

Now, when we see a plot of "the chromaticity gamut" of some particular RGB color space, what is that? Almost always, this is the chromaticity gamut that can be achieved only for very small luminances (such as that shown in red in the figure above). Thus such a plot is very misleading with regard to the overall "performance" of a system constrained by that color space.

Finally, note that 100% saturation as we have used it above has a special definition: the greatest saturation that can be represented in the RGB color model for the hue of interest. This is not the definition of saturation in formal colorimetric work, where 100% saturation is the greatest saturation that can exist (for the hue of interest) in the realm of human vision. Thus a color that has 100% saturation by the "RGB" definition has less than 100% saturation by the colorimetric definition.

In this article, since we are working almost exclusively with the gamut of the RGB color model, when we speak of saturation without further elaboration, we man the "RGB-basis" saturation.

IMPORTANT THINGS TO REMEMBER

Before we really examine the HSV and HSL color spaces, and then begin in earnest the quest for the elusive "hexcone", I would like to review a couple of important concepts.

1. A gamut may be established by a color space, but once this has happened, the gamut is just a "range" of colors—not a range of colors described under one color space or another. We may describe the gamut of the sRBG color space in the sRGB color space, or in the CIE xyY color space. The graphical portrayals will not be the same shape. But they will portray the same gamut. Any color (no matter how described) will either be in it or not.

2. We can choose to describe, or graphically portray, a certain gamut (such as the gamut of a certain color space) in any color space **within whose gamut the gamut we want to portray can completely fit**. We can portray the gamut of the Adobe RGB color space in the CIE xyY color space, or an L*a*b* color space[5]. But we can't portray the gamut of the Adobe RGB color space in the sRGB color space— the Adobe RGB color space includes colors that cannot be represented in the sRGB color space (that is, are outside **its** gamut).

3. Sometimes two color spaces will have the same gamut, usually when the one is derived from the other. We will soon learn that an HSV color space is a "transform", or "remapping", of an RGB color space. Its gamut is the same as the gamut of the RGB color space. The same is true of the HSL color space. And that is true regardless of the color space in which we might choose to "plot" that gamut.

THE HSV AND HSL COLOR MODELS

Introduction

Now we are ready for the real story.

---

[5] Strictly speaking, L*a*b is a color model, but when used for light (and not reflective color) and a reference white is associated with it (as called for by its defining equations), it becomes a color space.

Early work in computer graphics generally used a an RGB color space[6] to represent colors in computer memory and graphic files. But in "painting" programs, it was not usually ideal for the "artist" to have to choose colors by specifying its RGB coordinates. It was recognized that a color space based on the luminance-hue-saturation color model might be much better from that standpoint.

But any "scientifically valid" luminance-hue-saturation color space would involve complicated definitions of all three coordinates, with the result that fairly complicated calculations would be required to convert back and forth between images recorded in such a color space and the inputs needed to drive the display device, which is inherently RGB-oriented. At the time, processing power was at a premium (especially considering that these conversions had to be done for every pixel in an image).

As a result, various workers devised color models (and from them, specific color spaces) that used luminance-like, hue-like, and saturation-like measures with "simplified" definitions, which the human user could manipulate much like the "real thing", but which would allow much simpler transformation to and from an RGB color space.

An important contribution in this field was made in 1978 when Alvy Ray Smith of the New York Institute of Technology published a seminal paper, "Color Gamut Transform Pairs". In this paper he proposed two such "pragmatic compromise" color models, which he called "HSV" and HSL", from the symbols he used for their three coordinates.

In each case the coordinate labeled H represented hue in a direct way (not the way most used in color science, but having a consistent relationship to that definition).

The coordinate S represented "something like saturation" (different in the two schemes). However, in each case, it was actually fairly closely related to an suitable formal definition of saturation.

The coordinate V (or L) represented "something like luminance". The two had different definitions. Neither had a consistent relationship to the usual formal definition of luminance over the various colors that could be represented.

Smith's HSV color model has in fact come into wide use with almost the precise definition he proposed. (It is sometimes called "HSB",

---

[6] Since often the particular details weren't clearly specified, some may feel it is better described as only a color model.

where the coordinate "B" is the same as Smith's "V".) The S and V coordinates of this model represented (as I said above) "something like" saturation and luminance, but didn't track very closely with those actual colorimetric properties.

In Smith's HSL color space, the coordinates S and L tracked much more closely with real saturation and luminance, but the definitions would require far more computational effort in transforming colors expressed in that model to and from the RGB model. Smith himself said that accordingly it would not probably be as attractive as the HSV model in practical computer work.

Today, we have widespread use of a color model (often particularized to become a color space) called HSL. It is not the one proposed by Smith. In fact, it also has a "more pragmatic" definition of its S and V coordinates (much as in the HSV model), making them track less well with actual saturation and luminance than in Smith's HSL system, but simplifying the computations. (From here on, we will use "HSL" to mean the "modern" HSL color model, not Smith's.)

We'll look further into Smith's work later in the article.

Relationship to RGB

Both the HSV and HSL color models are "transforms" of the RGB model. That is, the standard definitions of their coordinates are given through equations that start with the R, G, and B coordinates of the color of interest.

"Model" vs. "space"

If we have a particular RGB color space (such as sRGB), then its transform into terms of the **HSV or HSL color model** produces a particular **HSV or HSL color space**.

Almost any actual use of the HSV or HSL model is in connection with a specific RGB color model, and thus constitutes an HSV or HSL color space.

Thus, and in order to best match the terminology we most often find in this matter, we will refer to the HSV and HSL <u>color spaces</u> in most of the remainder of this article.

DEFINITION OF THE COORDINATES

The coordinates

The coordinates of the HSV model are as follows:

H: Usually called “hue”. Essentially represents the hue of the color on a “color wheel” basis. It is not the usually scientific description of hue but still essentially a “genuine” hue descriptor.

S: Often called “saturation”. A property that is “something like saturation”, but which has a more manageable mathematical definition. I will describe it as a “pseudo-saturation” indicator.

V: Mnemonic for “value”. Often called “luminance” or “luminosity”. A property that is “something like (relative) luminance”, but which has a more manageable mathematical definition. I will describe it as a “pseudo-luminance” indicator.

The coordinates of the HSL model are as follows:

H: Usually called “hue”. Essentially represents the hue of the color on a “color wheel” basis. It is not the usually scientific description of hue but still essentially a “genuine” hue descriptor. It is essentially identical to H in the HSV model.

S: Often called “saturation”. A property that is “something like saturation”, but which has a more manageable mathematical definition. I will describe it as a “pseudo-saturation” indicator. It is defined differently from the S of the HSV model.

L: Often called “luminance” or “luminosity”. A property that is “something like (relative) luminance”, but which has a more manageable mathematical definition. I will describe it as a “pseudo-luminance” indicator. It is defined differently from the V of the HSV model.

Mathematical definitions

The three coordinates of either of these color models are defined as transforms of the R, G, and B coordinates of the color in the RGB color model.

H for either color space

In the HSV and HSL color spaces, the coordinate H is calculated by essentially the same principle, and its value is an angle (usually quoted in degrees). The result is that the three “primary” hues (red, green, and blue—defined as the hues of the RGB primaries) are given values of 0°, 120°, and 240°, respectively. The “secondary” hues, cyan, magenta, and yellow, are given the intermediate angles (180° from the “complementary” primary hue): 180°, 300°, and 60°, respectively. Other hues are given intermediate angles through a form of interpolation. We won’t worry about the actual formulas here.

V and S for HSV

In the HSV color space, V and S are defined this way:

$$V = \max(R, G, B) \tag{1}$$

$$S = \frac{\max(R, G, B) - \min(R, G, B)}{\max(R, G, B)} \tag{2}$$

where max(R,G,B) means the largest value among R, G, and B, and min(R,G,B) means the smallest value among R, G, and B.

V is a very poor indication of the actual relative luminance of the color. Colors for which V = 1.0 (if the specific RGB space implied were sRGB) could have relative luminances ranging from 0.072 to 1.0!

S here represents the saturation of the color of interest on an "RGB basis". S = 1.0 corresponds to 100% saturation, which is the maximum saturation for a color composed of a single primary or only two primaries. The actual scale from there down is hard to simply describe.

Note that S can have a value from 0-1 for any value of H.

V and S for HSL

In the HSL color space, L and S are defined this way:

$$L = \frac{\max(R, G, B) + \min(R, G, B)}{2} \tag{3}$$

$$S = \frac{\max(RGB) - \min(RGB)}{\max(RGB) + \min(RGB)} \qquad L \leq 0.5 \tag{4}$$

$$S = \frac{\max(RGB) - \min(RGB)}{2 - \left(\max(RGB) + \min(RGB)\right)} \qquad L > 0.5 \tag{5}$$

V is a better indication of the actual relative luminance of the color than H, but still pretty poor. At least, all colors for which L = 1.0 would have a relative luminance of 1.0! But colors for which V = 1.0 (if the specific RGB space implied were sRGB) could have relative luminances ranging from 0.072 to 0.783!

S here represents the saturation of the color of interest on an "RGB basis". S = 1.0 corresponds to 100% saturation, which is the maximum saturation for a color composed of a single primary or only two primaries. The actual scale from there down is hard to simply describe (but is different from the one for S in HSV).

Note that S can have a value from 0-1 for any value of L.

The gamuts of HSV and HSL

Both HSV and HSL color spaces are "transforms" of the RGB color space. Accordingly, the gamut of an HSL or HSV space is the same as that of the "underlying" RGB color space.

What does a "3D" plot of that gamut look like? Well, it depends on the coordinate system in which we plot the gamut.

If we plot it in the RGB coordinate system, it looks like this:

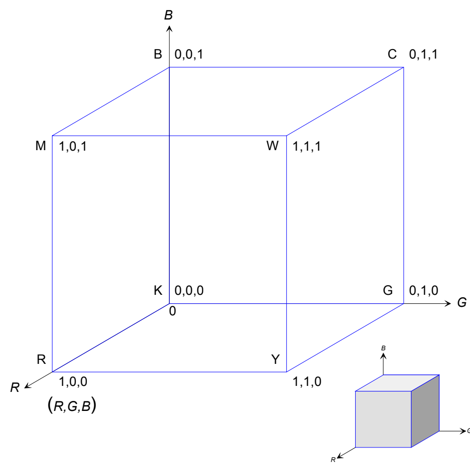

Figure 4. RGB/HSV/HSL gamut shown in RGB space

That is, the "gamut solid" is just a cube extending over all combinations of values of R, G, and B over the range 0-1.[7]

Some notes on the figure:

- The figure uses a form of the "oblique" projection to show this three-dimensional figure in our two-dimensional "page".
- I have shown the points that correspond to colors for which R, G, and B take on all combinations of the values 0 and 1, which are the corners of the "cube" representing the gamut of the color space.
- I have connected all pairs of points differing in only one coordinate by blue lines, which are the edges of the "cube". I have not suppressed nor distinguished the "hidden lines". Thus we have a basic "wireframe" representation of the cube.

---

[7] In practice, the R, G, and B coordinates operate on a scale of 0-255, reflecting the 8-bit representation that is customary (even prescribed by the standard for certain RGB color spaces), but here we will use the more mathematically fundamental expression of them on a scale of 0-1

- It is easy in this presentation to become confused about just what is toward us and away from us in actual 3-dimensional space (the well-known "cube inversion" optical illusion). To help us with this, in the lower-right of the figure I have shown the cube (in the same orientation) as an opaque solid with shaded faces.

HSV and HSL native coordinate systems

The native coordinate systems of the HSV or HSL color spaces are cylindrical coordinate systems. That is, the three coordinates are (in the conventional orientation of the coordinate system used for "academic" discussion):

- The height from the origin (called *h* in generic academic discussions of the cylindrical coordinate system)—used here for *V or L*.

- The radius from the *h* axis (*r*)—here used for *S*.

- The azimuth angle from some reference direction (Θ—upper-case Greek *theta*)—used here for *H* (and "red" is at the reference direction)

Now in either of these color spaces, the range of *V* or *L* is 0-1, the range of *S* is 0-1, and the range of *H* is 0°-360°. **Thus the gamut solid is a complete cylinder about the *h* (*V* or *L*) axis**.

We see that here:

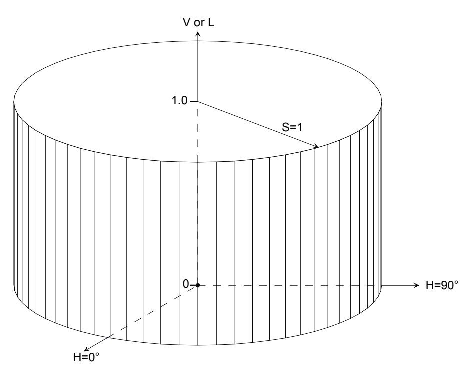

Figure 5. RGB/HSV/HSL gamut in HSV/HSL space

The HSV gamut and the "hexcone"

Almost always when the HSV color space is discussed, the author mentions something like:

"The HSV color space may be represented by a cone" (or "by a hexcone", a cute nickname for the solid figure properly known as a *hexagonal pyramid*).

What might this mean? It would seem that it means the gamut of the HSV space plots as a hexcone. But, as we saw just above, in the HSV coordinate system itself, the gamut plots as a cylinder. In the RGB coordinate system, it plots as a cube (recall that the HSV gamut is the same as the RGB gamut). In other color spaces we encounter in colorimetry, it plots as various other odd solid shapes. In fact, it does not plot as a bicone (I'll use the term for convenience) in any color space (or color model) you have ever encountered in theoretical or applied colorimetry.

Then what is the significance of the radial coordinate of the "hexcone" figure? It isn't S (or else the figure would be a cylinder).

And what is the significance of the tapering of the figure to just a point at its bottom? Clearly that is meant to visually covey that the available range of some property (represented by the radius of the figure) declines to zero as V declines to zero. But what property is that? Saturation? No, its available range doesn't decline as "lightness" decreases. Same for S. So we have here a serious mystery.

But the gamut does plot as a bicone in a very special coordinate system, one I'll describe in a moment. It is one you almost certainly have never heard of before. We never find it used in any area of colorimetry (except in this very specialized situation). In fact, its only purpose in life is this: to make the HSV gamut plot as a hexcone! Why would one want that? We'll find out a little later.

The special coordinate system

Let's develop this special coordinate system. We start with figure 4, which shows the gamut of the RGB color space (the gamut of the HSV space is exactly the same) in RGB coordinates (using an oblique projection to allow us to make a two-dimensional.

Because we use the same set of three letters for the three axes of the coordinate system and for the three primary chromaticities of the RGB color space, there is a risk of confusion. To avert that, I follow the usual mathematical convention of using italic letters for the names of axes and the corresponding coordinates.

The gamut as we plot it here is of course a cube, the locus of all points for which *R*, *G*, and *B* each take on all possible values from 0 to 1. At the origin (0,0,0 in *RGB* coordinates) we have the color "black" (which we symbolize "K", borrowed from the CMYK color space). At 1,1,1 we have the color "white" (W), by which we mean the *reference white color* for whatever specific RGB color space we may have in mind. If we don't have any particular one in mind, it is still "white" (we just don't know—or apparently care—exactly what chromaticity that is).

The edges of the cube that lie along the *R*, *G*, and *B* axes represent all colors in the gamut whose hues are those we consider to be red, green, and blue, respectively, fully saturated (in the "RGB" sense of fully saturated). Their luminance increases as we move further along the axis.

Lines running from the origin to the points 0,1,1, 1,1,1, and 1,1,0 (not drawn on the figure) represent all colors whose hues are those we consider to be cyan, magenta, and yellow, respectively, fully saturated. Their luminance also increases as we move further from the origin.

In fact, the three faces of the cube that touch the origin collectively carry all the fully saturated colors of the RGB color space (including for hues intermediate to the six we mentioned by name).

The six corners of the cube (other than black and white) represent the highest available luminance of the fully saturated versions of those six hues. They are (for simplicity) just labeled R, G, B, C, M, and Y. (R, G, and B are in fact the colors of the R, G, and B primaries (at "full-strength") as plotted in the RGB color space.)

The path R-M-B-C-G-Y-R ("broken field running" along six edges of the cube) carries all highest available luminance fully saturated colors (over the entire range of hues).

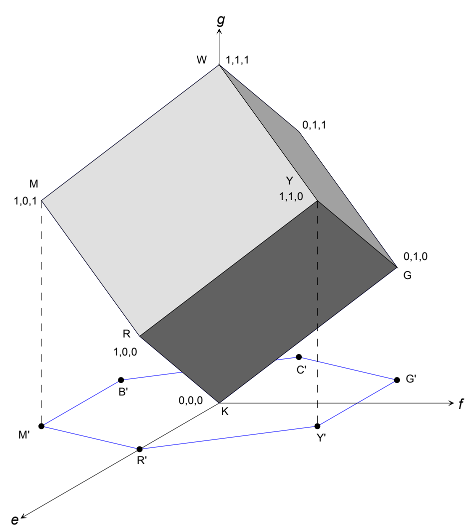

Figure 6. RGB cube rotated (*e*,*f*,*g* coordinate system)

Now, we will take this cube (and its associated set on R, G, and B axes and upend it, placing it within another set of three Cartesian axes. It would seem natural to perhaps call these *x*, *y*, and *z*, but to avoid any confusion with the *x* and *y* axes of the CIE xyY color space, I will arbitrarily call them *e*, *f*, and *g*. We will consider this color space with the *g* axis vertical (like *x* in our generic coordinate system), the *f* axis toward the right (like *y*), and the *e* axis toward us (but shown downward to the left in the oblique projection)—like *x*.

The cube (and its associated *R*, *G*, and *B* axes) is oriented with its original origin at the origin of the new coordinate system, its "1,1,1" corner upwards (the "major diagonal" of the cube lying along the *g* axis), and its original R axis lying directly above the *e* axis of the new coordinate system (but "headed upwards").If we consider the projections of the six color points R, G, B, C, M, and Y on the *e-f* plane (that is, plot their *e* and *f* coordinate values on that plane, as if the *g* coordinate were zero), we find that those projected points (I label them B′, G′, etc) form the vertexes of a regular hexagon. (I only show the projection lines for M′ and Y′ to save clutter.) This is a hint of something to come.

From a standpoint of colorimetry, what is the significance of the "*e-f*" plane? Is it a plane of chromaticity? No. Is it a plane of chrominance? No. It is just the plane that results from the manipulation I just performed.

Now we take the final step. We replace the scale of the vertical axis, initially the arbitrary geometrical coordinate *g*, with a scale running in V, the third coordinate of the HSV color space. We now have the "*efV*" coordinate system.

Note that V is defined thus:

$$V = \max(R, G, B) \qquad [1]$$

So any point for which either R, G, or B is 1.0 will now receive a V value (a "height" in the *ef*V coordinate system) of 1.0. The points in the cube for which either R, G, or B are 1.0 all lie on the path R-M-B-C-G-Y-R, which is the original "zig-zag up and down beltline" of the upended cube as seen in the *efg* coordinate system.

That is, in our *ef*V coordinate system, all the points on that ziz-zag beltline in the *efg* coordinate system are now elevated to a consistent "height" (V value) of 1.0. Thus they all now all lie in a plane parallel to the e-f plane that intercepts the V axis at V = 1.0.

Having done this, we find that our color gamut solid is a hexcone (base upwards, with its apex at the origin)!

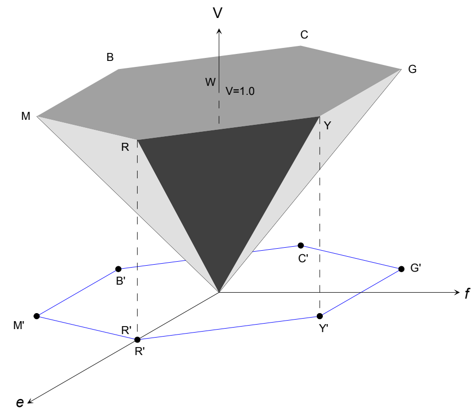

Figure 7. HSV hexcone

It is perfectly erect; it looks a little tilted in the figure, but this is an artifact of the oblique projection. It is more "squat" here than we usually see it portrayed. Here we have used consistent scales for e, f, and V, and this is the shape that results. If we wanted a more "svelte" hexcone, we could adopt a different scale for V.

The flat top of the hexcone is made up of the three faces of the RGB gamut cube that extend to the white point (now each warped to triangles, and seamlessly joined into a hexagon). This surface does not contain (except at its very outside edges) fully-saturated colors.

The rest of the hexcone surface is made of the original three "lower" faces of the rotated RGB cube (each of them now "creased" about a line running diagonally across it to become two hexcone faces). Each of them contains only fully-saturated colors, and thus the entire "outer" surface of the hexcone contains only fully saturated colors.

The luminance generally increases as we go up, but not in a direct and consistent way.

We see that the projections on the *e-f* plane of the six vertexes of the top plane form a hexagon. It is a regular (symmetrical) hexagon, although we can't tell this easily from the oblique projection view.

The *efV* coordinate system

Now, what is this *efV* coordinate system? Of what use is it?

The *e-f* plane doesn't really provide for the representation of any meaningful colorimetric quantity. The *efV* coordinate system is of no use that I can discern in grasping the significance of colors expressed in HSV form, or in grasping the scope of the HSV color gamut (which, after all, is just the RGB gamut).

Its only function today is to provide a premise for plotting the HSV gamut as a hexcone, and we cannot find why this is a useful tool either.

And what good is the hexcone? What facts about the HSV coordinate system, or the HSV color space, or its gamut, does it help us visualize? Beats me.

The HSL gamut and the "bi-hexcone"

Similarly, almost always when the HSL color space is discussed, the author mentions something like:

"The HSL color space may be represented by a bi-cone" (or "by a bi-hexcone").

What might this mean? It would seem that it means the gamut of the HSL space plots as a hexcone. But in the HSL coordinate system itself, the gamut plots as a cylinder. In the RGB coordinate system, it plots as a cube (the HSL gamut is the same as the RGB gamut). In other color spaces we encounter in colorimetry, it plots as various other odd solid shapes. In fact, it does not plot as a bicone (I'll use the term for convenience) in any color space you have ever encountered.

Then what is the significance of the radial coordinate of the "bi-hexcone" figure? It isn't S (or else the figure would be a cylinder).

And what is the significance of the tapering of the figure to just a point at its top and bottom? Clearly that is meant to visually covey that the available range of some property (represented by the radius of the figure) declines to zero as L declines to zero. But what property is that? Saturation? No, its available range doesn't decline as "lightness" decreases. And although the range of saturation declines as luminance increases above a certain point (a different point for each chromaticity), it doesn't decline with an increase in L. Same for S (whose range is 0-1 at any value of L). So here again we have a serious mystery.

But the gamut does plot as a bicone in a very special color space, one I'll describe in a moment. It is one you almost certainly have never heard of before. It is different from, but conceptually very similar to, the special coordinate system, described just above, that gives birth to the HSV "hexcone". Like the one we saw before, we never find it used in any area of colorimetry (except in this very specialized situation). In fact, its only purpose in life is this: to make the HSL gamut (or the RGB gamut) plot as a bi-hexcone!

The special coordinate system

The special coordinate system is developed this way. The early parts of the process are identical to those for the HSV case.

1. We start with the HSL gamut (which is the same as the RGB gamut), plotted as a cube in the RGB coordinate system.

2. We rotate the RGB coordinate system to form a new Cartesian coordinate system which I call *efg*. [This is precisely the same coordinate system I use at this step for the HSL situation.]

3. We then replace the inherent geometric scale of the g axis with the scale of the HSL coordinate L (which is defined differently than the V of HSV). This now gives the *efL* coordinate system.

4. This elevates the points R, B, and B to an "altitude" (L value) of 0.5, and depresses the points Y, C, and M to an altitude (L value) of 0.5. (As this happens, the faces of the cube all get "creased" about a diagonal across them, each then forming two faces of the resulting figure.) Note that all the points retain their *e* and *f* coordinates.

5. This has "warped" our cube into a "bi-hexcone". The part up to L=0.5 is just like the hexcone of HSV except that its height is only 0.5 units. The part above that (from L=0.5 to 1.0) is an identical hexcone but inverted.

6. Thus the entire plot of the HSL gamut **in the *efL* coordinate system** is a "bi-hexcone".

We see this result here:

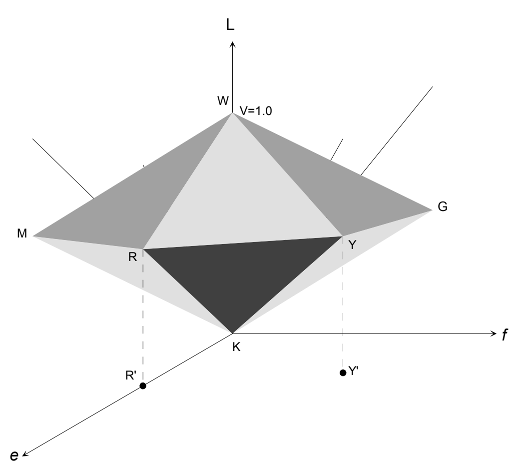

Figure 8. HSL bi-hexcone

The *efL* coordinate system

Now, as before. we ask, "What good is the *efL* coordinate system? Can we usefully express colors in it?

The answer, as before, seems to be that its only function is as a framework in which the HSL gamut (RGB and HSV gamuts, too) plots as a "bi-hexcone".

And what good is the bi-hexcone? What facts about the HSL coordinate system, or the HSL color space, or its gamut, does it help us visualize? Beats me.

Hexcones on the street

Discussions of the bi-hexcone (or bi-cone) are almost inevitably accompanied by a graphic presentation. Figure 9 is a popular one for the bi-hexcone, found (among other places) in the Wikipedia article on the HSV and HSL color spaces.

The writer of the relevant passage there doesn't hold to the "hexagonal" outlook for the figure. He claims that the "bi-cone" is a meaningful "conceptualization" of the actual cylindrical figure. Again, I cannot grasp that notion.

The colorful disks shown are cross-sections of the cone (as the author visualizes it) at various values of L. They reveal that (as we know), for values of L from 0.5 up to 1.0, the range of available saturation decreases with increasing L. So perhaps the upper portion of the cone is in effect the "envelope" of available saturation. But this means that the radius there would have to be saturation, not S. (The range of S remains 0-1 all the way to the pinnacle.)

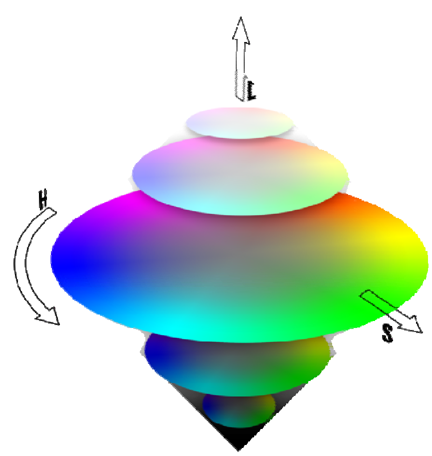

Figure 9. HSL hexcone shown in Wikipedia article

For values of L below 0.5, as seen in the disks, the range of saturation is not limited. So in that range, the cone isn't an envelope of available saturation. Neither does it reveal the available range of S (which remains 0-1 all the way to the "foot").

So the concept behind the bicone remains a mystery to me.

Sometimes illustrations of the hexcone are labeled with these coordinates (in cylindrical coordinate form):

- For the azimuth angle, *hue*
- For the radius, *saturation*
- For the "axis", *luminance*

So does that mean that the figure is actually the HSL gamut plotted in true hue-saturation-luminance form? No. That figure would have quite a different shape. For one thing, its lower portion would be a circular cylinder. And its "beltline" would not lie in a plane (owing to the different weightings given R, G, an B in reckoning true luminance).[8]

Of course, it is easy to say that, since this figure is attributed to the HSL color space, these terms are really just metaphors for the HSL coordinates H (which in fact tracks with hue), S (which is "something like saturation", and L ("something like luminance"), used to help the reader understand what is going on.

But, as we saw before, the hexcone is not plotted in H, S, and L coordinates. If it were, the figure would be a cylinder, not a bi-hexcone.

In fact, if we want to think of the *efL* color space (the birthplace of the hexcone) in terms of cylindrical coordinates, they would be:

- Angle (Θ): H
- Height (z): L

(Sounds good so far, doesn't it!)

- Radius (r)?

[8] We can dispose of this untidiness by adopting a "pseudo-luminance" in which all three RGB coordinates are weighted equally. Note that this is not L (or V). But this does not dispose of the other paradoxes of the bi-(hex)cone.

But what does the radius represent. Well, if you must know, the radius coordinate for the *e*-*f* plane turns out to be this:

$$r = \sqrt{\frac{(2R - G - B)^2 + 3(G - R)^2}{6}} \qquad (6)$$

But recall that S is defined (in the HSL system) this way:

$$S = \frac{\max(RGB) - \min(RGB)}{\max(RGB) + \min(RGB)} \qquad L \leq 0.5 \qquad [4]$$

$$S = \frac{\max(RGB) - \min(RGB)}{2 - (\max(RGB) + \min(RGB))} \qquad L > 0.5$$

So clearly the radius coordinate is not S (and certainly not saturation either).

A RATIONALE FOR THE CONES

A nice-sounding explanation sometimes given for the use of the “cone” and “bi-cone” metaphors (not their “hexagonal” cousins) is this (I will insert comments in square brackets):

For HSV: The bottom surface of the HSL cylinder contains only a single color [true—black, which doesn’t come in differing saturation and hue]. Thus we might as well shrink the bottom surface of the solid figure to just a point. This gives us a cone.”

For HSL: “The top surface of the HSL cylinder contains only a single color [true—black, which doesn’t come in differing saturation and hue]. Thus we might as well shrink the top surface of the solid figure to just a point. Similarly, the bottom surface of the HSL cylinder contains only a single color [true—black, which doesn’t come in differing saturation and hue]. Thus we might as well shrink the bottom surface of the solid figure to just a point. Doing both gives us a bi-cone figure.”

One problem with both these notions is that an infinitesimal distance above the “floor” of the two cylinders (even at V or L=0.0001), there is a full range of hue and pseudo-saturation (S) contained by the gamut. So the rationale for shrinking the bottom half of the cylinder to a cone doesn’t play.

But in any case, we are not working on shaping a child’s toy. We are using a three dimensional figure to tell us something. If this “reshaping” is to be useful, it must improve what the figure tells us. And just what is it that the radius of this new figure represents (even in a “sort-of” way)? It’s not S. Its not saturation (which does not decline at low levels of luminance or low levels of L either).

COLORFULNESS

There is an outlook that at first seems as if it could lead to a meaningful bicone representation of the HSL gamut (keeping L as the vertical coordinate).

It involves the property of "colorfulness" of a color. Here's the concept.

A color with a high luminance but low saturation has a low *colorfulness* (pink, for example, doesn't seem as "colorful" as red). But also a color with a high saturation but a low luminance has a low *colorfulness* (a very low luminance fully saturated red doesn't seem as "colorful" as a higher-luminance fully saturated red).

Quantitatively, colorfulness is often defined as twice saturation times luminance. [9]

But of course in the world of HSV and HSL, we don't insist on coordinates that actually correspond to real colorimetric properties. Instead, we use proxies that are easier to use in calculations (like L instead of luminance and S instead of saturation).

Could we define a metric that was something like colorfulness, but which, used as the radius of our figure, would produce a bi-cone?

We can start with the notion of a pseudo-saturation, a property that was "something like saturation" but had a simple mathematical definition. This definition seems reasonable in that vein:

$$PS = \frac{\max(R,G,B) - \min(R,G,B)}{\max(R,G,B)} \qquad (7)$$

Now, since (real) colorfulness is twice (real) saturation times (real) luminance, perhaps we could reasonably define pseudo colorfulness as twice pseudo saturation (PS) times pseudo luminance (L). That would give us:

$$C = \frac{(\max + \min)(\max - \min)}{\max} \qquad (8)$$

where C represents our pseudo colorfulness and, for conciseness, "max" means "max(R,G,B)" and "min" means "min(R,G,B)" (I'll use this latter convention from time to time hereon when space is tight).

---

[9] Video folks will recognize this as essentially the same thing is the *magnitude of the chrominance*, or *chroma* (although there is a distinction regarding nonlinear representation).

If we now plot the HSL (or RGB, or HSV) gamut in HCL coordinates, do we get a bicone? No. We get a figure that is a cone at the bottom and a dome-like figure at the top (not a hemisphere—more nearly a *convex oblate paraboloid of revolution*).

THERE MUST BE A BICONE IN THERE SOMEPLACE

Is there any property we could choose for the radius of our figure that would produce a figure that was at least close to a hexcone?

Well we could adopt this definition of the radius:

$$r = \max(R,G,B) - \min(R,G,B) \tag{9}$$

If we do that, then if we plot our gamut in the HrL coordinate system, we get a *bona fide* bi-cone. (I'll leave the proof up to the reader.)

But what is the significance of that measure? Well, it is very vaguely reminiscent of colorfulness—perhaps pseudo pseudo colorfulness. But of course its only claim to power is that it produces the fabled bi-cone.

SMITH AND THE HEXCONE

Alvy Ray Smith, who introduced the HSV color model in his famous 1978 paper, actually introduced the notion of the hexcone representation of the HSV gamut. He was so taken by it that he called the HSV color model "the hexcone model".

Smith developed the hexcone using essentially the process I described earlier. Then, working in that coordinate system (which he never really described as such), he validated (or perhaps even initially conceived) his equations defining S and H. It was apparently a "working space" that inspired him.

He thought of the plane I call the "*e-f*" plane in the equivalent circular coordinate form, with an angle (which was H) and a radius, which he never gave a designation. (It is, for example, not S, as we have seen.) He pointed out that, if we "slice" the figure at a height corresponding to some arbitrary value of V, the boundary of that section is a hexagon, and for all points on it S = 1. Thus S = 1 over the entire surface of the hexcone (except for its "lid").

But of course S = 1 over the entire surface of the HSL gamut plotted in HSL coordinates, where the plot is a cylinder, (except for its "lid" and "bottom"). So the "S = 1 all over the surface" property, while interesting, is something we do not have to go to the hexcone paradigm to attain.

In any case, Smith never said, "the HSV color space is represented by a hexcone". It was an intermediate tool, like falsework used while framing a complicated building.

THE BI-HEXCONE—SPAWN OF THE HEXCONE

Smith did not introduce the HSL color model we discuss here. He did introduce another color model whose coordinates he called H, S, and L. Those coordinates actually tracked very well with the properties hue, saturation, and luminance. But they were necessarily defined by equations that would have led to significant computational load in transforming between a color space based on this model and one based on RGB. And thus, Smith noted in his paper that it would perhaps not be attractive in the actual computer practice of the time.

No notion of a "bi-hexcone" emerged in Smith's derivation of this color model. But he did utilize a "working space" in which the "base plane" was in fact parallel (in the geometric sense) to the plane I call "*e-f*". He defined this as passing through the RGB points 0,0,1, 0,1,0, and 1,0,0 (the points R, G, and B). Because those points of the plane formed a triangle, Smith thought of the plane as a triangle, and called his HSL model the "triangle model".

But evidently some "student" of Smith, taken by his hexcone outlook for the HSV color model, followed a similar route (evidently paralleling what I did above) to develop the "bi-hexcone" outlook (and the "*efL* coordinate system). It may even have been whoever developed the "modern" HSL color model—I don't know the story of that.

As with Smith's hexcone, the bi-hexcone would have been, at best, a "working space" within which the HSL model might have been incubated. But, just as for the hexcone, it would only at best have been "falsework", and really had no significance once the HSL model was defined.

CONCLUSION

Hexcone and bi-hexcone

The infamous "HSV hexcone" and "HSL bi-hexcone" are three-dimensional solids that represent the RGB, HSL, and HSL color gamuts in two peculiar "hybrid" coordinate systems. The only value of these coordinate systems seems to be that they produce the hexcone and bi-hexcone solids often spoken of.

Although it is often said that these figures facilitate visualization of important properties of these two color spaces, or of their gamuts, no explanation of just how this works, or what these properties are, has been advanced.

I believe that all of these notions emerge from unthinking adoption, extension, and promulgation of the "hexcone" metaphor used by Smith in his paper as a "working space" in which he confirmed (perhaps even initially conceived) his equations for the HSV coordinates.

I urge authors not to cite these figures in discussions about the HSV and HSL color spaces.

Cone and bi-cone

The cone we often see cited as "representing the HSV color model" exists within another specialized coordinate system, which I call HCV. It can serve to graphically illustrate the fact that the available range of *colorfulness* (not an HSV property) within the gamut declines as L approaches its minimum.

The bi-cone we often see cited as "representing the HSL color model" exists within yet another specialized coordinate system, which I call HCL. It can serve to graphically illustrate the fact that the available range of *colorfulness* (not an HSL property) within the gamut declines as L approaches either its maximum or minimum.

I urge authors who wish to exploit these cone-based graphic metaphors to be sure to describe what they represent. Introducing them by saying only "they represent the HSV/HSL color models" is meaningless and does not equip the reader to understand their significance.

#

APPENDIX A

The Canon "HSL" color Model

The Canon Picture Style Editor software package allows the user to selectively make changes in the color representation of an image and to generate custom "picture styles" that apply certain "recipes" for such changes.

In this program one may select colors (for various purposes) in one of three color models (color coordinate systems): "Lab" (meaning L*a*b*), RGB, and "HSL". They designate the three coordinates of the "HSL" system as hue, saturation, and luminosity.

But the HSL color model here is not the same as what I will now call the "standard HSL" color model (as discussed at length in the body of this article.

The description of this system in Canon's reference documentation for the program points out that in this "HSL" system, the coordinate S can only be set over its full range of 0-1 (actually 0-100%) when L is 0.5 (50%). As L rises from 0.5 to 1.0, the maximum available S declines, reaching zero range (S can only be 0) when L=1.0. As L decreases from 0.5 to 0, the maximum available also S declines, reaching zero range (S can only be 0) when L=0.[10]

This is clearly inconsistent with the properties of the "standard HSL" system, in which the range of S is 0-1 for any value of L. Thus this model must be defined by different equations than the standard HSL model.

Canon's figure illustrating this relationship implies that the reduction of range is linear, as a result of which the gamut of this "HSL" system, plotted in its native coordinate system, would be a bicone (two right circular cones joined base-to-base).

However, reverse engineering of the relationship between the RGB and the HSL representations of various colors shows that such is not quite the case. The decline in the maximum available S is not linear with the variations of L above and below 0.5.

In fact, the gamut solid of this HSL color model (plotted in its native coordinate system) has this cross-section:

---

[10] Ah! Smells like colorfulness!

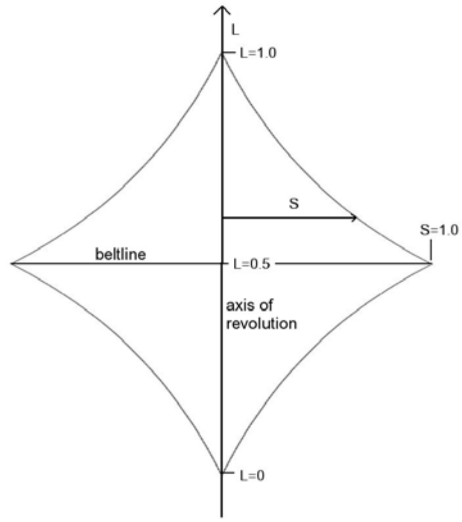

Figure 10. "Canon HSL" color space gamut in its native coordinate system

This figure is broadly evocative of a bi-cone, but its surfaces are "concave". (The "beltline" is at L=0.5)

Through reverse engineering, I have determined the equations defining S the "Canon HSL" color model. It appears that the definition of L is precisely the same as for standard HSL. It seems as if the definition of H is very much the same as for standard HSL.

For comparison, I will first restate the equations defining L and S for the "standard" HSL color mode.

L and S for standard HSL

In the standard HSL color space, L and S are defined this way:

$$L = \frac{\max(R,G,B) + \min(R,G,B)}{2} \qquad (1)$$

$$S = \frac{\max(RGB) - \min(RGB)}{\max(RGB) + \min(RGB)} \qquad L \leq 0.5 \qquad (4)$$

$$S = \frac{\max(RGB) - \min(RGB)}{2 - (\max(RGB) + \min(RGB))} \qquad L > 0.5 \qquad (5)$$

L and S for Canon HSL

In the Canon HSL color space, L and S are defined this way:

$$L = \frac{\max(R,G,B) + \min(R,G,B)}{2} \qquad (10)$$

From here on, for conciseness, I will use “max” and “min” to mean max(R,G,B) and min(R,G,B), respectively.

For S in the “lower hemisphere”:

$$S = \frac{\max - \min}{2 - (\max + \min)} \qquad L \leq 0.5 \qquad (11)$$

For S in the “upper hemisphere”:

$$S = \frac{\max - \min}{\max + \min} \qquad L > 0.5 \qquad (12)$$

Note that this color space is very much like conceptually like the HCL color space I introduced in the body of the paper, where the radial component, C, is a pseudo-colorfulness. There, I mentioned that there were in use various quantitative definitions of (actual) colorfulness. My definition of pseudo-colorfulness (my “C”) was broadly based on a particular one I chose.

It may be that Canon had in mind that their coordinate S (which they call “saturation”; it isn’t even close to that) would be a “pseudo-colorfulness”, but based on a different definition of (actual) colorfulness proper than I used.

The Mind of Canon!

#