Representing Colors as Three Numbers

Editor: Frank Bliss

Maureen C. Stone
StoneSoup Consulting

RGB in graphics is both a way of specifying color and a way of viewing color. Graphics algorithms manipulate RGB colors, and the images produced by graphics algorithms are encoded as RGB pixels and displayed on devices that render these pixels by emitting RGB light. Colored images are also used to specify color in graphics. These images may be captured by cameras or scanners, interactively drawn using tools such as Adobe PhotoShop, or algorithmically generated. But, what do all of these RGB values mean with respect to color perception? How does the RGB triple captured by a digital camera relate to the RGB pixels displayed on a monitor? How does the RGB triple selected with an interactive color tool relate to the RGB triple used to color an object in a 3D rendering?

How do three numbers, such as RGB or XYZ, represent color perception, and how are these representations related to each other and to physical color? When do they fail?

Most computer graphics texts and tutorials provide a description of human color vision and measurement as defined by the CIE tristimulus values, XYZ. Often missing, however, is an in-depth discussion of the relationship between the different applications of RGB and XYZ, and any discussion of color models beyond trichromacy. The goal of this tutorial is to provide a complete, concise analysis of RGB color specification and its relationship to perceptual and physical specifications of color, and to introduce some models for color perception beyond tristimulus theory.

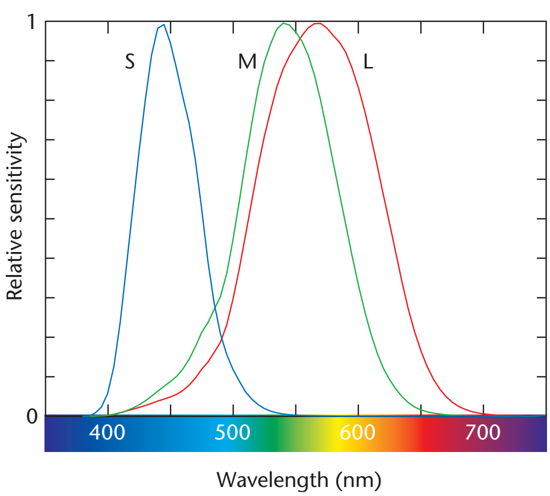

1 Spectral sensitivity curves for the short (blue), medium (green), and long (red) cones. The colored band shows approximate wavelength colors. (Reprinted by permission from A K Peters Ltd.)

Representing color as three numbers

That color can be represented by three numbers—whether RGB or XYZ—is a direct result of the physiology of human vision. Electromagnetic radiation whose wavelength is in the visible range (370 to 730 nanometers) is converted by photopigments in the retinal cones into three signals, which correspond to the response of the three types of cones. This response is a function of wavelength and is described by the spectral sensitivity curves for the cones, as Figure 1 shows.

Colored light can be represented as a spectral distribution, which plots power as a function of wavelength. (Other fields, such as signal processing, plot spectra as a function of frequency, which is the inverse of wavelength.) The cones convert this to three cone response values (*L, M, S*)—that is, the cone sensitivities in the long, medium, and short wavelength regions—defined by integrating the product of the spectral sensitivity curves and the incoming spectrum. Figure 2 shows this process.

Two important principles follow from this process:

- *Trichromacy*: all spectra can be reduced to precisely three values without loss of information with respect to the visual system.
- *Metamerism*: any spectra that create the same trichromatic response are indistinguishable.

This means that two different spectra will look the same if they stimulate the same cone response. Figure 3 shows two metameric spectra.

It's important at this point to distinguish between the perception of color and the creation of color. In practice, both can be described by three values, but the discus-

Published by the IEEE Computer Society

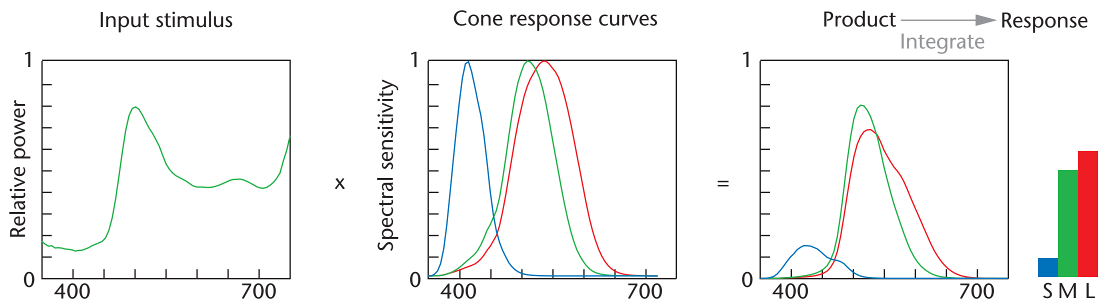

2 **Multiplying a spectrum times the cone response curves and integrating the result creates the basic color signal to the brain. The height of the bars reflects the strength of the three signals. (Reprinted by permission from A K Peters Ltd.)**

sion up to this point has only covered perception. To use a digital color analogy, converting spectra to three cone response values corresponds to using a camera to create image pixels. We have not yet discussed how to display these pixels in a way that recreates the original sensation of color. The bridge to creation is provided by the color-matching experiments that underlie the CIE standards for measuring and specifying color.

Color-matching experiments

Let's construct a color-matching experiment as follows. Choose three primary lights different from each other (such as red, green, and blue), which can vary in intensity. Then, take a set of reference colors such as the monochromatic colors of the spectrum, or colors generated by filtering a white light. An observer combines and adjusts the primary colors to create a result that matches the reference color. Figure 4 shows this schematically.

Once the match is made, a color can be defined by describing the amount of each primary needed to match it. These are called the *tristimulus values* for the color and must be defined with respect to a specific set of primary lights and a specific observer. Anyone with the same set of primary lights can recreate the color from the tristimulus values, and it's guaranteed to appear the same to the specified observer.

The analogy to color display systems—where RGB values are used to create colors—should be clear. Color is created by mixing the light emitted by the red, green, and blue primaries (for example, the phosphors of a CRT). RGB pixel values define the mixture, and hence uniquely define the color. Given an *identical display*—one whose phosphors produce the same spectra—the same RGB value will produce the same color.

Tristimulus values would be of limited use if they only applied to specific colors matched by a specific observer. To create a more general result, we need to apply the tristimulus values to a broader range of colors and observers.

Color-matching experiments in the 1920s and 1930s established that a large percentage of test subjects created color matches that were similar enough to establish a model of a statistically derived standard observer.[1] This doesn't mean that there aren't measurable differences, but that the deviation is small enough that the standard is useful. Therefore, any color matching performed by someone whose perception is normal can be usefully applied by all other normal observers.

Expanding the tristimulus definition beyond specifically matched colors follows from the additive nature of

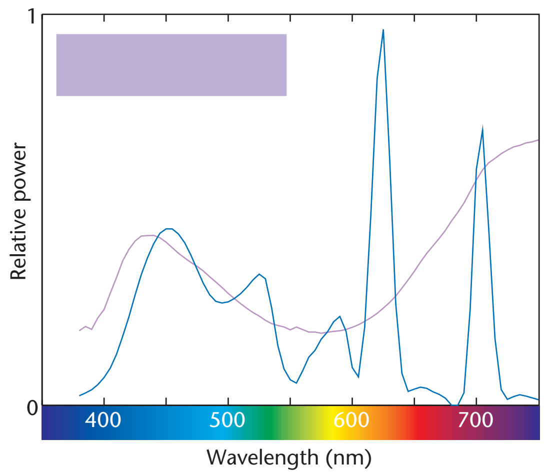

3 **Two different spectra that create the same color, illustrating the property of metamerism. The resulting color is approximately that of the purple bar in the upper corner. (Reprinted by permission from A K Peters Ltd.)**

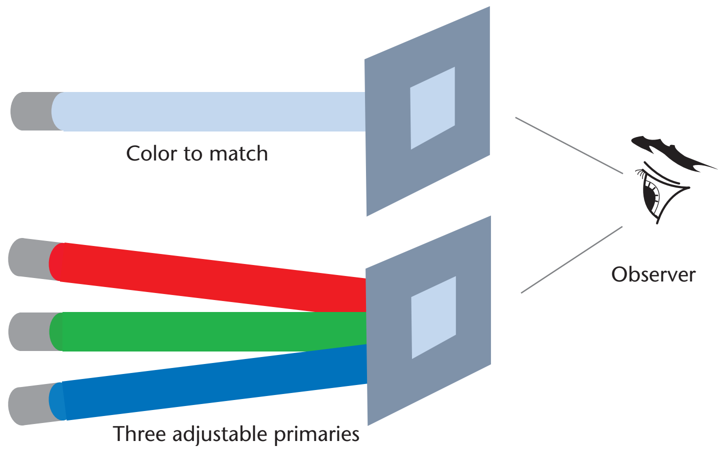

4 **Color-matching experiment. An observer adjusts three primary lights to match each sample color. (Reprinted by permission from A K Peters Ltd.)**

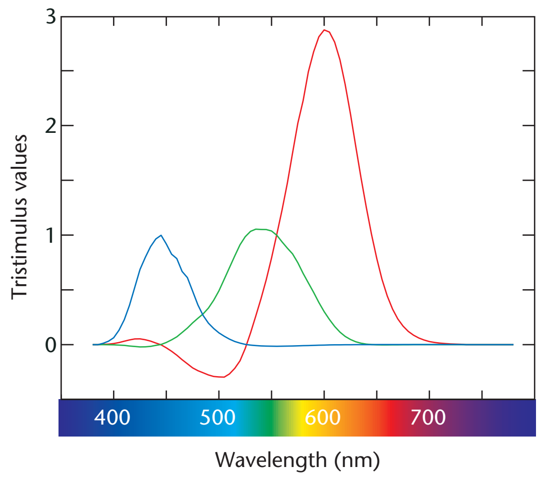

5 Color-matching functions defined by Stiles and Burch[1] using three monochrome light sources: R is 645 nm, G is 526 nm, and B is 444 nm. (Reprinted by permission from A K Peters Ltd.)

colored light and the additive nature of tristimulus values. The spectrum created by combining several light sources is the sum of their individual spectra. Suppose an observer defines tristimulus values for each of the individual lights. Can we predict the tristimulus values of their sum? The answer is yes. The tristimulus values of the summed lights are the sum of their tristimulus values. To spell this out mathematically, for a given spectrum, S, let the tristimulus values be R, G, and B (or RGB). If RGB_1 matches S_1, and RGB_2 matches S_2, then $RGB_1 + RGB_2$ will match $S_1 + S_2$. Grassmann first formalized this principle, now called Grassmann's additivity law.[1]

This is an extremely important result, which makes it possible to use a finite set of color matches to specify an infinite set of colors. Any spectral distribution can be modeled as a weighted sum of monochromatic colors, or colors specified by light energy at a single wavelength. (In practice, we use an extremely narrow band of wavelengths, usually created by splitting white light spatially into its different colors with a prism or diffraction grating. We then use a slit to select as narrow a band of wavelengths as desired.) Given the tristimulus values for each of these colors, the tristimulus value for the spectrum is the weighted sum of the tristimulus values for the monochromatic colors.

How is this applied? An observer uses three primary colors to match a set of monochromatic colors across the visible spectrum. These results create three color-matching functions, such as those shown in Figure 5, where the value of each function represents the relative amount of each primary color.

What about colors that cannot be matched by the set of primary lights? The solution is to use "negative light." For example, if the blend of primaries is too reddish, even with no red shining on the blend, the observer must shine some of the red primary on the reference to make the match. The amount of red added to the reference color is the negative red included in the match. This makes it possible to match all colors with any set of three distinct primary colors. The curve for the red primary in Figure 5 shows this use of negative light.

To create the tristimulus values for an arbitrary spectrum, first multiply it by the color-matching functions, then integrate the result to get the total relative weights of the primaries, which are the tristimulus values for that color.

By definition, the original spectrum and the spectrum created by combining the primary colors weighted by the tristimulus values match—they look the same color. This means that the two spectra are *metamers*—they produce the same cone response. If different spectra create the same set of tristimulus values, they must also be metamers. We therefore get the powerful conclusion that two spectra that produce the same tristimulus values are indistinguishable—they look the same color. This is the foundation for color measurement, or *colorimetry*.

To summarize, this is what we know so far about representing color as three numbers. On the perception (input) side, we have the cones, which convert spectra to the three cone response signals in the visual system. On the creation side, we have the observation that all colors can be matched by combining weighted sums (including negative weights) of any three distinct primary colors. We can use this matching process to create a set of color-matching functions, which we can apply to any spectrum to generate the matching weights, or tristimulus values. Color-matching functions and tristimulus values are strikingly similar to the spectral response functions of the cones and the cone response values. Their application is similar, and in both cases, if the integrated response is the same, the perceived color is the same. This is not a coincidence, but let's first look at the CIE standard observer, which is the basis for modern colorimetry.

CIE standard observer

In 1931, the Commission Internationale de l'Eclairage (CIE, or International Commission on Illumination) standardized a set of color-matching functions that form the basis for most color measurement instruments used today. They averaged experimental work from two independent sets of color-matching experiments performed on small visual fields (2 degrees), to create the 1931 standard observer, or the 2° observer. This standard observer statistically represents the average color-matching results for the human population having normal color vision. The 2-degree field is an important part of this specification. There is a second CIE standard observer, called the 10° observer, which was standardized in 1964, and should be used for fields larger than 4 degrees. Most digital imaging applications, however, use the 1931 standard. Berns provides a full discussion on the derivation of the CIE standards and their application.[2]

Figure 6 shows the color-matching functions for the 1931 standard observer. These functions have a number of convenient features for color measurement, which are used to generate the CIE tristimulus values, XYZ. They are positive throughout their range, making it possible to implement them using three physical filters. The middle one, $\bar{y}(\lambda)$, was crafted by the CIE committee to match the CIE standard luminous response function, whose integral corresponds to the perception of brightness. Therefore, its integral, Y, is equivalent to perceived brightness, or *luminance*.

The curves in Figure 6 are not the direct result of color-matching experiments, but are mathematical transformations of the experimental curves. How do we transform a set of color-matching functions? As hinted in the discussion of Grassmann's law, tristimulus values follow the rules of linear algebra. They are additive, and they also scale: For spectrum S and its tristimulus values RGB and a scalar value k, the tristimulus values for kS are $kRGB$. Therefore, the mathematics of linear algebra can transform color-matching values. Tristimulus values are vectors in a linear, 3D space defined by the primary lights. Transforming to a new set of primaries is simply a change of basis; it requires only the value of the new primaries with respect to the old ones. This is the link that ties XYZ to the physical RGB values of a computer display and is the foundation for all quantitative characterization of RGB color spaces.

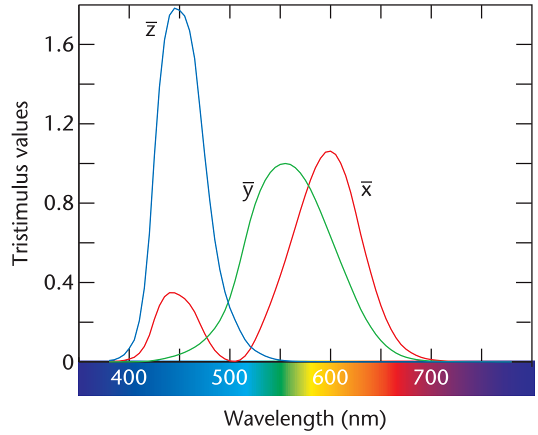

6 CIE recommended color-matching functions (1931), also called the CIE 1931 standard observer. (Reprinted by permission from A K Peters Ltd.)

Transformations between RGB and XYZ

Characterizing an RGB color space with respect to XYZ is straightforward. First, measure the R, G, and B primaries to determine their values in XYZ. Then, set up the corresponding linear transformation, which can be represented by a 3 × 3 matrix as follows:

$$\begin{bmatrix} R & G & B \end{bmatrix} M = \begin{bmatrix} X & Y & Z \end{bmatrix}$$

$$M = \begin{bmatrix} X_R & Y_R & Z_R \\ X_G & Y_G & Z_G \\ X_B & Y_B & Z_B \end{bmatrix}$$

Figure 7 shows both the linear transformation from RGB to XYZ, and the nonlinear transformation from XYZ to the familiar CIE chromaticity representation, xy (see the "CIE Chromaticity Diagram" sidebar, next page).

This ideal characterization assumes that the tristimulus values for $RGB = (0, 0, 0)$ are also $XYZ = (0, 0, 0)$, and that R, G, and B specify light values that are linear with respect to intensity. That is, equal steps in the color primaries create equal steps with respect to intensity, or, equivalently, that scaling the digital value for a primary color simply scales its spectral output by the same amount.

To convert between two RGB color spaces, convert from RGB_1 to XYZ using the RGB_1 characterization matrix, and from XYZ to RGB_2, using the inverse of the RGB_2 characterization matrix.

Most RGB color spaces are bounded, with RGB values limited to the range 0 … 1, as shown by the unit color cube in Figure 7a. This describes a volume of colors, or *color gamut*. Colors transformed from one RGB space to another might end up outside the gamut of the target color space and must be approximated. In computer graphics applications, the colors are either scaled or clipped to bring them inside the gamut boundary. In color reproduction applications, the process of transforming out-of-gamut colors involves perceptually defined 3D transformations and projections, and is called *gamut mapping*.[3]

RGB, XYZ, and cones

The previous sections lead to the conclusion that all color-matching functions (for the same observer) are linear transformations of one another. What about the cone response functions? Are they also linearly related to the color-matching functions? The answer is a qualified yes. The curves in Figure 1, which are the measured response of the cone cells, are not precisely linear transformations of the color-matching functions in Fig-

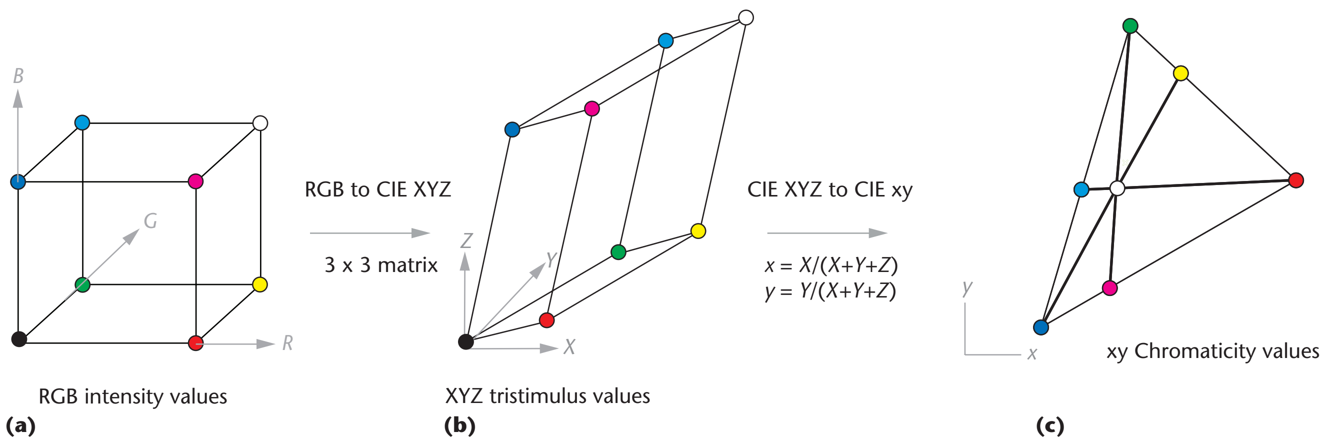

7 (a) Converting the RGB color cube to XYZ. (b) This transformation scales, rotates, and skews the cube. The XYZ values can be mapped to the chromaticity diagram via the usual formulas. (c) The result is a triangle, whose vertices are red, green, and blue. White and black lie on the same point near the center. (Reprinted by permission from A K Peters Ltd.)

CIE Chromaticity Diagram

The CIE chromaticity diagram, shown in Figure A, is the most common way of visualizing CIE tristimulus values. The chromaticity coordinates, (x, y), are computed from the XYZ tristimulus values by $x = X/(X + Y + Z)$ and $y = Y/(X + Y + Z)$. The result of this transformation is to factor out the brightness of the color. A spectral distribution plots as a point, independent of its power (brightness). Also, all its metameric matches, which by definition have the same tristimulus values, will plot to the same point. While this particular diagram, which corresponds to the tristimulus values generated by the 1931 standard observer, is the most familiar, any tristimulus values can be plotted this way.

The monochromatic spectral colors lie along the horseshoe-shaped path, which is called the *spectrum locus*. All visible colors lie within the shape bounded by this path and the line that connects the two ends, which is called the *purple line*. The closer a color lies to the spectrum locus, the more saturated—vivid, colorful, pure—it is. Colors that appear white lie near the center of the chromaticity diagram, near the black-body curve, which plots the colors created by heating a black-body radiator.

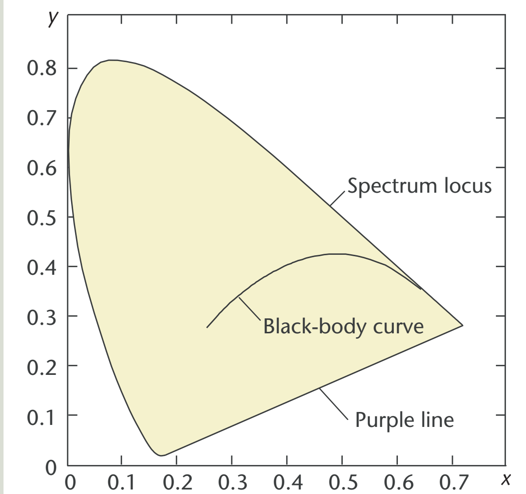

A CIE 1931 chromaticity diagram showing the spectrum locus, the purple line, and the black-body curve. (Reprinted by permission from A K Peters Ltd.)

Color on Displays

The mapping from pixel values to the color seen by the viewer is defined primarily by the light emitted by the display. In an ideal display system, the tristimulus values for this light can be calculated by decoding the pixels using the display's nonlinear transfer function (often called a *gamma function*, whether or not it is strictly a power function), then multiplying by a 3×3 matrix calculated from the tristimulus values of the primary colors. The appearance of the colors, however, will also be affected by ambient light reflected from the display, and by more subtle aspects of visual adaptation that impact the viewing of displays. Additional complications are introduced for many non-CRT display devices, where

- black ($RGB = (0, 0, 0)$) is bright enough to affect the appearance of colors (especially noticeable on projection displays), and/or
- the spectral distribution of the primary colors changes as the colors become dark, shifting the chromaticity and invalidating the 3×3 matrix (primarily a problem for LCD displays and projectors[1]).

In spite of these problems, the use of a 3×3 RGB characterization provides a useful foundation for making color in computer graphics more predictable and controllable. Inexpensive instruments that measure displays (http://www.ecolor.com) provide an affordable way to accurately transform between RGB and XYZ.

Reference

1. G. Marcu and K. Chen, "Gray Tracking Correction for TFT-LCDs," *Proc. Soc. for Information Science & Technology (IS&T)/Soc. for Information Display (SID) 10th Color Imaging Conf.*, IS&T, 2002, pp. 272-276.

ure 6. However, once proper compensation is made for the optical properties of the eye, such as wavelength absorption in the cornea, there is a linear relationship.[4]

Research in color vision has defined linear transformation matrices to map from XYZ to the cone response values, LMS. While several published transformations exist—reflecting the ongoing research in this area—the one recommended for the latest CIE color appearance model, CIECAM02,[5] is

$$\begin{bmatrix} L \\ M \\ S \end{bmatrix} = \begin{bmatrix} 0.7328 & 0.4296 & -0.1624 \\ -0.7036 & 1.6975 & 0.0061 \\ 0.0030 & 0.0136 & 0.9834 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

Nonlinear RGB

One of the few fields in digital color where digital RGB values are assumed to be linear with respect to intensity is 3D computer graphics. Display technology, digital video, and image encoding standards map pixel values to nonlinear functions of intensity, which are perceptually more uniform, to provide more efficient use of finite encoding values.[6] For displays, these transfer functions are often called *gamma functions* (see the "Color on Displays" sidebar).

The colorimetry of a nonlinear RGB color encoding is only slightly more complicated than that of a linear one. The transformation from RGB to XYZ requires the addition of three 1D functions (usually tables) that map from pixel value to intensity. The transformation from XYZ to RGB requires the inverse functions. Figure 8 shows a pair of typical encoding and decoding functions.

Figure 7b shows the shape that a linear RGB color cube transformed into XYZ becomes. If the cube represents a nonlinear RGB color space, the resulting shape will be identical—only the spacing of the pixel values within the shape will be different. In a linear space, equal steps remain equal. In a nonlinear space, they do not.

In most nonlinear RGB spaces, the shapes of the three pixel-to-intensity functions are identical—they vary only in scale. If the pixel-to-intensity functions are not scalar multiples of each other (which can easily occur in uncalibrated color systems, such as displays), then equal values of *R*, *G*, and *B* will not create a gray that is simply a darker value of white, but will have some tint depending on

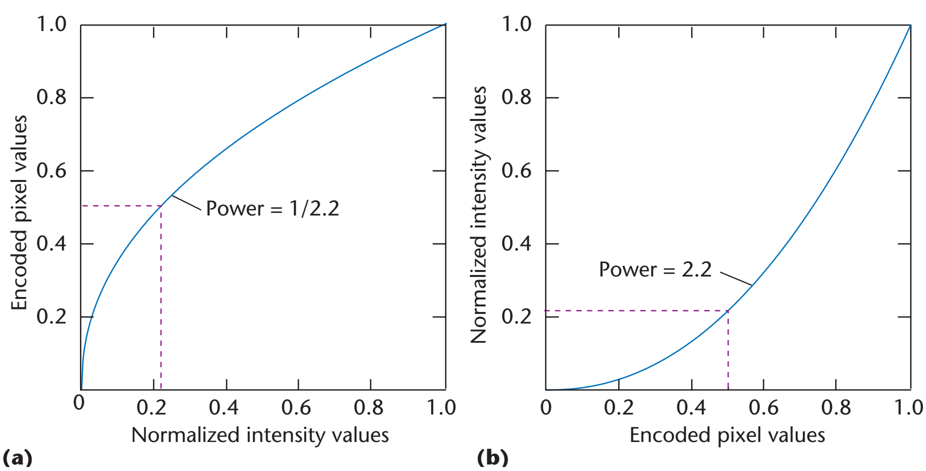

8 Typical (a) encoding and (b) decoding functions, based on a simple power function. (Reprinted by permission from A K Peters Ltd.)

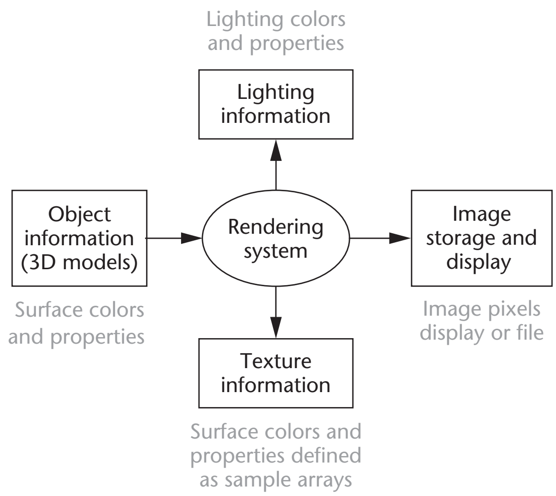

9 Basic components of a graphics rendering system. Objects, lights, and textures combine to create a rendered image. (Reprinted by permission from A K Peters Ltd.)

which primary is dominant in the blend. This color will often vary over the range of gray values. For example, dark gray could appear bluish while a middle gray might appear brownish.

This means that the set of colors between black and white defined by equal values of the primaries will not run along the diagonal of the color cube and its transformation, nor will they plot to a single point on the chromaticity diagram as they do in the ideal case. The transformation between RGB and XYZ is accurate, however, as long as the pixel-to-intensity transfer functions are correctly measured.

Color in computer graphics

Computer graphics algorithms manipulate digital representations of lights and surfaces to create an image, as shown in Figure 9. Most graphics systems represent all colors as RGB triples. In some cases, surface and light colors are represented as spectra to provide more accurate modeling of lights and surfaces. In all cases, the resulting image is an array of RGB pixels.

Computer graphics systems generally treat all RGB values as if they described points in a common color space. This works in a closed system, where the only way RGB values are viewed is on the system's monitor, whose physical characteristics create the visual appearance of the colors. All renderings are adjusted to make colors and textures look right on that particular display. However, even viewing a rendered image on a different display begins to expose the weakness of this simple model.

A better solution is to define a standard RGB specification for the rendering system, characterized with respect to XYZ. The standard should be a linear RGB color space, as a linear space is needed for rendering. Given such a standard and its characterization, it's possible to map all other color specifications into and out of this space. This provides a quantitative link between the graphics system, physically defined color, and other characterized RGB color spaces.

During rendering, an object's color is calculated as the product of its color with the color of the light shining on it. If both are represented as spectra, then their product is another spectrum whose color must be converted to RGB values to be stored in the image. This is accomplished by converting to XYZ, and from there to the standard RGB. If the objects and lights are colored with RGB triples, their product is also an RGB triple. While the color of a light can be reduced to three tristimulus values, the color of the surface cannot, so there is no direct physical model for this type of calculation. (Surface reflectances are sometimes expressed as XYZ, but the physical meaning of this is the color of the surface illuminated by an equal energy white light.) Accepting the approximation inherent in this form of rendering, all RGB values are simply defined with respect to the standard color space.

Color is also applied during rendering using textures and environment maps. The RGB colored images used for these purposes are either scanned images, digital photographs, or are images created interactively using an image editing system like Adobe Photoshop, and stored as an image file. Such RGB values are most often nonlinearly encoded, so at minimum, they should be decoded to be linear RGB values before being used in rendering. In addition, a full RGB-to-RGB transformation can be applied, if known.

Depending on the goals of the graphics system, there are two basic choices for standard color space: one that matches the display and one optimized for a wider range of media. A display-centric approach—where the standard color space is a linear space constructed from the same RGB primaries as the display—is the most conve-

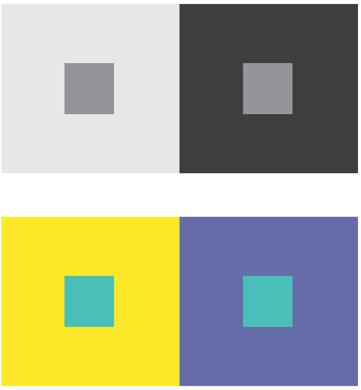

10 **In each pair, the inner patch is the same color, but appears different because of simultaneous contrast. (Reprinted by permission from A K Peters Ltd.)**

nient. Some of the colors used in film and printing, however, cannot be produced on a typical computer display. In these applications, a larger RGB color space that includes the color gamuts for all media of interest is more appropriate. A more complete discussion of RGB color management for graphics and its relationship to the color management systems common in the graphic arts is available in Stone.[3]

Limits of trichromatic encoding

A spectrum can be encoded as three numbers for comparison to other spectra (as in color measurement), or to generate a metameric match to the spectrum (as in color display). Such a representation is limited in two fundamental ways. It does not capture the full perception of color, nor can it be used in applications where spectra must be multiplied, such as computing the light reflected from an object or filtering light from a scene with a digital camera to create RGB pixels.

A trichromatic encoding, whether RGB or XYZ, specifies the appearance of a single colored light (either direct or reflected) viewed on a neutral background by a viewer with normal color vision who is fully adapted to the viewing environment. This is a precise but narrow specification. If any of these conditions are violated, two colors specified by the same tristimulus values can appear different. For example, each pair of identical patches in Figure 10 appear different because they are displayed on different backgrounds, a color appearance phenomenon called *simultaneous contrast*.

Beyond tristimulus values

The CIE recommendations that define CIE XYZ, and the standard functions needed to implement it, were first defined in 1931. They are the foundation for all modern color measurement and specification. Perception beyond tristimulus theory is an ongoing area of research, and progress over the past decade or so has provided some practical ways to model color perception beyond simple color matching. There is not space in this tutorial to do more than introduce a few of the key findings, but further detail is available in Fairchild's book.[7]

One key finding in this area is the precise specification of the cone response functions and the mathematics that link them to tristimulus measurements, such as the XYZ to LMS transformation previously presented.

The cone response can be used to model the way the cones adjust to changes in illumination, a process called *adaptation*. The visual system adapts to both the brightness and the color of the ambient lighting, redefining "white." The von Kries model for adaptation is a simple weighting of the cone response values in response to the overall illumination, or white. Think of it as independent gain controls for the L, M, and S cone responses. The weighting can be computed from the response to the ambient white. If L_1, M_1, and S_1 are the cone response to a color under one state of adaptation (white = L_{w1}, M_{w1}, and S_{w1}), the cone response to the same stimulus under a different state of adaptation (white = L_{w2}, M_{w2}, and S_{w2}) is L_2, M_2, S_2, which can be computed as:

$$L_2 = (L_{w2}/L_{w1})L_1$$
$$M_2 = (M_{w2}/M_{w1})M_1$$
$$S_2 = (S_{w2}/S_{w1})S_1$$

Suppose we have a set of tristimulus values, X_1, Y_1, Z_1, defined for one state of adaptation, such as pixels on a monitor viewed in a dim environment. What are the tristimulus values that would appear the same color in bright light (for example, a print viewed in sunlight)? Given the LMS values for white in the two different viewing environments (L_{w1}, M_{w1}, S_{w1} and L_{w2}, M_{w2}, S_{w2}) we can use the von Kries adaptation transformation to compute X_2, Y_2, Z_2 as follows:

$$\begin{bmatrix} X_2 \\ Y_2 \\ Z_2 \end{bmatrix} = \boldsymbol{A}^{-1} \begin{bmatrix} L_{w2}/L_{w1} & 0.0 & 0.0 \\ 0.0 & M_{w2}/M_{w1} & 0.0 \\ 0.0 & 0.0 & S_{w2}/S_{w1} \end{bmatrix} \boldsymbol{A} \begin{bmatrix} X_1 \\ Y_1 \\ Z_1 \end{bmatrix}$$

Where $\boldsymbol{A}$ is a matrix that transforms between XYZ and LMS, as previously described.

Another application of LMS is to simulate color blind vision, which is caused by the absence of a signal from one of the cones. The pixel colors in an image are transformed to LMS space, then projected along the axis that corresponds to the missing cone.[8]

The LMS signals define the opponent color channels, which are the first transformation from the cone response (effectively RGB encoding signals) to a perceptual encoding (hue, lightness, and colorfulness). The achromatic channel (A) is defined as the weighted sum of the cone response signals. The red-green opponent channel (R-G) is computed from the difference between the red (L) and green (M) channels, and the yellow-blue opponent channel is the difference between the yellow (L+M) and blue channels (S). Figure 11 shows a perceptual color space defined by these axes. Hue is defined as the angular dimension, saturation as the radial one.

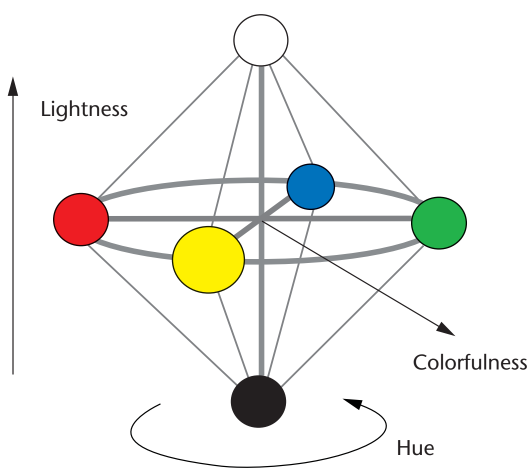

11 **Perceptual color spaces are organized by lightness, hue, and colorfulness. (Reprinted by permission from A K Peters Ltd.)**

Many perceptually organized color spaces exist.

Some, such as the Munsell color ordering system (http://www.munsell.com), were designed as physical color chips. CIELAB and CIELUV are computationally derived from tristimulus values, plus the tristimulus values for a reference white. All perceptual models have a similar lightness axis and place hues in the same order around the hue circle, though the spacing of the hues varies. Most perceptual color systems are designed so that distance in the color space is proportional to perceptual distance. That is, they make it possible to compute how different two colors appear. This is in contrast to tristimulus representations, which can only specify whether two colors match exactly. For more information on the design and history of perceptual color spaces and color difference formulations, see Berns.[2]

A common feature of all perceptually organized spaces is a unique specification for white and black, in contrast to tristimulus values, where a wide range of values can appear white (depending on adaptation). In a physically defined color space, white is defined by the light illuminating the colored chips. In CIELAB and CIELUV, unique white is defined by dividing the tristimulus values of the stimulus by the tristimulus values of the reference white. For example, L*, which is the lightness axis for both spaces, is defined as a function of Y/Y_{white}. Mathematically, this is similar to the von Kries model for adaptation, but computed in XYZ space rather than LMS space. It is tempting to convert tristimulus values to CIELAB or CIELUV, then use this specification to define colors in a different viewing environment, such as the display-to-print example previously mentioned. Perceptually, however, this is not as accurate a way to model adaptation as using LMS. The color spaces RLAB and LLAB have been defined as simple extensions to CIELAB that include a von Kries adaptation transformation.[7]

sCIELAB[9] is an extension of CIELAB that uses spatial filtering in LMS to more accurately model the color of image pixels, which are both small and surrounded by pixels of various colors. Classic colorimetry assumes a 2-degree sample (about 5/8-inch across when viewed at a distance of 18 inches), viewed on a neutral background. Using sCIELAB rather than CIELAB for pixel colors gives a more accurate way to evaluate how similar two images appear. sCIELAB has been combined with the LMS projection algorithm previously described to create a tool for simulating on a display how colored images would appear to a person with color blindness (http://www.vischeck.com).

The goal of the CIE color appearance models, CIECAM97s and CIECAM02,[5] is to create models for color appearance that accurately predict perception, yet are computationally practical enough to apply to color reproduction problems such as gamut mapping and image quality assessment. These models take as input CIE XYZ for the stimulus and the reference white, plus parameters that describe the immediately surrounding color, the overall level of illumination, and to what degree the observer is adapted to the illumination. The output of these models are quantitative values for hue, lightness, brightness, chroma, saturation, and colorfulness, all of which are precisely defined as part of the modeling process. In applications, the goal is to preserve these quantities across different transformations of media and viewing environments.

Conclusion

Using CIE colorimetry and a bit of linear algebra, the RGB color values used in computer graphics can be defined in a way that provides a direct link to perception. This creates a quantitative foundation for manipulating these values with respect to physical specifications of color such as color displays. It creates a link to scientific specifications of color, such as those created by modeling real surfaces and lights. It provides a way to integrate computer-generated imagery with color management systems, such as those used in the graphic arts. Finally, it provides an opportunity to integrate with current research in color appearance modeling, to provide renderings optimized for human color perception. ■

Acknowledgments

I thank A K Peters Ltd. for permission to use the figures in this article, which are from my book, *A Field Guide to Digital Color*.

References

1. G. Wyszechki and W.S. Stiles, *Color Science*, 2nd ed., John Wiley & Sons, 1982.
2. R.S. Berns, *Billmeyer and Saltzman's Principles of Color Technology*, 3rd ed., John Wiley & Sons, 2000.
3. M.C. Stone, *A Field Guide to Digital Color*, A K Peters, 2003.
4. B. Wandell, *Foundations of Vision*, Sinauer Associates, 1995, pp. 96-97.
5. N. Moroney et al., "The CIECAM02 Color Appearance Model," *Proc. Soc. for Information Science & Technology (IS&T)/Soc. for Information Display (SID) 10th Color Imaging Conf.*, IS&T, 2002, pp. 23-27.
6. C.A. Poynton, *Digital Video and HDTV Algorithms and Interfaces*, Morgan Kaufmann, 2003.
7. M.D. Fairchild, *Color Appearance Models*, 2nd ed., Wiley-IS&T, 2005.
8. H. Brettel and F. Viénot, "Web Design for the Colour-Blind User," *Colour Imaging, Vision and Technology*, L. MacDonald and M. Ronnier Luo, eds., John Wiley & Sons, 1999, pp. 55-72.
9. X. Zhang and B. Wandell, "A Spatial Extension of CIELAB for Digital Color Image Reproduction," *Proc. Soc. for Information Display 96 Digest*, IS&T, 1996, pp. 731-734.

Maureen C. Stone *is an independent consultant at StoneSoup Consulting. Prior to that, she was a principal scientist at Xerox Palo Alto Research Center. Her research interests include digital color, information visualization, and related topics in perception, graphics, and interaction. Stone has an MS in electrical engineering from the University of Illinois and an MS in computer science from the California Institute of Technology. She is an associate editor in chief of* IEEE Computer Graphics and Applications. *Contact her at stone@stonesc.com.*