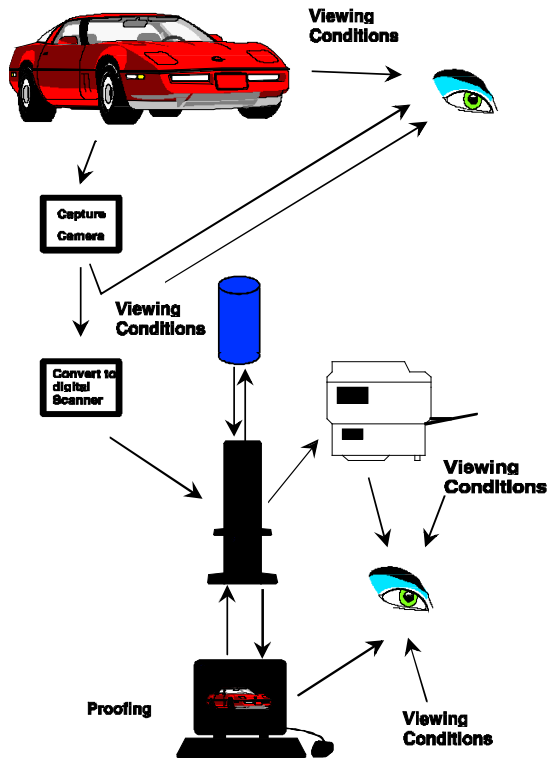


The Color Reproduction Problem


Consider a digital system for reproducing images of the real world:



An observer views an original scene under some set of viewing conditions: a certain illuminant, a state of visual adaptation, a specific surround, with the subject occupying a particular portion of the visual field, etc. This scene is captured by photochemical means, providing an alternate reference¹. In fact, because of the evanescence of real-world scenes compared with the handy persistence of their photochemical replicas, a photographic reproduction is usually treated as the primary reference. A scanner converts the reproduction to digital form, which is stored, manipulated, displayed on monitors of many types; possibly converted to a photographic negative or transparency; and printed on a wide range of printers. When a viewer observes each of the reproductions, he does so under a set of viewing conditions that are unlikely to duplicate the reference conditions.

When simulated images are produced or reproduced, little changes but the references:

¹Electronic cameras that can directly deliver digital representations bypass the photochemical step, but leave the rest of the diagram unaffected.



If the colors are defined in the simulation, they, along with a set of virtual viewing conditions, become the standard of comparison. However, in many simulations, the simulated image is not defined colorimetrically, and the reference is the scene as viewed by its designer on a monitor and the conditions under which the viewing takes place.

The goal of both kinds of color reproduction systems is the same: to have the displayed and printed versions of the image reproduce visual sensations caused by the reference (success is particularly difficult to measure if the reference is the abstract simulation, since it cannot actually be seen by anyone). Before discussing the theory and practice of color reproduction, we should consider exactly what “reproduce” means in this context. Hunt (Hunt, *RofC*, pages 177-197, Hunt 1970) defines six possible objectives for color reproduction, which are paraphrased here:

Spectral color reproduction, in which the reproduction, on a pixel-by-pixel basis, contains the same spectral power distributions or reflectance spectra as the original.

Colorimetric color reproduction, in which the reproduced image has the same chromaticities as the original, and luminances proportional to those of the original.

Exact color reproduction, in which the reproduction has the same chromaticities and luminances as those of the original.

Equivalent color reproduction, in which the image values are corrected so that the image appears the same as the original, even though the reproduction is viewed in different conditions than was the original.

Corresponding color reproduction, in which the constraints of equivalent color reproduction are relaxed to allow differing absolute illumination levels between the original and the reproduction; the criterion becomes that the reproduction looks the same as the original would have had it been illuminated at the absolute level at which the reproduction is viewed.


Preferred color reproduction, in which reproduced colors differ from the original colors in order to give a more pleasing result.

Spectral color reproduction, were it practically achieved, would have undeniable advantages over the other approaches: for a reflection print, viewing the print in any given illuminant would yield the same result as viewing the original in that illuminant, and reproduced images would match the original not only for people with normal color perceptions, but also for color-deficient observers. However, there are a host of implementation problems with this approach: the necessary sensors do not exist outside of a handful of research laboratories; spectral information occupies far more storage than tristimulus data; and practically output devices and media are not obtainable. There have been attempts to deal with the storage requirements (some Wandell paper), but practical systems embodying this approach remain many years away. This chapter will concentrate on the tristimulus techniques for color reproduction.

For all but the most restricted choices of originals, exact color reproduction requires an impractical dynamic range of luminance; we will not consider it further. Colorimetric color reproduction is a practical goal for originals with a limited range of colors, and will result in a satisfying match if the viewing conditions for the original and the reproduction are similar. Many computer imaging systems meet this criterion, such as those in which all images are viewed on CRT monitors in dim rooms. However, in general, a robust system will produce acceptable results over some broad range of viewing conditions, so equivalent or corresponding color reproduction becomes the necessary goal. Unfortunately, the psychology of human vision is not sufficiently well understood to do a perfect job of correcting for changes in viewing conditions, but we do have enough information to achieve adequate results in many circumstances.

Organization of digital color reproduction systems

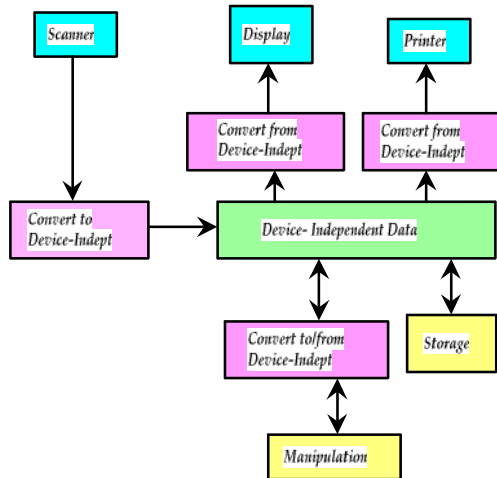
Up until the 1990s, almost all digital color reproduction systems used a fairly simple system model:



In this model, the image data is transformed in the scanner or in associated processing to the native form of the output device, and thus consists not so much as a collection of colors, but a set of recipes that the output device uses to make the colors desired. The transform between the native scanner representation and that of the output device is performed on a pairwise basis; changing either device means computing a new transform. Good results can be achieved in this manner, but successful application of this simple system model requires performing only standardized processing upon the image, having only one kind of output device, knowing the characteristics of that device, and adopting standard viewing conditions. The strength of this approach is control: the operator of the scanner knows exactly what will be sent to the printer. However, this method is extremely

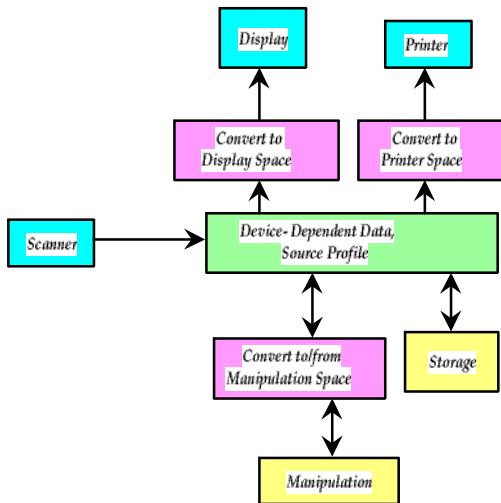
inflexible: if the output device or the viewing conditions change, the original must usually be rescanned.

The widespread use of small computer systems has promoted an appetite for more flexible approaches, and the increasing processing power available can provide the means to slake the hunger. A consensus system model has yet to emerge, but there are two commonly-proposed alternatives. The first is presented below:



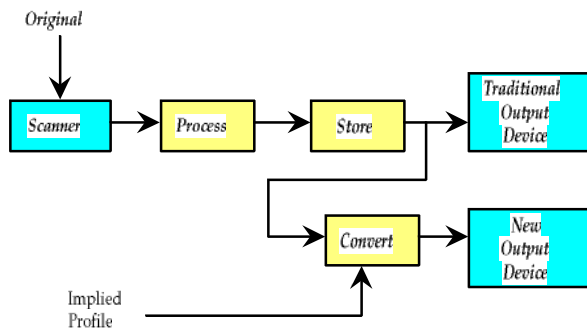
In this approach, image data is stored in a form independent of the device from whence it came, and the (usually unknown) device upon which it will be rendered: such an image encoding is called a *device-independent* representation. In the most-common implementations of this approach, the descriptors of the colors of the image, as perceived by a color-normal human observer, not the colorants of some particular output device, form the basis for the encoding. The conversions between device-dependent and device-independent color representations, since it must be performed on a pixel-by-pixel basis, requires considerable processing. In return for this overhead, we obtain flexibility: the scanner software need know nothing about the end use of the image, images from more than one scanner can easily be incorporated into a single document, and images printed on different comparable printers at different times and locations will appear alike. If the conversions to and from the device-dependent representations can correct for the effects of disparate viewing conditions, an image may be proofed on a monitor and sent to an offset press with an acceptable visual match.

In order to allow systems with modest processing capabilities to perform acceptably rapidly, conversions between device-dependent and device-independent color representations are usually performed approximately. In order to reduce the number of potentially-damaging approximations to a minimum, some system designers have proposed the following modified model:



In this approach, the data is stored in the native form of the originating device. In order to allow the flexible conversion of this device-dependent data into the native color space of arbitrary output devices, sufficient information about the originating device to allow the conversion of the device colorants to colors must be made available. This is the same information the conversion module between the scanner and the device-independent representation in the previous figure uses, but the conversion itself is postponed. Now, whenever the information needs to be converted to the colorants of a particular output device, a new approximation is constructed, one that embodies both the effects of the conversion of the scanner space to a device-independent representation and the conversion of that representation to the native space of the output device. The pixels in the image are run through that approximation, which should damage the image less than the two conversions performed in the previous system model. If slavishly followed, this model requires that processing of the image after initial capture be performed in the native space of the input device. Such a rigorous approach may not be practical: the native color space of the input device may be unknown to the image processing program; even if known, at the precision chosen, it may not allow desired changes to be made without introducing artifacts; and it may not have the right characteristics for some kinds of processing algorithms, such as gamut mapping, which usually requires a luminance-chrominance color space for good results.

Since for some time to come, systems designed along the lines of the first model will coexist and communicate with systems employing the second or third model, an extension of the first model is worth considering:



This approach brings the methods of the third approach to allow greater flexibility for the first. Knowledge of the characteristics of the traditional output device and its associated viewing conditions allows us to construct an approximation module to convert the image information to a form appropriate to any new output device and viewing conditions. A similar module can be constructed to convert the device-dependent data of the closed system to a convenient device-independent representation.

Television is an example of a system organized along the lines of the first model: processing in or after the camera transforms the output of the light sensors into the color space of a standard output device, in this case a self-luminous display viewed in a dim environment. In order to take an image so encoded and produce a reproduction on a different output device such as a printer, we need to know the nature of the canonical display.

Device-Independent Color Spaces

No matter which of the three models she uses, the designer of a flexible system must know how to represent colors in a device-independent fashion: one system model explicitly requires this, and such knowledge is necessary to construct the conversion modules of the other systems.

Storing spectral data for each pixel is impractical; even if we could conveniently measure the spectral data, storing and transmitting so much information would produce an unwieldy system. Most device-independent color representations (note Wandell) reduce the amount of data required by describing the colors in the image only in terms of their effect upon a color-normal observer. The key to this reduction is the color matching experiment described in the previous chapter. We have seen that for each color to be matched, the color matching experiment yields three numbers, corresponding to the quantity of light from each of the three projectors; this set of three numbers is called the tristimulus value of the matched color. Infinitely many spectral distributions can be described by the same tristimulus values, and hence will look the same; instances of such spectra are called *metamers*.

The conversion of spectral to color information is performed using a standard observer derived from a color-matching experiment like that described in the previous chapter. This standard observer is not a person, but merely a set of tristimulus values which match a series of spectral sources. The most commonly used standard observer, adopted in 1931

by the CIE, is the result of averaging the results obtained from 17 color-normal observers with the color matching carried out a 2-degree visual field. These curves are sufficient to compute the tristimulus values of an arbitrary spectrum. For each wavelength in the spectrum, the standard observer gives the tristimulus values required to match that wavelength. Because of additivity, the tristimulus value matching the entire spectrum is the sum of the tristimulus values matching the energy at each wavelength.

We can consider the three numbers comprising a tristimulus value as a column vector:

$$\begin{bmatrix} U \\ V \\ W \end{bmatrix}$$

We can produce an alternate representation of this column vector by multiplying it by a matrix as follows:


$$\begin{bmatrix} Q \\ R \\ S \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \begin{bmatrix} U \\ V \\ W \end{bmatrix}$$

If the matrix is nonsingular, we can retrieve our original representation by multiplying the new tristimulus value by the inverse matrix. The tristimulus values produced by the color matching experiment are linear, so multiplication by a constant matrix can transform the tristimulus values corresponding to any set of projector filters to the tristimulus values of any other set of such filters. Taking the values of the original color matching curves at each wavelength as a tristimulus value, we can construct the color matching curves for the new color space by multiplying each tristimulus value by the matrix. We call all such linear transforms of color spaces defined in the color matching experiment *RGB color spaces*.

One particular RGB color space, CIE 1931 XYZ (referred to in the rest of this chapter simply as XYZ), has become the most common reference for color matching. XYZ is derived from CIE 1931 spectral RGB (define or refer to earlier description) as follows:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.49 & 0.31 & 0.20 \\ 0.17697 & 0.81240 & 0.01063 \\ 0.0 & 0.01 & 0.99 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

The color matching curves for XYZ are given below:



XYZ has several interesting properties. The color matching curves are everywhere positive, and the Y dimension is proportional to the luminous-efficiency function. However, the XYZ primaries are not physically realizable.

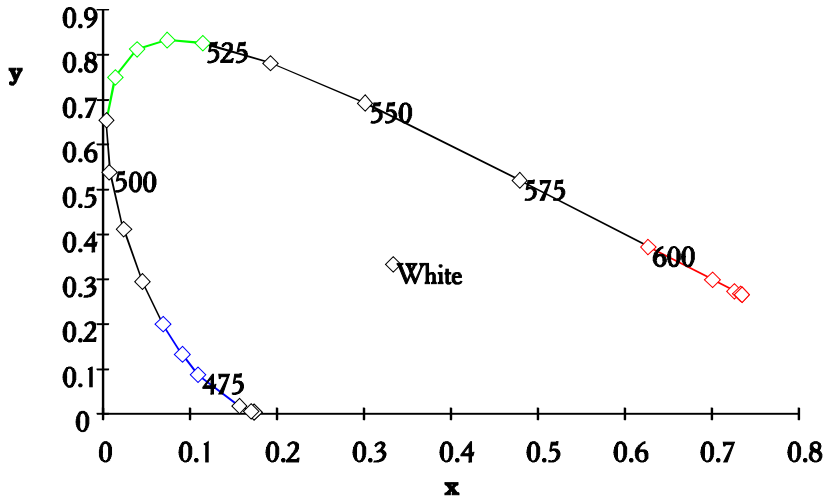
Three dimensional color representations are much more tractable than spectra, but even three-dimensional information is difficult to represent clearly on paper or white boards. One popular simplification is to remove the information pertaining to the luminosity of the color by normalizing each of the tristimulus values to their sum; this operation produces a measure of the color of a stimulus without regard to its intensity or its *chromaticity*. When XYZ is subjected to this operation, three normalized values, represented by lower case characters, are formed as follows:

$$x = \frac{X}{X + Y + Z}$$


$$y = \frac{Y}{X + Y + Z}$$

$$z = \frac{Z}{X + Y + Z}$$

Since the three values add to unity, one of them is superfluous and may be discarded. Following the tradition of discarding the z, we are left with x and y, which together define *xy chromaticity space*. Figure xx shows the visible spectrum plotted in xy chromaticity space.




This linear chromaticity representation has some useful properties. If we connect the two extremes of the visible spectrum by a straight line, we form a horseshoe-shaped figure with a one-to-one mapping to the visible colors: all the visible colors lie within the figure, and all colors that lie within the horseshoe are visible. The addition of two colors to form a third obeys what is called the *center-of-gravity rule*: if A_1 of the first color is added to A_2 of the second, the chromaticity of the new color lies on a straight line between the two original colors, $A_2 / (A_1 + A_2)$ of the way from the first color to the second color, as shown below for $A_1 = 2A_2$:



When working with chromaticity diagrams, one shouldn't lose track of the fact that a great deal of information has been lost in achieving the convenience of a two-dimensional representation. A blazing bright red and a dull fire-brick color, a royal blue and an inky black, or a brilliant yellow and a somber brown can have identical chromaticities.

A problem with xy chromaticity space is that equal steps at various places on the diagram correspond to different perceptual changes: a large numerical change in the chromaticity

of a green color may be barely noticeable, while a small change in that of a blue could dramatically change the perceived color. In 1942, David MacAdam performed a study in which he measured the amount of change in color that produced a just-noticeable difference in a set of observers. He presented his results in the form of a set of ellipsoids in XYZ. Shortly afterward, Walter Stiles predicted the shape of a set of ellipsoids based on other testing. The two sets of ellipsoids are similar, but not identical. If Stiles' ellipsoids are enlarged by a factor of ten and converted to xy chromaticities, they become ellipses. Plotting the major and minor axes of these ellipses results in the following diagram:




A color representation in which equal increments corresponded to equal perceptual differences would be called a *perceptually-uniform* representation. If such a representation possessed a chromaticity diagram, Stiles' ellipsoids would plot a circles of constant radius. Unfortunately, such representations do not exist, so the term perceptually uniform is extended to encodings that are close to the desired property. In 1976, the CIE standardized a modification of xy chromaticity space called u'v' chromaticity space, with the following definition:

$$u' = \frac{4x}{-2x + 12y + 3}$$

$$v' = \frac{9y}{-2x + 12y + 3}$$

Plotting Stiles' ellipsoids (magnified by a factor of ten, as before) in on the u'v' chromaticity diagram yields the following:



The ellipses are more nearly the same size, and the major and minor axes are closer to the same length; the worst-case departure from uniformity is about 4:1.

Instead of the Cartesian chromaticity diagrams shown above, a polar representation is sometimes used. If the origin is set to a convenient white, the phase corresponds to the hue of the color under discussion, and the positive distance to its chroma.

Perceived brightness is nonlinearly related to luminous intensity. The exact nature of the response varies according to the nature of the surround and the absolute luminance, but people can always distinguish more dark levels and fewer light levels than a linear relationship would predict. For a light surround, the cube root of luminous intensity is a good approximation to perceived brightness.

Now that we have derived an approximation to a perceptually-uniform chromaticity space, and have a way to convert luminance to a perceptually-uniform analog, we are almost ready to define a three-dimensional perceptually-uniform color space. We need to take into account two additional pieces of information. The first is that people judge both chromaticity and lightness not in absolute terms, but by comparison with a mentally-constructed color (which may or may not appear in the scene) which they refer to as “white”. The value of this color is part of what is called the state of adaptation of the viewer, and will in general vary depending on where in the image the viewers attention is placed. The second is that, as the luminance decreases, the subjective chroma of a color decreases.

In 1976, the CIE standardized a color space called CIE 1976 ($L^*u^*v^*$), less formally known as CIELUV, which takes the above psychophysics into account to construct an

approximation to a perceptually-uniform three-dimensional color space. To calculate the coordinates of a color in CIELUV, we begin with the XYZ coordinates of the color, $\begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$,

and those of the white point, $\begin{bmatrix} X_n \\ Y_n \\ Z_n \end{bmatrix}$. One axis of CIELUV is called the CIE 1976 Lightness,

or L^* , and it is defined using a cube root function with a straight line segment near the origin:

$$L^* = 116 \left(\frac{Y}{Y_n} \right)^{1/3} - 16, \quad 1 \geq \frac{Y}{Y_n} > 0.008856$$

$$L^* = 903.3 \left(\frac{Y}{Y_n} \right), \quad 0 \leq \frac{Y}{Y_n} \leq 0.008856$$

CIELUV uses a Cartesian coordinate system. A color's position along the L^* axis contains only information derived from its luminance in relation to that of the reference white. The other two axes of CIELUV are derived from the chromaticity of the color and that of the reference white:

$$u^* = 13L^*(u' - u_n')$$

$$v^* = 13L^*(v' - v_n')$$

Multiplying the difference between the chromaticities of the color and the reference white by value proportional to L^* mimics the psychological effect that darker colors appear less chromatic.

The Cartesian distance between two similar colors is a useful measure of their perceptual difference. CIE 1976 ($L^*u^*v^*$) color difference, ΔE , is defined as:

$$\Delta E = \sqrt{\Delta L^2 + \Delta u^2 + \Delta v^2}$$

where ΔL is the difference in the L^* values of the two colors, Δu the difference in their u^* values and Δv the difference in their v^* values. The constants of proportionality in the CIELUV definition were chosen so that one just-noticeable-difference is approximately one ΔE . As the colors get farther apart, this measure becomes less reliable; it is most useful for colors closer together than 10 ΔE or so.

The cylindrical representation of CIELUV is also useful. The L^* axis remains the same, with the phase corresponding to hue and the radius associated with chroma.

chroma: $C^* = (u^{*2} + v^{*2})^{1/2}$

hue angle: $huv = \tan^{-1} \left(\frac{v^*}{u^*} \right)$

In 1976, the CIE also standardized another putative perceptually-uniform color space with characteristics similar to CIELUV. CIE 1976 ($L^*a^*b^*$), or CIELAB, is defined as follows:

$$L^* = 116f\left(\frac{Y}{Y_n}\right) - 16$$

$$a^* = 500\left(f\left(\frac{X}{X_n}\right) - f\left(\frac{Y}{Y_n}\right)\right)$$

$$b^* = 200\left(f\left(\frac{Y}{Y_n}\right) - f\left(\frac{Z}{Z_n}\right)\right)$$

where

$$f(x) = 116(x)^{1/3} - 16, \quad 1 \geq x > 0.008856$$

$$f(x) = 903.3x, \quad 0 \leq x \leq 0.008856$$

chroma: $C^* = (a^{*2} + b^{*2})^{1/2}$

hue: $hab = \tan^{-1}\left(\frac{b^*}{a^*}\right)$

The two color spaces have much in common. L^* is the same, whether measured in CIELAB or CIELUV. In each space the worst-case departure from perceptual uniformity is about 6:1. CIELUV is widely used in situations involving additive color, such as television and CRT displays (the chromaticity diagram is particularly convenient when working with additive color), while CIELAB is popular in the colorant industries, such as printing, paint, dyes, pigments, and the like. Both spaces have fervent proponents, and it appears unlikely that one will supersede the other in the near future, in spite of their similarities.

Desirable characteristics for Device-Independent Color Spaces

A device-independent color space should see colors the way that color-normal people do; colors that match for such people should map to similar positions in the color space, and colors that don't appear to match should be farther apart. This implies the existence of exact transforms to and from internationally-recognized colorimetric representations, such as CIE 1931 XYZ. Defining transforms between a color space and XYZ implicitly defines transforms to all other spaces having such transforms. A further implication is that a device-independent color space should allow representation of most, if not all, visible colors.

A device-independent color space should allow compact, accurate representation. In order to minimize storage and transmission costs and improve performance, colors should be represented in the minimum number of bits, given the desired accuracy. Inaccuracies will be introduced by quantizing, and may be aggravated by manipulations of quantized data. In order to further provide a compact representation, any space should produce compact results when subjected to common image-compression techniques. This criterion favors perceptually-uniform color spaces; non uniform spaces will waste precision quantizing the parts of the space where colors are farther apart than they should be, and may not resolve perceptually-important differences in the portions of the color space where colors are closer together than a uniform representation would place them.

Most image compression algorithms are themselves monochromatic, even though they are used on color images. JPEG, for example, performs compression of color images by compressing each color plane independently. The lossy discrete cosine transform compression performed by the JPEG algorithm works by discarding information rendered invisible by its spatial frequency content. As we saw in the preceding chapter, human luminance response extends to higher spatial frequency than chrominance response. If an image contains high spatial frequency information, only the luminance component of that image must be stored and transmitted at high resolution; some chrominance information can be discarded with little or no visual effect. Effective lossy image compression algorithms such as DCT can take advantage of the difference in visual spatial resolution for luminance and chrominance, but, since they themselves are monochromatic, they can only do so if the image color space separates the two components. Thus, a color space used with lossy compression should have a luminance component.

The existence of a separate luminance channel is necessary, but not sufficient. There also should be little luminance information in the putative chrominance channels, where its presence will cause several problems. If the threshold matrices for the chrominance channels are constructed with the knowledge that those channels are contaminated with luminance information, the compressed chrominance channels will contain more high-frequency information than would the compressed version of uncontaminated chrominance channels. Hence, a compressed image with luminance-contaminated chrominance channels will require greater storage for the same quality than an uncontaminated image. If the threshold matrices for the chrominance channels are constructed assuming that the channels are uncontaminated, visible luminance information in these channels will be discarded during compression. Normal reconstruction algorithms will produce luminance errors in the reconstructed image because the missing luminance information in the chrominance components will affect the overall luminance of each reconstructed pixel. Sophisticated reconstruction algorithms that ignore the luminance information in the chrominance channels and make the luminance of each pixel purely a function of the information in the luminance channel will correctly reconstruct the luminance information, but will be more computationally complex.

A device-independent color space should minimize computations for translations between the interchange color space and the native spaces of common devices. It is unlikely that the interchange color space will be the native space of many devices. Most devices will have to perform some conversion from their native spaces into the interchange space. System cost will be minimized if these computations are easily implemented.

Linear RGB Color Spaces Revisited

We have already encountered linear RGB color spaces, which are linear transforms of color spaces defined in the color matching experiment. It is worthwhile spending some time to understand some of the details of such color spaces, not because they are themselves widely used in computer graphics and imaging systems, but because they

form the basis for the more complex color representations used in such systems. CIE 1931 XYZ itself is the most common *reference color space*, the color space in which most others are defined.

We shall consider that the range of an RGB color space is the interval [0,1]. When working with RGB values with other ranges, they can be scaled appropriately. The white point of an RGB color space is defined to be that color emitted when all three RGB values are set to 1. The light sources in an additive color system are called its *primaries*. Most commonly, linear RGB color spaces other than XYZ are defined in terms of the xy chromaticities of their primaries and those of their white point. However, to perform conversions among color the matrix that transforms CIE 1931 XYZ to them.

Given the chromaticities of each of the primaries, $x_R, y_R, x_G, y_G, x_B, y_B$, and of the white point, x_W, y_W , we first calculate the chromaticities that were discarded as redundant in the definition of xy chromaticity:

$$z_R = 1 - (x_R + y_R)$$

$$z_G = 1 - (x_G + y_G)$$

$$z_B = 1 - (x_B + y_B)$$

$$z_W = 1 - (x_W + y_W)$$

Next we compute a set of weighting coefficients implied by the white point:

$$\begin{bmatrix} a_1 \\ a_2 \\ a_3 \end{bmatrix} = \begin{bmatrix} x_R & x_G & x_B \\ y_R & y_G & y_B \\ z_R & z_G & z_B \end{bmatrix}^{-1} \begin{bmatrix} x_W / y_W \\ 1 \\ z_W / y_W \end{bmatrix}$$

Then we compute the RGB to XYZ conversion matrix:

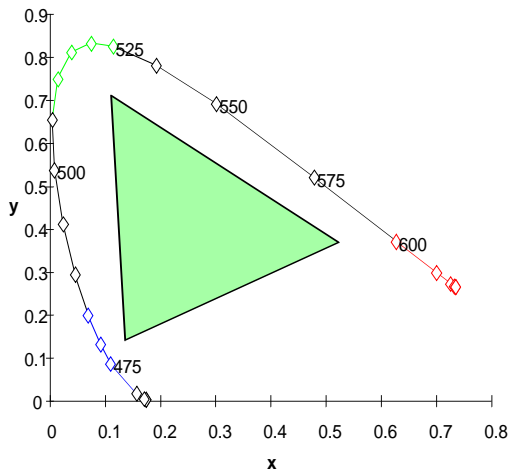
$$\mathbf{M} = \begin{bmatrix} x_R & x_G & x_B \\ y_R & y_G & y_B \\ z_R & z_G & z_B \end{bmatrix} \begin{bmatrix} a_1 & 0 & 0 \\ 0 & a_2 & 0 \\ 0 & 0 & a_3 \end{bmatrix}$$

and finally,

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \mathbf{M}^{-1} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

If we plot the xy chromaticities of the primaries of an RGB color space, we obtain a diagram like the one below. Because of the center-of-gravity rule, mixing any two of the primaries in various proportions allows us to construct any color on the line between them. Adding varied amounts of the third primary allows the construction of any color between an arbitrary point on the line and the chromaticity of the third primary.

Therefore, the chromaticity gamut of an RGB color space is the interior of the triangle formed by the primaries.



Knowing how to find the gamut of an RGB color space, we can now simply explain why, for some colors, we must add light to the sample side of the screen to achieve a match in the color-matching experiment. The chromaticity representation of the gamut of visible colors is convex. In order for a primary to be visible, it must lie within the gamut of visible colors. Any triangle constructed within the convex gamut will not contain some colors in the gamut. The omitted colors are the ones that cannot be matched with positive amounts of the primaries.

Nonlinear RGB Color Spaces

The primaries of a color cathode ray tube (CRT) are the visible emissions of three different mixtures of phosphors, each of which, can be independently excited by an electron beam. When an observer views the CRT from a proper distance, the individual phosphor dots cannot be resolved, and the contributions of each of the three phosphors are added together to create a combined spectrum, which the viewer perceives as a single color. To a first approximation, the intensity of light emitted from each phosphor is proportional to the electron beam current raised to a power:

$$L_e \propto i^\gamma$$

Thus a CRT with constant gamma for all three phosphors, viewed in a dark room, produces a color which can be described, in the color space of its primaries, as

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} \propto \begin{bmatrix} (I_R)^\gamma \\ (I_G)^\gamma \\ (I_B)^\gamma \end{bmatrix}$$

The beam current of computer monitor usually bears a linear relationship to the values stored in the display buffer. With the RGB values scaled into the range [0,1], if we wish to

produce a color with tristimulus values in the color space of the monitor's primaries of

$\begin{bmatrix} R \\ G \\ B \end{bmatrix}$, then the beam current should be $\begin{bmatrix} R^{1/\gamma} \\ G^{1/\gamma} \\ B^{1/\gamma} \end{bmatrix}$. This nonlinear RGB encoding is often

called *gamma-corrected RGB*. In general, nonlinear RGB encodings are derived by starting with linear XYZ, multiplying by a three-by-three matrix, and applying a one-dimensional nonlinearity to each component of the result.

Most CRTs have gammas of a little over 2, which means that they significantly compress dark colors and expand light ones. An image which has been gamma-corrected for a typical CRT thus compresses the light colors and expands the dark ones, like the eye's response, which can be reasonably modeled with a power law of three for a light surround (the requisite power is about 3.75 for a dim surround and 4.5 for a dark surround). Thus the gamma correction usually performed for television and computer displays brings the representation closer to perceptual uniformity. The jargon of gamma correction has a confusing quirk: an image that has been gamma corrected for a monitor with a gamma of 2.2 is referred to as itself having a gamma of 2.2, even though the gamma correction applied to a linear representation to correct it amounts to raising each component of each pixel in the image to the power 0.45 or 1/2.2.

Gamma-corrected RGB color spaces are often reasonable choices for storing color images: they can be made device-independent by specifying their primaries and white point, they offer simplified decoding to CRT monitors, the nonlinearities used usually allow acceptable accuracy with eight bits precision per primary. However, if limited to positive primaries they have limited gamut, and they are not as easily compressible as other representations.

Derivatives of Nonlinear RGB Color Spaces

As seen above, RGB color spaces have advantages. Perhaps their greatest disadvantage is that they have no luminance axis, and thus cannot directly benefit from bandwidth compression schemes that send chromaticity information at lower effective resolution than luminance information. Throughout the world, television signals are encoded in color spaces derived from nonlinear RGB spaces in a manner that approximates a luminance-chrominance space. YUV and YCrCb

France and the former Soviet Union employ a television standard called *Sequential Couleur avec Mémoire* (SECAM). Britain, Germany, and many other European countries use a system termed Phase Alternation Line (PAL). Both these systems use the YUV encoding, which is defined as follows:

$$\begin{bmatrix} Y \\ U \\ V \end{bmatrix} = \begin{bmatrix} 0.30 & 0.59 & 0.11 \\ 0.70 & -0.59 & -0.011 \\ -0.30 & -0.59 & 0.89 \end{bmatrix} \begin{bmatrix} R' \\ G' \\ B' \end{bmatrix}$$

(double check this. Where is YUV specified?)

Recently, CCIR 601-2 has specified an encoding for studio digital television systems, to be employed in situations which make it applicable to all three of the above standards. This encoding, termed YCrCb, is defined as follows:

$$\begin{bmatrix} Y \\ Cr \\ Cb \end{bmatrix} = \begin{bmatrix} 0.299 & 0.587 & 0.114 \\ 0.5 & -0.419 & -0.081 \\ -0.169 & -0.331 & 0.5 \end{bmatrix} \begin{bmatrix} R' \\ G' \\ B' \end{bmatrix} + \begin{bmatrix} 0.0 \\ 0.5 \\ 0.5 \end{bmatrix}$$

YIQ

In the United States, Canada, Japan, and Mexico, the National Television System Committee (NTSC) standard is used. NTSC specifies an encoding called YIQ, defined as follows:

$$\begin{bmatrix} Y \\ I \\ Q \end{bmatrix} = \begin{bmatrix} 0.30 & 0.59 & 0.11 \\ 0.60 & -0.28 & -0.032 \\ 0.21 & -0.52 & 0.31 \end{bmatrix} \begin{bmatrix} R' \\ G' \\ B' \end{bmatrix}$$

YIQ's two chrominance axes are better aligned with red/green and blue/yellow than is the case for YUV and YCrCb. This allows taking advantage of chromatic differences in visual spatial frequency response by encoding the blue/yellow channel at lower spatial resolution than the red/green channel. NTSC television transmission takes advantage of the chrominance axis alignment by encoding the Q (blue/yellow) signal at one-third (check this) the bandwidth of the I (red/green) signal.

In the equations for YIQ, YUV, and YCrCb, the primes are added to the RGB designators to emphasize the fact that these matrix operations are performed on the gamma-corrected (nonlinear) RGB signals. Because of this fact, the Y component in each of these color spaces bears the desired relationship to luminance only along the neutral axis; as chroma increases, the Y signal contains values lower than the actual luminance. This does not mean that the reproduced colors are in error, but only that luminance information is being carried in the putative chrominance components.

Although treated as such in some of the literature, *YIQ, YUV, and YCrCb, are not themselves gamma-corrected RGB color spaces*. They are derived from linear XYZ as follows: first multiply by a three-by-three matrix, apply a one-dimensional nonlinearity to each component, and multiply the result by a three-by-three matrix. Thus there are three or four components to the definition of each of these color spaces: the linear RGB color space from which they are derived, the nonlinearity, the matrix defining how the nonlinear components are combined, and, in possibly, the column vector to be added at the end of the calculation. This set of transformations is not specialized to television use, but is quite general; for example, CIELAB can be derived from XYZ using this sequence of operations.

Kodak PhotoYCC is a color space of great commercial interest. It is defined as YCrCb, using the CCIR 709 primaries, a white point of D₆₅ (x = 0.3127, y = 0.3290), and the CCIR 601 nonlinearity, which is a power law with a gamma of 2.2 over most of its range.

PhotoYCC is unusual in that it allows for negative values of the primaries, thus increasing the gamut over what would otherwise be possible. For each component, plus and minus full scale values are specified.

HSV

HSV, proposed by Alvy Ray Smith in 1978 (ref Smith), is another luminance-chrominance color space; it defines colors in terms of a *hexcone*, a roughly conical color space with the tip at the origin, a luminance axis up the middle with hue angles arranged around it, and chroma increasing away from the luminance axis. HSL is named for its axes: hue, saturation, and value (a synonym for lightness). The geometry is similar to CIELAB and CIELUV, but the RGB/HSV and HSV/RGB computations are simpler than the equivalent calculations in the perceptually-uniform CIE spaces. HSV was not originally intended as a device-independent color space; it is defined in terms of transforms from some unspecified RGB, but HSV can obtain device-independent status if the RGB upon which it is based is colorimetrically defined. HSV was originally defined in terms of linear RGB, but in the usual practice, the basis for the RGB to HSV conversion is whatever RGB happens to be around, and that is most often gamma-corrected RGB.

HSV is defined algorithmically:

```
v := max(R,G,B);
Let X := min(R,G,B);
S := (V-X)/V; if S=0 return;
Let r := (V-R)/(V-X);
Let g := (V-G)/(V-X);
Let b := (V-B)/(V-X);
If R=V then H := (if G=X then 5+b else 1-g);
If G=V then H := (if B=X then 1+r else 3-b);
           else H := (if R=X then 3+g else 5-r);
H := H/6;
```

The reverse algorithm is also defined.

HSL

HSL (ref Graphics Standards Planning Committee) has similar objectives to HSV, but employs a different geometry: a *double hexcone*. This arrangement takes the cone balanced on its point from HSV and adds another similar, but inverted, cone on top of it. Thus the lightest and the darkest color in HSL are achromatic, with maximum saturation obtained at a lightness of one-half.

HSL is also defined algorithmically:

```
M := max(R,G,B); m := min(R,G,B);
If M=m go to Step 1;
r := (M-R)/(M-m); g := (M-G)/(M-m); b := (M-B)/(M-m);
L := (M+m)/2;
If M=m then S := 0;
           else if L <= 0.5 then S := (M-m)/(M+m);
```

```

        else S := (M-m)/(2-M-m);
    If S=0 then h := 0;
        else if R=M then h := 2+b-g;
        else if G=M then h := 4+r-b;
        else h := 6+g-r;
    H := h*60 mod 360;

```

The reverse algorithm is somewhat slower, but a modification exists that is about as fast (ref. Fishkin).

HSL, HSV, and similar easy-to-calculate luminance-chrominance color spaces enjoyed great popularity during the 1980s, but are not commonly used for device-independent color, which usually employs approximation instead of direct calculation, rendering unimportant their chief advantage.

Calibrating CRTs

Primaries in practice.

Unfortunately for the designer of a device-independent color system, CRT phosphors found in commercial monitors vary. In 1953 the National Television System Committee (NTSC) specified a set of phosphor chromaticities for television use. Over the years, receiver manufacturers specified different phosphors, sacrificing saturation in the greens in order to achieve brighter displays. Two standard sets of phosphor chromaticities are now in use: the European Broadcasting Union (EBU), and the Society of Motion Picture and Television Engineers (SMPTE) standards. However, many receiver and monitor manufacturers employ phosphors which meet neither standard.

A similar situation exists with respect to white point. The NTSC initially specified a white point of CIE (?) illuminant C ($x = 0.3101$, $y = 0.3162$), but illuminant D₆₅ ($x = 0.3127$, $y = 0.3290$) is prevalent in studio television equipment today. The difference between these two values is not great, but television receivers for home use and computer monitors have generally used white points much bluer than either of these values in order to obtain brighter displays.

Gammas in practice.

When displaying images on uncalibrated CRTs, by far the greatest source of objectionable error is the variation in the nonlinear response of the various displays. Users can accommodate to a moderate change in white point and the color shifts caused by the usually-encountered differences in phosphors. However, the lack of standardization of CRT gammas coupled with user sensitivity to relative luminance errors often create unacceptable results. An image displayed on a monitor with higher gamma than intended will suffer from overly-dark midtones, while a CRT with lower than intended gamma will display a washed-out image.

The simple expression, $L_e \propto i^\gamma$, is an adequate model for the relationship of beam current to CRT luminous output for many purposes, but accurate model-based calibration of a

CRT requires a slightly more complex function. Motta and Berns have shown that the relationship

$$L_e = (K_1 i + K_2)^g, \quad (K_1 + K_2) = 1$$

can predict colors to within 0.5 CIELAB ΔE over the complete CRT color space. K_1 , termed the gain factor, is greater than one, making K_2 , the offset, negative. This relationship accurately models the CRT's typical lack of output until the input reaches a certain level. In general, K_1 , K_2 , and g will be different for each primary. The coefficients for this model may be derived by measuring the CRT output energy in response to a series of known inputs, but Motta has proposed an alternative method that requires no instrumentation. In Motta's visual calibration approach, the user adjusts a slider until he first sees a noticeable change from a dark background, then matches a constant dithered pattern with a variable non-dithered pattern. These six (two for each primary) measurements provide enough information to characterize the CRT.

White Points in Practice

Calibrating Printers

The interaction of light with the dyes and pigments of practical printers to form colors is more complex than the color-forming mechanisms of CRTs, making it more difficult to construct an accurate mathematical model for a printer. A theoretical model can be used alone (ref. Neugabauer) or modified to reflect testing (ref. Viggano, ref Yule) or an empirical transfer function can be derived from print samples. (ref Nin, Kasson, & Plouffe)

Calibrating Scanners and Cameras

Both cameras and scanners convert spectral data into tristimulus values. If the spectral response curves of these devices were linear transforms of the human color matching functions, then calibration would be fairly simple. Many electronic cameras, especially those based on television technology, attempt to relate their spectral responses to color matching functions; the television industry has a long association with a colorimetric approach to image processing. Unfortunately, with one or two exceptions (ref Yorktown Scanner Paper), scanner spectral responses bear no easily-decoded relationship to human color matching functions. Thus, most scanners suffer from metamerism: in general some colors encoded as identical will look different to people, and some metamers (colors with different spectral compositions that appear identical) will be encoded as different colors. However, in the case of devices that scan photographic materials, the universe of possible spectra which must be converted to color is constrained. Consider a color transparency; at each point on the film, the spectrum is the wavelength-by-wavelength product of the spectrum of the illuminant and the transmission spectra of each of the three (cyan, magenta, and yellow) dye layers. Scanner calibration methods can take advantage of the constrained spectra of photographic materials to produce accurate results, especially if

calibrated for the particular dye spectra in each type of scanned material (ref Stockham paper).

Efficient Conversions to and from Device-Independent Color Spaces

Mechanisms for converting data from one color space to another may be constructed using either model-based or measurement-based techniques. The mathematics of conversion between well-defined device-independent color spaces straightforwardly defines a model which can be used to convert arbitrary colors from one space to the other. As discussed above, the mechanisms that govern the conversion of cathode ray tube (CRT) beam current into visible colors are reasonably easy to model, so the conversion to monitor space can be described mathematically. Some have found mathematical characterization inadequate, and have used measurement and interpolation to produce more accurate results. (ref Post and Calhoun)

In most device-independent image processing systems, conversion from device-independent to device-dependent form dominates the conversion cost budget. For all but the simplest models, performing this conversion by directly implementing the underlying mathematics is not computationally attractive, and approximate methods more suitable. There are two classes of approximation commonly used; those employing simplified models, and those relying on some kind of interpolation. An understanding of three-dimensional interpolation techniques properly begins with trilinear interpolation, although other interpolation methods are capable of similar accuracies with less computation. The form of trilinear interpolation described here produces a continuous output from a continuous input, which is an advantage where the tables are populated so coarsely that errors can exceed the just-noticeable difference.

In trilinear interpolation, the sample function values are arranged into a three dimensional table indexed by the independent variables [33], e.g., x , y , and z . The range of each input variable is evenly sampled. For example, let x be in the range $x_0..x_a$, y be in the range $y_0..y_b$, and z be in the range $z_0..z_c$. One possible sampling would be (x_i, y_j, z_k) where

$$0 \leq i \leq a$$

$$0 \leq j \leq b$$

$$0 \leq k \leq c$$

and

$$x_i = x_0 + i \times \frac{x_a - x_0}{a}$$

$$y_j = y_0 + j \times \frac{y_b - y_0}{b}$$

$$z_k = z_0 + k \times \frac{z_c - z_0}{c}$$

Then the function, F , is approximated for the target point (r,s,t) as follows:

$$F(r, s, t) = (1.0 - d_r) \left\{ \begin{aligned} & \left[(1.0 - d_s) \left[(1.0 - d_t) F(x_i, y_j, z_k) + d_t F(x_i, y_j, z_{k+1}) \right] \right] \\ & + d_s \left[\left[(1.0 - d_t) F(x_i, y_{j+1}, z_k) + d_t F(x_i, y_{j+1}, z_{k+1}) \right] \right] \end{aligned} \right\} \\ + d_r \left\{ \begin{aligned} & \left[(1.0 - d_s) \left[(1.0 - d_t) F(x_{i+1}, y_j, z_k) + d_t F(x_{i+1}, y_j, z_{k+1}) \right] \right] \\ & + d_s \left[\left[(1.0 - d_t) F(x_{i+1}, y_{j+1}, z_k) + d_t F(x_{i+1}, y_{j+1}, z_{k+1}) \right] \right] \end{aligned} \right\}$$

where

$$x_i \leq r < x_{i+1}$$

$$y_j \leq s < y_{j+1}$$


$$z_k \leq t < z_{k+1}$$

$$d_r = \frac{r - x_i}{x_{i+1} - x_i}$$

$$d_s = \frac{s - y_j}{y_{j+1} - y_j}$$

$$d_t = \frac{t - z_k}{z_{k+1} - z_k}$$

The fineness of the sampling and the choice of the color spaces strongly affects the accuracy that trilinear interpolation achieves. The following graph illustrates the error, measured in CIELAB ΔE , of converting from the indicated color spaces to the color space of a display using the CCIR 709 primaries, a white point of D₆₅, and a gamma of 2.2.



So far, we have only addressed the techniques required for colorimetric color reproduction, in which the reproduced image has the same chromaticities as the original, and luminances proportional to those of the original. As we have seen from the preceding

chapter and the beginning of this one, this is not enough to produce images which look the same under various viewing conditions; for this we must go beyond simple colorimetric accuracy to equivalent or corresponding color reproduction.

Gamut Mapping

Our first departure from colorimetric reproduction is not occasioned by a change in the viewing conditions, but by the likelihood that a given device-independent image will contain colors that our output device can't render. In this situation, we must map the colors in the image into the gamut of the output device. This is not an optional step. Gamut mapping *will* take place, whether via an explicit algorithm or by attempting to specify out-of-range values in printer space.

Virtually all successful gamut mapping algorithms strive to minimize apparent shifts in hue angle, instead reducing chroma or changing the luminance of out-of-gamut colors. Most algorithms pass each pixel in the image through processing that does not depend on the value of nearby pixels, although the processing may be affected by global image characteristics. Two popular approaches are *compression* and *clipping*. Clipping algorithms map out-of-gamut values to points on the gamut surface, leaving in-gamut values unaffected. Clipping algorithms have the advantage of retaining all the accuracy and saturation of in-gamut colors, but there is at least a theoretical possibility that adjacent visually distinct out-of-gamut colors will merge, or that smooth color gradients will terminate as they cross the gamut edge, creating visible artifacts. Compression algorithms scale image color values in a possibly nonlinear fashion so as to bring out-of-gamut colors within the device gamut, affecting at least some in-gamut colors. Compression algorithms better preserve the relationship between colors in the image and avoid the two disadvantages of clipping algorithms, but do so at the expense of reducing the saturation or of changing the luminance of in-gamut colors. In a series of preference tests using photographic images, Gentile, et al, found that observers preferred clipping techniques, particularly those which preserved luminance and hue angle, over compression algorithms (ref Gentile).

Correction for Viewing Conditions

Surround

There are many valuable rules of thumb for correcting for viewing conditions. To compensate for the apparent loss of contrast as the surround illumination is decreased, Hunt (RofC p 56-7) suggests increasing the gamma of an image originally viewed with a light surround by 1.25 if it is to be viewed with a dim surround and by 1.5 if it is to be viewed with a dark surround.

White Point

When an observer views a reflection print, he usually accepts as a white point a color near that of the paper illuminated by the ambient light. When he views a projected transparency in a dark room, he accepts the projector illuminant as the white point. When he views a monitor in a dim room, the white point accepted is often quite close to the

white point of the monitor. If the chromaticity of the original scene illuminant is substantially different from the reproduction white point, we need to make a correction. For example, if a viewer observes an image captured in incandescent light (color temperature about 2800 degrees Kelvin) on a CRT with white point D₆₅ in a dark room, the scene will appear much too orange. This occurs because the viewer is adapted to the bluer white point of the CRT instead of the illumination of the original scene, as would have been the case had he been present.

Size

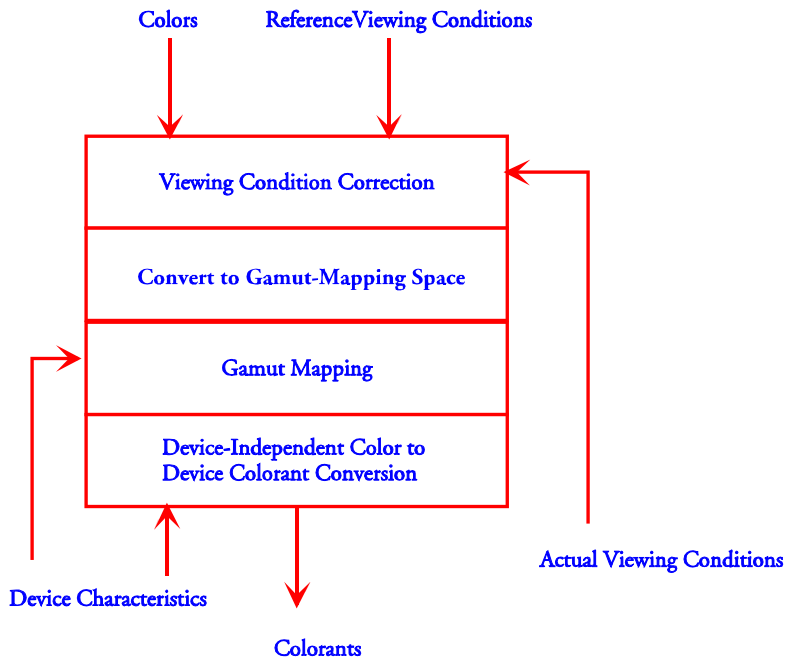
In the previous chapter, we saw that the eye's contrast sensitivity function for blue/yellow and red/green variation peaks at spatial frequencies much lower than does it does for luminance variation. This effect can translate to an apparent decrease in chroma as the size of an image decreases. (Reference Rogowitz Human Vision in Display Design)

Color Appearance Models

Correcting for viewing conditions would probably be more accurate, and would certainly be more satisfying, if, rather than assembling together a collection of *ad hoc* corrections, we possessed a model for human color appearance that operated over a broad range of viewing conditions. Such *color-appearance models* exist, but are largely unproven (ref Hunt, Guth and Nay) . Fairchild and Berns (ref Fairch) have extended CIELAB to correct for chromatic adaptation and surround, calling the result RLAB. Report on comparison of color appearance models.

Putting it all together

Armed with the knowledge of how to perform color calibration, gamut mapping, and correction for viewing conditions, we are now ready to examine the processing contained in the innocent-looking boxes labeled "Convert from Device-Independent" in the system diagrams at the beginning of the chapter. Not all systems are implemented in the same manner, and not all applications will require all three functions, but the diagram below describes a reasonable approach.



The device-independent image is defined both by the colors in it and the reference viewing conditions -- the conditions under which those colors look right. The first thing we do is correct for the viewing conditions under which the output will be observed. The result is an image specified in a device-independent color space. Next we map the out-of-gamut colors of the image into colors within the gamut of the output device, possibly converting to a color space convenient for the process. The result is an image specified in a device-independent color space, which is then converted to the colorants of the output device. As discussed earlier, a practical system will not put each pixel of an image through this kind of processing; rather, the operations defined in the diagram will be used to populate a single three-dimensional lookup table, which will then be used by a three-dimensional interpolation algorithm to process the image.

Adobe Systems, Inc., *PostScript Language Reference Manual*, Addison-Wesley, New York (1990).

American National Standards Institute, *Color Prints, Transparencies, and Photomechanical Reproductions - Viewing Conditions*, ANSI PH2.30-1989 (1989).

Berns, R. S. and Motta, R. J., "Colorimetric Calibration of Soft-Copy Devices to Aid in Hard-Copy Reproduction", *Proc SPSE 41st Annual Conference*, pp. 266-269 (1988).

Berns, R. S., Motta, R. J., and Gorzynski, M. E., "CRT Colorimetry, Part I: Theory and Practice," *Color Research and Application*, vol. xx, no. x, pp xx-xx (199x).

Berns, R. S., Gorzynski, M. E., and Motta, R. J., "CRT Colorimetry, Part II: Metrology," *Color Research and Application*, vol. xx, no. x, pp xx-xx (199x).

Billmeyer, F., and Saltzman, *Principles of Color Technology*, second edition, Wiley, New York (1981).

Bingley, F.J., "Colorimetry in Color Television", *Proceedings of the IRE*, vol. 41, pp. 838851 (July 1953).

Boynton, R.M., *Human Color Vision*, Holt, Rhinehart, and Winston, New York (1979)

Brainard, D.H. and Wandell, B.A., "Calibrated Processing of Image Color", *Color Research and Application*, vol 15, No. 5, pp 266—271 (October, 1990).

CCIR, "Basic Parameter Values for the HDTV Standard for the Studio and for International Programme Exchange," CCIR Recommendation 709, International Radio Consultative Committee (ITU) (April, 1990).

CCIR, "Encoding Parameters of Digital Television for Studios", CCIR Recommendation 601-2, International Radio Consultative Committee (ITU) (1990)

CCIR, Characteristics of Television Systems, Report 624-4, Reports of the CCIR, 1990, Annex to Volume XI - Part 1, page 9, 1990.

CIE, *Colorimetry*, CIE Publication No. 15.2, 2nd Edition, pp 1920, 5658, Commission Internationale de L'Eclairage, Vienna (1986).

Cornsweet, T.N., *Visual Perception*, Academic Press, New York (1970).

Eastman Kodak Company, "Kodak Photo CD Products: A Planning Guide for Developers," Kodak Part Number DCI-200R, Rochester, New York, 1992.

Fairchild, M. D., and Berns, R. S., "Image Color-Appearance Specification Through Extension of CIELAB," *Color Research and Application*, vol. 18, no. 3, pp 178-190 (June 1993).

Fishkin, K., "A Fast HSL-to-RGB Transform," *Graphics Gems*, Glannser, A. S., editor, Academic Press, Boston, pp 448-449 (1990)

Foley, J. D., van Dam, A., Feiner, S. K., and Hughes, J. F., *Computer Graphics, Principles and Practice*, second edition, Addison-Wesley, Reading, Massachusetts, Chapter 13.

Gentile, R. S., Walowit, E., and Allebach, J. P., "A Comparison of Techniques for Color Mismatch Compensation," *Journal of Imaging Technology*, vol. 16, no. 5, pp 176-181 (October 1990).

Graphics Standards Planning Committee, "Status Report of the Graphics Standards Planning Committee," *Computer Graphics*, vol. 13, no. 3, Appendix B, pp III-37-38, (1979)

Guth, S. L., "Unified Model for Human Color Perception and Visual Adaptation II," *Human Vision, Visual Processing and Digital Display IV*, Allebach, J. A., and Rogowitz, B. E., Editors, SPIE/IS&T, vol. 1913, pp 440-448 (1993).

Hunt, R.W.G., *The Reproduction of Colour in Photography, Printing, and Television*, 4th Edition, pp 98-110, Fountain Press, Tolworth, England (1987).

Hunt, R.W.G., *J. Phot. Sci.*, 18, p. 205 (1970) (flesh out)

Hunt, R.W.G., *Measuring Colour*, Ellis Horwood Limited, Chichester, West Sussex, England, 1989, page 70.

Hunt, R.W.G., "Revised Colour Appearance Model for Related and Unrelated Colours," *Color Research and Application*, vol. 16, pp 146-165 (1991).

Hunt, R.W.G., "Model of Color Vision for predicting the Appearance of Colors Under Different Viewing Conditions," *Device-Independent Color Imaging and Imaging Systems Integration*, Motta, R. J., and Berberian, H. A., Editors, SPIE/IS&T, vol. 1909, pp 2-14 (1993).

Kasson, J.M., and Plouffe, W.E., "An Analysis of Selected Computer Interchange Color Spaces", *ACM Transactions on Graphics*, vol 11, no. 4, pp 373-405 (October, 1992).

Loughlin, B.D., *Principles of Color Television*, McIlwain, K. and Dean, C, Eds, pp. 200256, ch 11, Wiley (1956).

Mahy, M., Van Mellaert, B., Van Eycken, L., and Oosterlinck, A., "The Influence of Uniform Color Spaces on Color Image Processing: A Comparative Study of CIELAB, CIELUV, and ATD," *Journal Of Imaging Technology*, Vol. 17, No. 5, pp 232 - 243

Motta, R. J., and Berns, R. S., "The Colorimetric Calibration of a CRT Imaging System for Color Appearance Research," *Proceedings, ISCC/SID Conference, CRT to Hard Copy in Color*, Baltimore, pp 28-31 (1988).

Motta, R. J., "Visual Characterization of Color CRTS," *Device-Independent Color Imaging and Imaging Systems Integration*, Motta, R. J., and Berberian, H. A., Editors, SPIE/IS&T, vol. 1909, pp 224-228 (1993).

Nayatani, Y., Takahama, K., Sobagaki, H., and Hashimoto, K., "Color-Appearance Model and Chromatic Adaptation Transform," *Color Research and Application*, vol. 15, pp 210-221 (1990).

Neugebauer, H. E. J., "Die theoretischen Grundlagen des Mehrfarbenbuchdrucks," *Zeitschrift für wissenschaftliche Photographie, Photophysik, und Photochemie*, Vol. 36 No. 4, pp. 73-89, April, 1937; reprinted in *Neugebauer Memorial Seminar on Color Reproduction*, SPIE Vol. 1184, pp. 194-202 (December 1989).

Nin, S. I., Kasson, J. M., and Plouffe, W. E., "Printing CIELAB images on a CMYK printer using trilinear interpolation," *Color Hard Copy and Graphic Arts*, Bares, J., Editor, Proc SPIE 1670, pp. 164-174, 1992.

NTSC, "NTSC Signal Specifications," Proc. Institute of Radio Engineers, Vol. 42, No. 1, pp. 17-19 (January, 1954).

Pointer, M.R., "The Gamut of Real Surface Colours," *Color Research and Application*, vol. 5 No.3, pp. 145-155 (Fall 1980).

Post, D.L., and Calhoun, C.S., "An Evaluation of Methods for Producing Desired Colors on CRT Monitors," *Color Research and Application*, vol. 14, no. 4, August 1989, pp 172-186.

Rodriguez, M. A., and Stockham, T. G., "Producing Colorimetric Data from Densiometric Scans," *Human Vision, Visual Processing and Digital Display IV*, Allebach, J. A., and Rogowitz, B. E., Editors, Proc SPIE 1913, pp 413-418 (1993).

Smith, A. R., "Color Gamut Transform Pairs," *Computer Graphics, ACM/SIGGRAPH*, pp 12-19. (1978).

SMPTE, "Color Monitor Colorimetry," SMPTE Recommended Practice RP 145-1987, Society of Motion Picture and Television Engineers, White Plains, NY (June, 1987).

SMPTE, "Color Temperature for Color Television Studio Monitors," SMPTE Recommended Practice RP 37-1969, Society of Motion Picture and Television Engineers, White Plains, NY (July 1969)

Society of Motion Picture and Television Engineers, *Signal Parameters - 1125/60 High-Definition Production System*, SMPTE 240M-1988, 1988.

Stone, M, C., Cowan, W. B., and Beatty, J. C., "Color Gamut Mapping and the Printing of Digital Color Images," *ACM Transactions on Graphics*, Vol. 7, No. 4, pp. 249-292 (October 1988)

Von Kries, J.A., *Handbuk der Physiologisches Optik*, Vol. II (W. Nagel, ed), pp. 366-369, Leopold Voss, Hamburg (1911).

Wyszecki, G., and Stiles, W. S., *Color Science, Concepts and Methods, Quantitative Data and Formulae*, second edition, Wiley, New York.

Viggiano, J. A. S., "Modeling the Color of Multi-Colored Halftones", *Proceedings of the Technical Association of the Graphic Arts*, pp. 44-62 (1990).

[35] Xerox Systems Institute, *Color Encoding Standard*, Xerox Corporation, Xerox Systems Institute, 475 Oakmead Parkway, Sunnyvale, CA 94086 (March 1989)

Yule, J. A. C., and Neilsen, W. J., "The Penetration of Light into Paper and its Effect on Halftone Reproduction," *Proceedings of the Technical Association of the Graphic Arts*, pp. 65-76 (1951).