

CS417 Color Science notes

Steve Marschner

May 8, 2003

Color science is the study of how the human visual system (HVS) responds to light of different wavelengths. These notes are to write down the perspective I gave in lecture, which is a little different than what you find in most books, in a bit more detail than is in the slides.

Light detectors

The human eye contains a large number of light detectors. Generally a light detector works by detecting photons of light energy that land on it, and for every photon that comes in there is a probability that it will be detected. This probability depends on the wavelength of the photon: the eye is very likely to detect photons in the green range, less likely to detect photons in the red and blue ranges, and not at all likely to detect photons in the infrared and ultraviolet ranges (or any other wavelengths outside the visible range). The function that tells you the probability of detecting a photon as a function of wavelength is a characteristic of any particular light detector; let's call it $p(\lambda)$ for "probability."

A detector that has detected some photons will produce a scalar signal—a single number—that depends only on the number of photons that were detected. All information about the wavelength is "forgotten," in that there is no way to tell by looking at the signal what wavelength or wavelengths went into it. The number of photons we'll collect at a wavelength λ , then, is just $p(\lambda)$ times the number $n(\lambda)$ of photons of that wavelength that fell on the detector.¹ Adding up the photons for all wavelengths results in an integral:

$$X = \int p(\lambda)n(\lambda) d\lambda$$

Here X just stands for the detector response.

¹Really, I should say "the number of photons per unit wavelength" but I think you get the idea.

Integrals of this form will occur now and then in what follows, so I will use the abbreviation

$$\langle f, g \rangle = \int f(x)g(x)dx.$$

This operation, called the inner product, is a way of taking two functions and producing a single number that is analogous to the dot product, which takes two vectors and produces a single number. In fact, if we approximate p and n by sampling them, say at every nanometer across the range of visible wavelengths, then we can approximate the integral above with a sum that turns out to be exactly a dot product:²

$$\int p(\lambda)n(\lambda) d\lambda \approx \sum_i \tilde{p}[i]\tilde{n}[i] = \tilde{p}\tilde{n}$$

I am using the notation $\tilde{p}[i]$ for the i^{th} sample of the function $p(\lambda)$ and \tilde{p} for the column vector consisting of all the $\tilde{p}[i]$ s. I am also adopting the convention that sensitivity functions will always be row vectors and spectra will always be column vectors.

[Skip this paragraph on first reading and just substitute s for p to r for n in your head.] Before I go on I want to make a minor change of units that will bring these notes into line with the usual practice in color science. The functions n and p are the number of photons and the probability of detecting those photons, but in color science (and in radiometry in general) it's usual to think in terms of the power rather than the number of photons. The photon density $n(\lambda)$ is then replaced by the spectral power density (SPD) $s(\lambda)$, and the detection probability $p(\lambda)$ is replaced by a sensitivity, or signal per unit power, $r(\lambda)$ (r stands for responsivity, since s is already in use). It turns out that $s(\lambda) \propto n(\lambda)/\lambda$ and $r(\lambda) \propto \lambda p(\lambda)$, so everything cancels out and $\langle r, s \rangle = \langle p, n \rangle$. So we needn't have worried about this problem with units—it's just important to be using the right convention when it comes time to look up data and compute things from it.

The eye actually contains three distinct types of light sensors with three different spectral responses. The different types of sensors are generally called S, M, and L for short, medium, and long wavelengths. Let's call these response functions $r_s(\lambda)$, $r_m(\lambda)$, and $r_l(\lambda)$. When it sees a spectrum $s(\lambda)$, your eye then produces

²By using 1 as the sample spacing I am sweeping a global scale factor under the rug. This scale factor is not important in the basic computations of color science so I will not mention it again.

three values that determine the color of the spectrum:

$$\begin{aligned} S &= \langle r_s, s \rangle \approx \tilde{r}_s \tilde{s} \\ M &= \langle r_m, s \rangle \approx \tilde{r}_m \tilde{s} \\ L &= \langle r_l, s \rangle \approx \tilde{r}_l \tilde{s} \end{aligned}$$

If we look just at the samples approximations here, then we have a matrix equation:

$$\begin{bmatrix} S \\ M \\ L \end{bmatrix} = \begin{bmatrix} - & \tilde{r}_s & - \\ - & \tilde{r}_m & - \\ - & \tilde{r}_l & - \end{bmatrix} \begin{bmatrix} | \\ \tilde{s} \\ | \end{bmatrix}$$

or,

$$V = M_{SML} \tilde{s}.$$

(The lines in this equation are just to emphasize the shape of the vectors: the matrix M_{SML} is short and wide; the vector \tilde{s} is very tall.) Here V stands for “visual response.” This equation very succinctly says that what the eye detects is the projection of the spectrum onto a three-dimensional subspace.

Additive color

Consider an RGB computer monitor on which we plan to display some color (R, G, B) . We know that these three numbers control the intensities of the three different colored phosphors in the monitor, and that the colors blend together as long as we sit far enough from the monitor that we can't resolve the individual phosphor dots. Mathematically, this means that the spectrum that is effectively presented to the eye by a monitor displaying the color (R, G, B) is

$$s_a(\lambda) = R s_r(\lambda) + G s_g(\lambda) + B s_b(\lambda).$$

(The a in s_a stands for “additive.”) The functions s_r , s_g , and s_b are the spectra of the three individual phosphors; I am assuming that turning the intensity of one of the colors up and down just scales the spectrum without changing the balance of wavelengths it contains.

If we go to the sampled representation, we can write this equation as a matrix equation, too. Remember that all spectra are treated as column vectors, so \tilde{s}_r , \tilde{s}_g , and \tilde{s}_b are columns of the matrix:

$$\begin{bmatrix} | \\ \tilde{s}_a \\ | \end{bmatrix} = \begin{bmatrix} | & | & | \\ \tilde{s}_r & \tilde{s}_g & \tilde{s}_b \\ | & | & | \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} =$$

or,

$$\tilde{s}_a = M_{RGB} C.$$

Here C stands for “color.” Notice that the matrix M_{RGB} is tall and narrow, meaning that it maps from a small number of dimensions (the 3D space of colors the monitor can display) to a large number of dimensions (the space of all spectra). The matrix M_{SML} is short and wide, meaning that it maps from a large number of dimensions (spectra) to a small number of dimensions (the 3D space of colors detectable by the human eye).

Color reproduction

The problem of color reproduction, which is central to color science, at least as it applies to computer graphics, is this: Given a spectrum s , what color C do I put in the framebuffer to achieve a spectrum s_a that looks the same as s ?

Let me rephrase that question in the terms of the math in the previous sections. When we put the color C in the framebuffer, we will get the spectrum

$$\tilde{s}_a = M_{RGB} C.$$

We want the eye’s response to s and s_a to be the same:

$$M_{SML} \tilde{s} = M_{SML} \tilde{s}_a.$$

Substituting for \tilde{s}_a ,

$$M_{SML} \tilde{s} = M_{SML} M_{RGB} C.$$

The thing to notice about this equation is that the product $M_{SML} M_{RGB}$ is just a 3 by 3 matrix. This means that we can very easily find the right value of C :

$$C = (M_{SML} M_{RGB})^{-1} M_{SML} \tilde{s} \quad (1)$$

There you have it! The solution to the color matching problem, and the key to why color monitors work at all.

Grouping terms in this equation leads to a couple of interesting matrices. If we group it like this:

$$C = [(M_{SML} M_{RGB})^{-1}] [M_{SML} \tilde{s}]$$

$$C = TV$$

then T is a *color transformation matrix* for turning colors that are specified as (S, M, L) triples into equivalent (R, G, B) triples suitable for displaying on our monitor. If we instead group it as:

$$\begin{aligned} C &= [(M_{SML}M_{RGB})^{-1}M_{SML}] \tilde{s} \\ C &= M \tilde{s} \end{aligned}$$

then the matrix M is a *color matching matrix* for this monitor that tells us how much red, green, and blue we should use to match a given spectrum. The rows of this matrix are known as *color matching functions*. They are canonically named by putting bars over the names of the primaries: in this case,

$$M = \begin{bmatrix} -\bar{r} & - \\ -\bar{g} & - \\ -\bar{b} & - \end{bmatrix}$$

(I'm going to give up on the tildes at this point—you know that when I write matrices or vectors I'm talking about the sampled approximations of the actual functions.)

The standard color matching functions

One useful property of color matching functions is that, unlike the actual response functions r_s , r_m , and r_l of the eye, they can be measured just by asking questions of the observer (rather than by actually going and messing with the receptor cells). One column of M tells us how much R , G , and B are needed to match a spectrum that has light at only one wavelength. We can find out what (R, G, B) matches a particular wavelength as follows: Show an observer a pure spectral color (a single-wavelength spectrum) next to a mixture of R , G , and B . Let the observer adjust the intensities of the three primary colors until they match the spectral color.³ Write down the settings the observer chose in the column of M that corresponds to the wavelength you used. By repeating this experiment for a large number of wavelengths one can measure the whole matrix M for a given set of (R, G, B) primaries.

Because they are easily and accurately measured, color matching functions are used in practice as the basis of color computations, rather than the cone response functions r_s , r_m , and r_l . Since color matching functions are just transformations of these “real” response functions, they serve the same purpose in color reproduction.

³Sometimes one of the settings will have to be negative; we can achieve that by moving that light to the other side of the experiment, so that it is added to the spectral color and compared to a mixture of the other two.

To see this, think of applying some arbitrary (but non-singular) 3x3 matrix U to M_{SML} :

$$M_U = UM_{SML}.$$

The matrix M_X is just as good as M_{SML} for computing color reproductions. If we plug in M_X where M_{SML} goes in Equation 1, we get:

$$\begin{aligned} C &= (M_U M_{RGB})^{-1} M_U s \\ &= (UM_{SML} M_{RGB})^{-1} UM_{SML} s \\ &= (M_{SML} M_{RGB})^{-1} U^{-1} UM_{SML} s \\ &= (M_{SML} M_{RGB})^{-1} M_{SML} s \end{aligned}$$

This shows that, if all we want to do is color reproduction (which only depends on which colors look the same or different, not specifically what responses they generate in the eye) then any linear combination of the eye's response functions is as good as any other. So we may as well use the color matching functions, which are easy to measure.

The question remains, What particular color matching functions should we use? There is a particular set of functions, known as \bar{x} , \bar{y} , and \bar{z} , that was standardized by the CIE (*Comission Internationale de l'Eclairage*, or International Commission on Illumination) in 1931, and is universally used as the basis of colorimetry.

The argument that leads to these particular functions is a little involved, but here is the basic outline:

- Start with RGB color matching functions for spectral lights at 700 nm (red), 546.1 nm (green), and 435.8 nm (blue), averaged over a group of observers with normal color vision.
- Set the function $\bar{y}(\lambda)$ to the *luminous efficiency function*, which is a measure of the eye's overall sensitivity to light at each frequency. (This curve is determined by a separate experiment, but it is guaranteed to be a linear combination of the color matching functions because it is derived from the signals produced by the sensors in the eye.)
- Choose the two remaining functions $\bar{x}(\lambda)$ and $\bar{z}(\lambda)$ so that they are positive for all λ and are convenient (by some criteria).

We can think of these three functions as the rows of the canonical color matching matrix M_{XYZ} . The product of M_{XYZ} with a spectrum s is a vector whose

three entries are known as the *tristimulus values* X , Y , and Z of that spectrum. One of the convenient features of XYZ space is that for all spectra X , Y , and Z are positive.

Chromaticity

Three-dimensional spaces, such as color spaces, are awkward to draw pictures of. One common way of reducing a color space to 2D is to project out the overall brightness—that is, to work in a space where all scalar multiples of a particular color are considered to be a single point. Such spaces are known as *chromaticity* spaces because they focus on color to the exclusion of brightness.

A chromaticity space is defined by applying a normalization to a color space that results in numbers that are invariant to scaling the three components together. For XYZ space, the chromaticity coordinates are

$$\begin{aligned}x &= \frac{X}{X + Y + Z} \\y &= \frac{Y}{X + Y + Z} \\z &= \frac{Z}{X + Y + Z}\end{aligned}$$

Note that $x + y + z$ is always 1. For this reason it's conventional to ignore z and simply report x and y .

Chromaticity coordinates are a convenient way to make a plot of the range of possible colors. The most extreme colors are the spectral colors—those that come from single-wavelength spectra. The set of all spectral colors is known as the *spectral locus*. Since all spectra can be made by summing up positive multiples of spectral colors, the image of the spectral locus in chromaticity space is the boundary outside of which no colors exist (that is, there is no spectrum that will give rise to x, y pairs outside the spectral locus).

Transforming between color spaces

Any three color matching functions (that is, any linear combination of the canonical color matching functions $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, and $\bar{z}(\lambda)$) define what's known as a *visual color space*, or a coordinate system for representing colors as seen by the human visual system. We've seen RGB color spaces that correspond to particular primaries, and we've seen the XYZ color space, which does not have any primaries that go with it. In computers, colors are most often represented in RGB spaces that correspond to the primaries of a particular display device, or of a hypothetical standard

display device. The color space is normally specified by the x, y chromaticities of the three primaries and the *white point*, or the chromaticity of the color that results from turning on all three primaries to the same intensity. Together this is enough information to deduce the tristimulus values of the three primaries, up to a scale factor.

So assume we know the tristimulus values $X_r, Y_r, Z_r, X_g, Y_g, \dots$ of the red, green, and blue primaries of a monitor. Since the transformation from RGB color to spectrum to XYZ color is linear, it's just a 3×3 matrix, and since we know the RGB color $(1, 0, 0)$ has to map to the tristimulus values (X_r, Y_r, Z_r) , and similarly for green and blue, the matrix is:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} X_r & X_g & X_b \\ Y_r & Y_g & Y_b \\ Z_r & Z_g & Z_b \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

This gives us the ability to transform colors from RGB to XYZ by knowing where the three canonical basis vectors go—yet another instance of the change-of-basis transformation we've seen so many times in the course. To convert from one RGB color space, say $R_1G_1B_1$, to another, say $R_2G_2B_2$, we just use this matrix twice: once (with the XYZ values of R_1, G_1 , and B_1) to get from $R_1G_1B_1$ to XYZ, then again (with the XYZ values of R_2, G_2 , and B_2) to get from XYZ to $R_2G_2B_2$.