Friday, 2006 January 20, 01:01:08
JAW, Wednesday, 2004 November 17, 11:42:48

Vectorial Color

James A. Worthey
11 Rye Court, Gaithersburg, Maryland 20878-1901, USA.

Abstract

A set of orthonormal color matching functions is presented in which the first is an all-positive achromatic function, the second is red versus green, and the third can be loosely described as blue versus yellow. The achromatic function, proportional to the familiar $\bar{y}$, is a sum of red and green cones. The red-green function uses the same red and green cones, but subtracted, with coefficients so that it is orthogonal to the achromatic one. The third function involves all three cones, but is primarily a blue sensitivity. Using this basis to compute the tristimulus vectors of narrow-band lights at unit power gives Jozef Cohen's Locus of Unit Monochromats, an invariant shape now graphed in a space where the axes have intuitive meaning. The extreme points of Cohen's locus reveal the wavelengths that act most strongly in mixtures, a close approximation to William Thornton's Prime Colors. In effect, decades of research converge in three functions and a vectorial schema, demystifying such issues as color rendering and the selection of additive primaries.

<170 words>

Introduction

With regard to things that stimulate our senses, we all know what amplitude means—usually. A high-amplitude noise SOUNDS LOUD. Summer sunlight looks bright. When adaptation state is controlled, sensation is often a simple—if nonlinear—function of physical amplitude. Color mixing is different, and we all know this too. In a color match, lights add linearly, and the match does not depend on adaptation conditions. Within that context, lights have a direction and an amplitude in color space. Algebraically, we model the mixing of lights by adding triplets of numbers. In the usual system, when lights 1 and 2 are superimposed,

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix}_{mixture} = \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}_1 + \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}_2 . \qquad (1)$$

We may speak of the column matrices in Eq. (1) as *tristimulus vectors*, but vector diagrams and amplitudes in XYZ space are seldom mentioned.

Jozef Cohen saw that any light can be mapped to a unique function of wavelength, its fundamental metamer. In turn, fundamental metamers map to vectors in a 3-space. Cohen's analysis begins with any valid set of color matching functions, even the highly arbitrary $\bar{x}$, $\bar{y}$, $\bar{z}$, then leaps to an invariant formulation, retaining the facts of color matching while losing the arbitrary representation.

In this article, a set of orthonormalized opponent color functions is developed. The algebra will look different from Cohen's, but the results agree entirely with his. Like any set of color matching functions, the orthonormal ones permit a spectral power distribution (471 numbers perhaps) to be reduced to a tristimulus vector, 3 numbers. The components of the tristimulus vector will have intuitive meaning in terms of color names, and also mathematical meaning as components of the fundamental metamer.

The main result, the set of three functions, is simple in itself: a set of graphs, or a table of numbers, to be used in place of other tabulated color matching functions. A method for generating the functions could be stated in a few sentences, reducing this article to a paragraph and a set of graphs. In the event, many paragraphs will be used to put the orthonormal cmf's into the context of past research and future applications.

Background

Alternate sets of cmf's. As an arbiter of matches, a set of color matching functions (cmf's) is not unique; a new set made by linear transformation of the old set predicts the same matches so long as the transformation is reversible (see Appendix A). Figure 1 shows an assortment of cmf's, all equivalent to the CIE 2° Observer in the color matches that they predict. In Fig. 1a is a set of color matching functions in the root meaning of the phrase, the matches that the 2° Observer would make in a visual colorimeter, when the primary lights are narrow bands at 603, 538, 446 nm. In Fig. 1b is a set of cone sensitivities, while in Fig. 1c are the usual functions $\bar{x}$, $\bar{y}$, $\bar{z}$. Finally, in Fig. 1d are a set of opponent functions, essentially Guth's 1980 model with the functions normalized. Guth started with a set of color matching functions considered more accurate, but only a little different from the still-official 2° observer, Fig. 1c (see Appendices B and C). Again, if the only question is whether a light L_1 will match a light L_2, then the 4 sets of cmf's in Fig. 1 will agree on the answer, yes or no.

Superimposing two lights corresponds to adding their tristimulus vectors, and by extension other concepts such as vector amplitude or decomposition should apply to color stimuli. The benefit of vectorial ideas is hard to see in the XYZ schema, but will be more apparent with the right choice of cmfs.

Cone Sensitivities. Consider the sensitivities of the three cone systems[1], Fig. 1b. Whatever stimulus vector amplitude is, it should vanish at the ends of the spectrum and show a local minimum near 495 nm, where all cone sensitivities are low. We also see why the idea of amplitude in color mixing is not as simple as it might be. The cone sensitivities overlap! This is to our benefit in catching plenty of photons and in finely discriminating hues in the red to green range, but it complicates the discussion of color mixing. Consider, for example, the traditional instructions for transforming color mixing data to new primaries. If the 3 cone types had non-overlapping spectral sensitivities, the data would trace out those functions directly. The wavelengths of peak sensitivity could become the primaries in further experiments. Because of overlap, we find a more subtle reality in which changing one primary wavelength alters all 3 functions.

To put it another way, all of colorimetry is about overlap. Overlapping cone sensitivities can be added and subtracted to make the interesting and dissimilar graphs of Figure 1. If non-overlapping functions were added and subtracted, the results would show a more consistent resemblance to the underlying cone functions, a qualitatively different situation.

Opponent Colors, etc. Have these issues of amplitude and overlap been addressed before? Yes, though not with those words, usually. Opponent-color models emphasize the difference between the overlapping red and green sensitivities, literally subtracting one function from the other, with some coefficients. Cornsweet showed the role of overlap for a dichromatic eye, by finding a "spectral locus" of stimuli as wavelength is varied for a light of fixed quantum flux. Cohen's "locus of unit monochromats," though not exactly Cornsweet's locus, is another graphic whose shape comes from spectrally overlapping sensitivities.

Color as it is taught. Superimposed in Figure 2 are two sets of color matching functions, calculated data corresponding to the settings of the idealized 2° observer, in matching a narrow-band test light by 3 narrow-band primaries. For the solid lines, the primary set is {650, 530, 460 nm}, approximating Wright's primaries, while the primaries for the dotted functions are {629, 543, 461 nm}, similar to those used by Guild. The same facts lead to alternate sets of graphs, an awkward situation. Favoring monotony over mystery, teachers and students may rush to embrace the standardized XYZ system. However, despite its puzzling features, the mixing of 3 primaries is a model for television and indirectly for other color technologies.

Figure 3 shows the NTSC primaries for color television in relation to the wavelengths of peak absorption by cones. The peak for red cones is in the yellow, but the red primary—not surprisingly—is in the red, or at least in the orange. These facts are well known, but seldom discussed. The television primaries were chosen by trial and error, a reasonable approach for that well-defined problem.

In Figure 2, looking to the settings of the red primary, Guild's primaries had a possible advantage over Wright's. That is, Guild needed less power at 629 nm than Wright did in his red channel at 650 nm, Table 1. In the horizontal dimension, the red setting has its tall peak at about the same position for either primary set, revealing a wavelength at which the test light *acts strongly*. The primaries could perhaps be set to those wavelengths, permitting a matching experiment at minimum power. Thornton developed this idea in detail [1999 article].

Table 1. Observation concerning peaks of the red cmf in Figure 2. The exact locations of the peaks depend on the green and blue primary wavelengths as well as the red ones. Peak power is expressed as a multiple of the power in the test wavelength

scientist	red primary λ	peak power	λ of peak
Wright	650 nm	3.6	602 nm
Guild	629	1.5	604

Moment of Reflection: Primary Colors

We have seen a puzzle in various forms, which might be called "the question of primary colors." Why is the peak wavelength of the red cmf in Fig. 1a so far from the red cone sensitivity peak in Fig. 1b? In Fig. 3, why did the National Television System Committee, by its trial-and-error process, put the red television phosphor so far from the red cone peak? In the traditional presentation of additive color mixing, Fig. 4, why are the primaries red, green and blue, and not something else? If one is given the cone functions, Fig. 1b, it is easy to infer the cmf's of Fig. 1a, with their more-separated red and green peaks. The riddle of primary colors is this: by what process of cause and effect do the primary colors arise, especially the red primary that is so far from the red cone peak? If color matching data can be converted from one set of primaries to another, why are there in fact Prime Colors, a set of primaries that work best?

In Fig. 1a, for the 2-degree observer, the indicated primaries *are* the Prime Colors, because no more than unit power of each primary is needed to match the test light [Ref: new Brill & Worthey article.]. To make a color-matching apparatus, a three-band light, or a three-color video display, it is a leap in the right direction to set the working primaries close to these Prime Colors, and historically that is done. One description of Prime Colors is that they act strongly in mixtures, and in fact, "action in mixtures" is already what color matching experiments measure, Fig. 1a. The goal now is to make the concept of action in mixtures available for all color discussions, not only for mixtures of narrow-band primaries. Any tristimulus vector expresses action in mixtures, but the orthonormal basis will be especially helpful.

Orthonormal Functions

Consider just red and green cones as a system. Color depends on comparison, so we want to compare red and green cone sensitivities, to find the wavelengths at which they are *the most* *different*. Let the cone functions, Fig. 1b, be called r, g, and b, and compute $r-g$. That is, we want to compare the red and green signals, so with no further thought we subtract them. The resulting opponent color function, Fig. 5, has positive and negative peaks at 605 nm and 520 nm, about where prime colors are expected to fall. We are on the right track, but feel some remorse, because the arbitrary scaling of r and g does affect the result. We then recall that whiteness sensitivity, the familiar $\bar{y}$, is a sum of red and green sensitivities:

$$\bar{y} = 0.6372r + 0.3924g \quad . \tag{2}$$

(The constants can be found by inverting the matrix in Eq. C2, for example.) For later convenience, we can then normalize $\bar{y}$ and call that result ω_1. In fact, $\omega_1 = 0.11381\bar{y}$, but the important requirement is that $\langle \omega_1 | \omega_1 \rangle = 1$. Then a red-green function can be found that is orthogonal to the achromatic function. We find the function that is r minus the projection of r onto ω_1:

$$| \omega_2 \rangle = | r \rangle - | \omega_1 \rangle \langle \omega_1 | r \rangle \quad . \tag{3}$$

Then normalize $| \omega_2 \rangle$:

$$| \omega_2 \rangle \leftarrow | \omega_2 \rangle / \langle \omega_2 | \omega_2 \rangle^{1/2} \quad . \tag{4}$$

(Recall that $|f\rangle$ is function f as a column vector and $\langle g|$ is function g as a row vector, so that $\langle g|f\rangle$ is the inner product of g and f, a single number.) To check that we are on the right track, multiply Eq. (3) on the left by $\langle \omega_1 |$ to obtain:

$$\langle \omega_1 | \omega_2 \rangle = \langle \omega_1 | r \rangle - \langle \omega_1 | \omega_1 \rangle \langle \omega_1 | r \rangle \quad . \tag{5}$$

Since $| \omega_1 \rangle$ was normalized, meaning $\langle \omega_1 | \omega_1 \rangle = 1$, the RHS of Eq. (5) = 0, confirming that $| \omega_1 \rangle$, $| \omega_2 \rangle$ are orthogonal, and in fact orthonormal.

So, in any case, $| \omega_1 \rangle$ and $| \omega_2 \rangle$ are two linear combinations of the red and green cone sensitivities, orthogonal to one another. One is the achromatic function, proportional to $\bar{y}$, and the other is a red-green opponent function, Fig. 6. For concreteness,

$$\begin{bmatrix} \langle \omega_1 | \\ \langle \omega_2 | \end{bmatrix} = \begin{bmatrix} 0.07252 & 0.04466 \\ 0.2670 & -0.3100 \end{bmatrix} \begin{bmatrix} \langle r | \\ \langle g | \end{bmatrix} \quad , \tag{6}$$

though the concepts are more important than the numbers. We can then say that $| \omega_2 \rangle$ represents the other combination of $|r\rangle$ and $|g\rangle$, independent of the achromatic function, $| \omega_1 \rangle$. The remaining arbitrariness resides in the choice of $| \omega_1 \rangle$ as proportional to $\bar{y}$ for the first combination. Other choices for the first function would lead to other possibilities for the orthogonal second function. For the new pair of functions we can write:

$$\begin{bmatrix} \langle \omega_1 | \\ \langle \omega_2 | \end{bmatrix}_{\theta} = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix} \begin{bmatrix} \langle \omega_1 | \\ \langle \omega_2 | \end{bmatrix}_{\theta=0} \quad . \tag{7}$$

The square matrix expresses a rotation, and it would take only brief algebra to show that if the old $\langle\omega_1|, \langle\omega_2|$ (on the right) are orthonormal, so are the new $\langle\omega_1|, \langle\omega_2|$ on the left.

Combining the functions of Fig. 6 into a single plot of $\omega_2(\lambda)$ versus $\omega_1(\lambda)$ gives the spectrum locus shown in black in Fig. 7. Rotating by $\theta = 81.8°$ gives a pair of functions in which one is proportional to $-r(\lambda)$ (but normalized), and the other is orthonormal to that red function, leading to the locus shown in gray. The loci have the same shape, differing only by the rotation. In summary, if we ask where the overlapping red and green functions are the most different, an opponent-color function addresses that question but is somewhat arbitrary. Expressing red-green sensitivity by a pair of orthonormal functions, then plotting one function against the other, leads to a spectrum locus whose shape is invariant except for rotations and reflections. It is convenient to use the original ω_1 and ω_2, so that $\omega_1(\lambda)$ is the achromatic function.

To review, we sought the wavelengths where the red and green cones are the "most different." It then made sense to convert to *functions* that are the most different—a pair of mutually orthonormal functions. Plotting the orthonormal functions together—a so-called parametric plot—gives a 2-dimensional shape. Choosing one of the orthonormal functions fixes the other, within a minus sign. The remaining freedom is a rotation that does not alter the 2D shape.

The new functions are color matching functions like any of the examples in Fig. 1. A narrow-band light of wavelength λ and unit power maps to a vector $[\omega_1(\lambda), \omega_2(\lambda)]$, a point on the locus. The overlap of red and green sensitivities imposes a constraint which could be represented by graphing the green cone function versus the red cone function, more or less what Cornsweet does. Fig. 7 shows the same constraint from overlapping sensitivities, but the orthonormal representation spreads out the locus and provides us with meaningful axes. In Fig. 8 stimulus vectors add vectorially. The locus is not a gamut or boundary in the usual sense. When lights of unit power add, the total vector (using 2 units of power) can extend beyond the locus.

Eq. (3) and (4) above express the method of Gram-Schmidt orthonormalization. Taking blue cone sensitivity $b(\lambda)$ as the next independent starting function, Eq. (3) can be generalized to subtract from $b(\lambda)$ its projection onto each of $\omega_1(\lambda)$, $\omega_2(\lambda)$. Normalizing $|\omega_3\rangle$ as in Eq. (4) completes the calculation, and we have a set of 3 orthonormal color matching functions. In short,

$$\langle\omega_i|\omega_j\rangle = \delta_{ij} \quad , \tag{8}$$

where δ_{ij} is the Kronecker delta, = 1 if $i = j$, = 0 otherwise. These color matching functions, graphed in Fig. 9, can be given intuitive descriptions:

1. ω_1 = achromatic function = $0.11381\bar{y}$.
2. ω_2 = red-green opponent function.
3. ω_3 = a sort of blue-yellow opponent function, but it has contributions from all three cones, and as a practical matter is close to being just a blue cone function.

Let the 3 functions be written as columns of a matrix Ω,

$$\Omega = [|\omega_1\rangle\ |\omega_2\rangle\ |\omega_3\rangle] \quad . \tag{9}$$

(The ket notation, $|\rangle$, makes explicit that the functions of wavelength are column vectors.) Then the tristimulus vector V of a light with spectral distribution L is

$$V = \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = \mathbf{\Omega}^{\mathrm{T}} L \quad . \tag{10}$$

Implicit within the matrix multiplication are the expected three sums of products, so that the first element of V could be written $v_1 = \sum_{\lambda} \omega_1(\lambda) L(\lambda)$, for example.

Now consider a narrow-band light with variable wavelength and a constant one watt of optical power. When its wavelength is λ, the tristimulus vector is $[\omega_1(\lambda)\ \omega_2(\lambda)\ \omega_3(\lambda)]^{\mathrm{T}}$. Varying λ then generates a curve in 3-space. (Superscript T denotes transpose, just to make the vector a column vector.) Fig. 10 shows a static view of a graph that was made in color and 3D with the help of VRML, the Virtual Reality Modeling Language. The locus is the edge of the colored surface. Interactive 3D examples are on the author's web site. See especially the page http://www.jimworthey.com/jimtalk2004nov.html . Clicking on one of the large static examples will bring up an interactive picture. The 3-dimensional shape is in fact Jozef Cohen's "locus of unit monochromats," but found by different steps. Cohen showed that orthonormal primaries lead to the locus of unit monochromats, and his approach via the projector matrix **R** makes it clear that the shape is invariant.

The red, green, and blue extreme points of the locus indicate three "longest vectors" at 604, 536, and 445 nm, explaining why additive primaries usually fall at about those wavelengths. Thinking in terms of 3-phosphor video, for instance, the role of the green phosphor is to pull mixtures away from white and towards green, the red phosphor's role is to pull mixtures towards red. Within a certain power budget, the designer wants long vectors, and the listed wavelengths fill the need. The exact wavelengths given are not ideal primaries for all purposes. For example, the reader may note that 604 nm is an orangish color, not very red. Moving the red primary to a slightly longer wavelength gives a vector at a greater angle from white or green, a more saturated red. Moving the primary to 610 or 620 nm imposes a small cost in power that would increase if the wavelength were shifted further. The motivation for moving the primary beyond 604 nm can also be seen in Fig. 1a, where the red primary is in fact 603 nm. To the right of the red primary is a region where the green cmf goes negative, meaning that green is added to the test light to desaturate it. Moving the red primary to a longer wavelength makes it redder, reducing the amount of green "desaturant" that is needed. The effect can be seen in the animated color matching functions found at http://www.jimworthey.com/matchingprime.html .

The all-too-familiar chromaticity (x, y) indicates the direction of tristimulus vectors, losing their amplitude. Because the vector plot does not discard amplitude, the locus falls into the origin at the short and long ends of the spectrum. Seen in three dimensions, the surface is interesting but not intricately folded. Cohen called it "butterfly wings." Since the ω_1 axis measures whiteness, it is logical to think of the ω_2-ω_3 plane as the chromatic plane. Projecting a tristimulus vector into that plane loses only its whiteness component. Projecting the spectrum locus into the chromatic plane gives the boomerang shape of Fig. 11. As a map of the chromatic component of stimuli, this graph serves a function vaguely similar to that of a chromaticity diagram, but there is no line of purples connecting 400 nm to 700 nm. To make purples with practical technology, one would want to mix lights that have the needed stimulus amplitude within a power budget. The dashed "line of practical purples" indicates such practical mixtures, and the "line of practical blue-greens" shows hues that might better be made as mixtures rather than narrow-band lights. At this time, I am not giving exact definitions for the dashed lines, but suggesting a general idea that can be followed up as the need arises. If **b** and **g** are two vectors, then the constrained mixture $x\mathbf{b} + (1-x)\mathbf{g}$ plots along the straight line between them, so the dashed lines are the loci of unit power mixtures.

Relationship to Cohen's Approach.

Jozef Cohen emphasized that any light can be represented by its fundamental metamer, defined as the linear combination of color matching functions that is metameric to the light. If two or more lights have the same tristimulus vector—meaning that they match for the standard observer—then they have the same fundamental metamer. The light's actual spectral distribution, say $L_1(\lambda)$, minus the fundamental metamer, is a metameric black, a nonzero function whose tristimulus vector is $[0\ 0\ 0]^{\mathrm{T}}$. We shall now see that tristimulus vectors based on the orthonormal color matching functions are proxies for the fundamental metamers of the same lights.

Cohen would find the fundamental metamer by using the projector matrix **R**, which is appealing because **R** is invariant in the strongest possible sense. If any of the sets of cmfs in Fig 1 a-d, or Fig. 9 are called f_1, f_2, f_3, and **A** is a matrix whose columns are those vectors, $\mathbf{A} = [|f_1\rangle\ |f_2\rangle\ |f_3\rangle]$, then

$$\mathbf{R} = \mathbf{A}[\mathbf{A}^{\mathrm{T}}\mathbf{A}]^{-1}\mathbf{A}^{\mathrm{T}} \tag{11}$$

is the projector matrix and it is the same big array of numbers in every case. It is one array for the 2° observer and a different array for the 10° observer, for example, but otherwise it is a fixed arrangement of constants. If $|L\rangle$ is the spectral distribution of a light, and $|L^*\rangle$ is its fundamental metamer, then

$$|L^*\rangle = \mathbf{R}|L\rangle \quad . \tag{12}$$

The fundamental metamer can also be found as a linear combination of the orthonormal cmfs:

$$|L^*\rangle = c_1|\omega_1\rangle + c_2|\omega_2\rangle + c_3|\omega_3\rangle \quad . \tag{13}$$

Applying the usual methods for finding the coefficients c_j (See Appendix D) we find,

$$c_j = \langle\omega_j|L\rangle \quad , \tag{14}$$

leading to

$$|L^*\rangle = |\omega_1\rangle\langle\omega_1|L\rangle + |\omega_2\rangle\langle\omega_2|L\rangle + |\omega_3\rangle\langle\omega_3|L\rangle \quad . \tag{15}$$

In other words, the coefficients in the so-called orthogonal function expansion of $|L^*\rangle$, Eq. (13),

are the tristimulus values of L. The squared length of $|L^*\rangle$ is $\langle L^*|L^*\rangle$. In Eq. (15), multiply the RHS by its transpose, which gives 3 terms times 3 terms, then immediately apply orthonormality, Eq. (8), $\langle \omega_i|\omega_j\rangle = \delta_{ij}$. We find

$$\langle L^*|L^*\rangle = \langle \omega_1|L\rangle^2 + \langle \omega_2|L\rangle^2 + \langle \omega_3|L\rangle^2 \quad . \tag{16}$$

Therefore, the length of the fundamental metamer is the length of the tristimulus vector. For example, if the wavelength domain is 360 to 830 nm with 1 nm steps, there is a sum of 471 terms on the left, and a sum of 3 terms on the right. But the sums, the squared lengths of the vectors, are the same. Or if L_1, L_2 are 2 different lights, the inner product of fundamental metamers equals that of the tristimulus vectors:

$$\langle L_1^*|L_2^*\rangle = [\langle \omega_1|L_1\rangle + \langle \omega_2|L_1\rangle + \langle \omega_3|L_1\rangle]\begin{bmatrix}\langle \omega_1|L_2\rangle \\ \langle \omega_2|L_2\rangle \\ \langle \omega_3|L_2\rangle\end{bmatrix} \quad . \tag{17}$$

Further, the direction cosines are the same for the fundamental metamers and for the tristimulus vectors, and therefore the angular relationships are the same. Therefore, the locus of unit monochromats and related vector diagrams are the same whether they are intended to portray fundamental metamers or tristimulus vectors.

Amplitude and Meaning.

Any tristimulus vector has meaning: it quantifies the stimulus based on the eye's spectral sensitivities. Using the orthonormal cmfs gives clearer meaning to tristimulus vector amplitude. Referring to Eq. (15), the fundamental metamer, $|L^*\rangle$, is an approximation to the physical stimulus $|L\rangle$, and it is also that component of $|L\rangle$ to which color vision can respond. Consider $|\omega_1\rangle\langle \omega_1|L\rangle$ as a component of $|L^*\rangle$. That is, we have a function $|\omega_1\rangle$ scaled by the coefficient $\langle \omega_1|L\rangle$. The component has a sum-squared value $\langle \omega_1|L\rangle^2$ which adds to the sum-squared values of the other two similar components to give the sum-squared fundamental metamer. Therefore, the tristimulus values individually are scaled to the stimulus, and so is tristimulus vector as a whole. Adding the three components plus the fourth one, the metameric black, gives back the physical stimulus $|L\rangle$, and adding the sum-square values of the components plus that of the metameric black gives the sum-squared value of the physical stimulus, $\langle L|L\rangle$.

When tristimulus vectors are added vectorially, as in Fig. 10, an alert student might ask "How do you know what scaling the retina applies in adding the components, what multiple of red-green adds to achromatic?" The answer is, we don't know. The vectors are scaled to the stimulus, and vectorial color is about adding stimuli. Creating and adding stimuli is the goal of all color technology, and vectorial color is about making better use of color matching data. Guth indeed found vectorial addition in brightness experiments, but that is another matter.

Fig. 12 illustrates the notion of component functions. The solid black line is a standardized daylight, D65. The maroon dot-dash line is the fundamental metamer of D65, which can be called D65*. The green, red, and blue lines are the ω_1, ω_2, and ω_3 components which add up to the fundamental metamer. Finally, the gray dashed line is the metameric black, b = D65−D65*. So, D65 is the sum of 4 components, and also the sum-square of D65 is the sum of the sum-squared values of the 4 components. The result that the sum-square of a function, in this case a light, equals the totaled sum-squares of the components, works only for orthogonal components. In this case, if we call the tristimulus values c_j as in Eq. (12), then

$$\langle \mathrm{D65} | \mathrm{D65} \rangle = \sum_{j=1}^{3} c_j^2 + \langle b | b \rangle \quad . \tag{18}$$

The explicit summation in Eq. (18) is a version of the matrix product in Eq. (17). Eq. (18) is a specialized version of Parseval's theorem.

Wavelengths of Strong Action

In the introduction and development above, reference is made to primary colors, colors that act strongly in mixtures, and "longest vectors," concepts that are not completely synonymous. The idea that certain colors act strongly in mixtures can be traced to an article or two of MacAdam, and to articles of William A. Thornton in the 1970s. The idea did not at first lead to a precise definition, but various numerical experiments gave consistent results, showing that wavelengths near 450, 540 and 610 nm act strongly in mixtures. Some uncertainty in the wavelengths of strong action did not impair the practical importance of the idea. Thornton later coined the term Prime Colors for the three wavelengths, and ultimately defined the Prime Colors as the least-power primaries for a color matching experiment. The primaries in Fig. 1a are the Prime Colors for the 2° observer; the primary wavelengths coincide with the peaks of the color matching functions. At those wavelengths, the unit-power test light is matched by unit power of one primary.

Thinking of vectorial color and the locus of unit monochromats, Fig. 10, we can ask which wavelengths give a local maximum in radius from the origin. The question can be answered in 3 dimensions by considering the orthonormal cmfs as vector components, and finding the vector length, then the peaks. Alternatively, one can compute Matrix **R** and then the vector length is the square root of the diagonal of Matrix **R**, which again is vector length as a function of wavelength, from which the peaks are found. It is also interesting to seek the peaks of radius in the two-dimensional chromatic plane, Fig. 11. Table 2 shows results of such calculations.

Table 2: Wavelengths of Strong Action in Mixtures.			
2° Observer			
Longest vectors in 3D are at:	445 nm	536	604
Longest vectors in 2D are at:	445	525	608
Prime colors = least power primaries =	446	538	603
10° Observer			
Longest vectors in 3D are at:	445	535	600
Longest vectors in 2D are at:	445	521	606
Prime colors = least power primaries =	445	536	600

While the "longest vector" wavelengths come from simple manipulations with the orthonormal functions, the prime colors come from color-matching Gedanken experiments or related calculations.

Color Rendering

Most everyday objects do not emit light, so their tristimulus vectors result from a combination of a light source's spectral power distribution (SPD), the object's reflectance, and a set of color-matching functions. For example,

$$V_i = \begin{bmatrix} \langle \omega_1 L | s_i \rangle \\ \langle \omega_2 L | s_i \rangle \\ \langle \omega_3 L | s_i \rangle \end{bmatrix} , \qquad (19)$$

where L is the light, s_i is the spectral reflectance of one surface and V_i is that surface's tristimulus vector. It is realistic to assume a single light L, but numerous surfaces s_i. The light L, and the cmfs ω_j , apply to all objects. If one light is substituted for another, $L_1 \rightarrow L_2$, the color stimuli V_i in general will change. The problem of Color Rendering is to describe the systematic effects of the change $L_1 \rightarrow L_2$, based on facts about color matching, lights, and objects. A vectorial approach can cut through the seeming complexity.

First consider Eq. (19) and the single light L. The three groupings $\langle \omega_j L |$, the object color matching functions, are the same across all surfaces. Now think about this a little further, referring to Figures 10 and 11, and if possible to a VRML 3D drawing of the locus of unit monochromats, such as http://www.jimworthey.com/locusunitmonochvr.html . Table 2 reminds us of the strongly acting wavelengths, and the figures give a sense that the regions of strongest action are well-defined. If two narrow wavelength regions dominate the creation of red and green stimuli, then the power levels in those bands will be more important than other details of the light's spectrum L or the reflectances s_i.

It was noted above that the "locus of unit monochromats" is the locus of fundamental metamers of narrow-band lights, or the locus of tristimulus vectors. Yet another interpretation is that it is *the eye's vectorial sensitivity as a function of wavelength*. To call it "vectorial sensitivity" emphasizes that there is a single independent variable, namely wavelength, but the response is vectorial, having a direction and amplitude in color space. (To say "the eye's sensitivity" is an imprecise but hopefully intuitive reference to the stimulus calculation.) Rather than discuss the groupings $\langle \omega_i L|$ as three functions, we can think of them as a vector, scaled by the variable $L(\lambda)$:

$$\begin{bmatrix} \langle \omega_1(\lambda)| \\ \langle \omega_2(\lambda)| \\ \langle \omega_3(\lambda)| \end{bmatrix} L(\lambda) = \text{vectorial sensitivity} \times \text{SPD of light } L. \tag{20}$$

This product is a function of wavelength, a vector-valued function that could be graphed in three dimensions. Rather than make that graph, we jump ahead to another step.

In Fig. 13, the SPDs of two white lights are plotted. The two lights have the same tristimulus vector, in the orthonormal system or in XYZ for that matter. (Yes, a person could say that "the lights have equal illuminance and chromaticity," but let's think more in vector terms.) One is a light of poor color rendering, high pressure mercury vapor. The other is JMW daylight, adjusted to make its tristimulus vector equal that of the mercury light. (The JMW daylight is a little outside the stated domain of the JMW model.) The vertical dashed lines divide wavelength into 10 nm bands, except for a few wider bands at the ends of the spectrum. In most cases, the *middle* of the band is a round-number wavelength.

For each 1-nm step in λ, Eq. (20) gives a small stimulus vector, and these can be added vectorially within each wavelength band. The sum of those vectors is the tristimulus vector of the light, and it is now interesting to add them graphically, Fig. 14. The simulated daylight gives the smoothly arcing chain of arrows. Mixed long and short arrows distinguish the chain for the mercury-vapor light, with the long components showing the wavelength bands where that light radiates much of its power. The two chains end at the same point by design. A chain of three components, somewhat arbitrarily colored gray, red, and blue, denotes the tristimulus vector of either light, broken into its orthogonal components. The red-green component is positive, but approximately zero.

Although the two lights have the same total stimulus vector, the mercury light takes a shortcut to the final point, while daylight swings in the green direction, then back in the red direction. Now imagine a strongly chromatic object, a red apple for example. To express its redness, it absorbs blue, green and yellow, but reflects red. It selects just the reddest vectorial components from the light, and attenuates other components. If we would say "it selects the red wavelengths," that would leave out colorimetry. But if we say that the apple selects the red vectorial components, that statement *includes colorimetry*.

Discussions of color rendering are often burdened by hidden assumptions, leaving us to infer what people are thinking. All would acknowledge that a light's stimulus vector can be found, the three numbers plotted as gray, red, and blue. Then people seem to assume that the information from simple colorimetry has been used up. The method just developed, based on narrow bands and vectorial addition, shows the sameness of the two lights, that they give the same total stimulus vectors. Within the one presentation, it also shows how they differ: the mercury vapor light lacks red *and* it lacks green. To relate the vectorial graphs to vision of colored objects, one must make the popular assumption that object reflectances are spectrally smooth, so that the 10-nm steps do not conceal an intricate interaction between the light's spectrum and the reflectance. The graphs could be drawn with 1-nm jumps, and the computer programming would be easier, but the arrowheads would be too numerous. In any case, the vector sums are not a rigid "method" to be used blindly, but an explanation of color rendering in terms of basic science.

Moment of Reflection. What is a vectorial method for colorimetry? Why have we not seen anything like the addition of stimulus vectors in Fig. 14? One may recall how vectors are used in basic physics or engineering. Vectors may be added, but a vector may also be *decomposed* into components. The numerical work is simple, but there is not a single "vector method" for the student to memorize. Vector operations are used as appropriate, relating what's known to what isn't. A vectorial treatment of color stimuli exploits the linearity of color matching—Grassmann's laws. We routinely add vectors (X, Y, Z), but there are problems:
1.) The functions $\bar{x}$, $\bar{y}$, $\bar{z}$ are not orthogonal. Stimuli which plot at right angles have fundamental metamers that are not orthogonal. Angles between stimuli in XYZ space are arbitrary, and in fact vector directions are arbitrarily squeezed together. (The overlap of cone sensitivities brings stimuli together, as modeled by the locus of unit monochromats. The XYZ system applies a further arbitrary squeeze.)
2.) Component amplitudes X, Y, and Z are arbitrarily scaled.
3.) The XYZ axes have a murky relationship to the intuitive color concepts of red, green, and blue.
4.) By tradition, we rush to calculate and graph (x, y), which is not a vector.
The orthonormal basis solves these problems, allowing and encouraging us to work with stimulus vectors.

Discussion

Working-class summary. An alert reader may ask, "So, you are selling a new set of color matching functions. At one point you seek the wavelengths where receptor sensitivities are the 'most different,' and that is interesting. But in the end, what do you have that is new?" If we consider $\bar{x}$, $\bar{y}$, $\bar{z}$, Fig. 1c, as the "old" functions, then ω_1 is a multiple of $\bar{y}$, so it is not new. The third orthonormal function, ω_3, is really a blue function with a little extra variation, therefore it is little different from the old $\bar{z}$, Fig. 15. So, 2/3 of the "new" system is similar to the old system. The remaining old function is $\bar{x}$, an arbitrary magenta primary. The orthonormal system replaces $\bar{x}$ with the red-green opponent function, ω_2, allowing the axes to have intuitive meanings, namely white, red or green, and blue or yellow. Fig. 15 also compares $\bar{x}$ to ω_2. The new system is an opponent-color scheme whose greatest benefit is to resolve the overlapping red and green cones into functions that are more independent, and indeed orthogonal. Only a few narrow-band stimuli will actually plot in the yellow (minus blue) direction, but the range of possibilities from yellow through shades of white to blue is nicely spread out in the v_1-v_3 plane. Any set of color-matching functions can map stimuli to vectors. Thanks to orthonormality, the new cmf's spread out the vectors as much as possible.

Locus of Unit Monochromats. Jozef Cohen started with Wyszecki's idea, to break any SPD into a metameric black, orthogonal to all three cmf's, and the fundamental metamer, which is a linear combination of the cmf's. A light and all its metamers share one fundamental metamer, which is a non-arbitrary proxy for the group. Although the fundamental metamers are functions of wavelength, 471-vectors perhaps, they map to 3-vectors. Each unit-power narrow-band light, or "monochromat" in Cohen's usage, thus maps to a point on the Locus of Unit Monochromats. The reference to "unit power" calls attention to the vector lengths: wavelength governs direction *and* amplitude. The too-familiar chromaticity (x, y) by itself quantifies direction *but loses amplitude*. Choosing the orthonormal functions also establishes axes, and it was found above that a three-component stimulus vector in the orthonormal system has the same direction and amplitude as the related fundamental metamer. Vector amplitude has the meaning of "strength of action in mixtures," a concept that is near the surface in color-matching experiments, but was particularly called to our attention by Thornton (as well as MacAdam).

Vectorial color. In Fig. 14, the vector chain gray-red-blue stands for the familiar operation in which three sums over the visible spectrum give a light's tristimulus values. That's very well, but summing every time over the whole spectrum loses the peaks and valleys that distinguish the one light from another. The other chains in the figure come from rearranging the sum: choose a wavelength band, collect the terms for that band from each larger sum, treat three partial sums as a vector. The vector then has a direction in color space and an amplitude; from the abstraction of "wavelength bands" comes something more concrete: stimulus vectors. Isaac Newton said the rays are not colored, but he didn't have computer graphics. The linearity of retinal transduction permits us to draw colored arrows, if not rays, and see exact meanings in the results.

Cause and effect. The best cause-and-effect reasoning is at work when a phenomenon is described by simple laws, especially linear laws like $F=ma$ and methods that assume linearity, such as vectors. One then computes "the answer" and words such as "effect" are seldom used. The facts of color vision present themselves in a certain way, as functions of wavelength, Fig. 1. Linearity is implicit in any calculation when we compute (SPD)×(sensitivity) at each wavelength, then *add* those contributions. The XYZ formulas exploit linearity to predict matching lights or paint chips, but the arbitrariness of X, Y, Z deters the use of vector diagrams. Losing vectors means losing cause and effect. Consider the problem of color rendering. Where does the problem arise, at the linear stage or later? Fig. 14 confirms what has been said before: the problem involves systematic differences between lights, and arises in the linear stage where cause and effect can be made clear. The traditional color rendering method is full of arbitrary elements such as the 8 color chips, but its problems go deeper. It is based on an assumption that the problem is hopelessly complex: color shifts might be random or might be systematic; maybe the object colors are important; maybe nonlinearity of perception is important. We don't know and it's beyond hope that we could find out, so we had better allow for all possibilities. The XYZ system generates this pessimism by cutting short the discussion of linear events.

Terminology. In Fig. 1b, we see the great overlap of the red and green cone sensitivities, and the lesser overlap of blue with green or red. Adding and subtracting such functions gives the surprising variety seen in Fig. 1 and Fig. 9. Assuming that the retina somehow compares red and green signals, overlapping sensitivities permit the gradual variation of hues that we experience, even with narrow-band lights. College professors, finding their students confused about hue variation from three receptors, and other subtleties, often call the receptors long, middle, and short. I take a different view. The names red, green, and blue appeal to intuition at any educational level. The complexity of color vision begins with overlap, a concrete fact that can be discussed and taught to students. The orthonormal basis deals with overlap by use of an opponent-color scheme, leading to the Locus of Unit Monochromats, an intuitive embodiment of trichromatic theory. The orthonormal basis is **not** a physiological model, but it offers a rational color space in which color science concepts such as unique hues, or Guth's functions, Fig. 1d, can be represented by vectors.

Lighting. Lighting books begin with discussions of luminous flux, treating light as a liquid and tossing aside the roles that optics and color play in making a light useful. Eschewing even the 1931 color-vision observer, they base their discussion on the one-dimensional 1924 observer: the $\bar{y}$ function alone. Summarizing Cohen's and Thornton's ideas as evolved in this article, we can acknowledge color vision by saying that light is *three liquids*. Thornton invented names for them: rubinosity, verdinosity and bluminosity. If the light supplies less than a normal dose of power in the regions near the prime colors (Table 2), color contrasts are diminished. Calling for three liquids is a clumsy simplification, but less clumsy than the one-liquid model. Apart from color, the over-riding optical issue is the light's area or luminance. Area and luminance must vary inversely for a given luminous flux. Optics teaches that luminance can diminish but not increase; no "optical funnel" can focus light from a fluorescent tube into a narrow spotlight beam.

Chromaticity. A version of chromaticity could be defined for the orthonormal space. The vectors comprising the locus of unit monochromats would be projected outward to meet a certain plane, creating a spectrum locus that maps direction but not vector amplitude. Other stimulus vectors could be projected to that plane, again losing their amplitude. The task of defining such a diagram is left for the future. For now, we have an opportunity to forget chromaticity and think about vectorial color.

Math. The orthonormal basis resembles prior opponent color models, but is not the same. One may ask how the orthonormal system, relying heavily on mathematical ideas, can presume to supersede a physiological model based on "real data." A two-part answer can be given:

1. A physiological model and the orthonormal basis are answers to different questions. There is no need for them to be the same.
2. Choosing an orthonormal set of color matching functions allows perpendicular stimulus vectors to represent independent variations of the stimulus, conforming to the usual notion of vectors. The invariant locus of unit monochromats emerges, revealing the strong-acting red, green and blue Prime Colors. Thus, linear matching data point to a special status for red and green, which in turn points the way toward opponent colors.

Critical tests. When a version of vectorial color was presented orally, one audience member asked what would be a critical test of the theory. The answer I gave then was "television phosphors." A vectorial picture (Figs. 10, 11 and 14) shows why additive primaries are red, green, and blue (Figs. 1a, 3 and 4). That is one critical test. The test must involve more than a single color match, because the orthonormal functions are a transform of the CIE's 2-degree observer functions, and predict the same matches; see Appendix A.

In the XYZ scheme, no two of the cmfs are orthogonal, and in particular when X and Y are measured for a population of lights or paint chips, they will tend to be correlated because of the overlap of $\bar{x}$ and $\bar{y}$. (Direction cosine of $\bar{x}$ and $\bar{y}$ = 0.760 .) The components of V, Eq. 10, will tend to be less correlated and therefore more helpful in the context of measurement and quality control. If derived quantities involve adding or multiplying measured values, simple formulas will apply for error propagation so long as the measurements can be assumed independent; otherwise there is an onerous requirement to estimate various measures of correlation. [NIST/SEMATECH e-Handbook of Statistical Methods, http://www.itl.nist.gov/div898/handbook/ , 2006.]

Cohen's space is a logical one for adding and graphing vectors, as in Fig. 14. The color rendering example above, and similar ones, will aid understanding of lights that are similar in their total tristimulus vector, but different in their detailed makeup. The opponent feature of the orthonormal cmfs helps with color naming, but it has other less obvious uses also. A separate short article presents realistic "Applications of Vectorial Color."

Historical development. This article shows that expressing the facts of color matching in orthonormalized opponent functions yields benefits for the understanding of color. Readers may ask how such an idea was developed. The answer is very much in the open literature. Cohen studied dependency of spectral reflectances, which led him to think about dependency of color matching functions, and ultimately about Matrix R. Thornton started out with engineering calculations to improve fluorescent lamps, which led him to the Prime Colors, which he then studied through many published articles. Cornsweet showed the role of overlap and emphasized discussing color without the CIE's concepts. Michael H. Brill contributed theorems and collaborations. Sherman Lee Guth offered the idea of a simple opponent model. I began by combining Thornton's and Guth's ideas. I found uses for opponent colors as a calculation tool. Buchsbaum derived a set of orthonormal opponent cmf's. In a recent article on color rendering, I sketched the outlines of a relationship of Prime Colors, opponent colors, and Matrix R. In the next article I used orthonormal opponent functions, but only as a tool, an important middle step. The final sentences of that article reveal a last-minute discovery: that Matrix R is the same as a unity operator made from an orthonormal basis. See **Appendix D** below. After that insight, some work with computer graphics was important in moving from the abstract idea of orthonormal color matching functions to a better understanding of practical benefits.

In 1998 I was contemplating the work that became Reference __. Thornton suggested that I not use Guth's 1980 model directly, but make a version consistent with one of the CIE observers. I have followed that advice since and now it blossoms into a larger idea. Figure 1, for example, shows cmf's that are linear transformations of each other, but have differing interpretations. Even Fig. 2 shows pseudo-historical graphs, similar to the Wright and Guild data, but based on the 2° observer. The use of consistent functions is not new, and indeed it is part of the 1980 Guth model. Still, basing all development on the CIE's smooth functions helps to separate measurement issues from algebraic ones. The opponent-orthonormal schema can then incorporate revised data, Appendix C, or even be applied to alternate systems such as cameras.

In order to use the orthonormal basis, it is not necessary to review all the historical development. To the contrary, a stimulus vector is calculated from the orthonormal cmf's as from any other set of color matching functions, Eq. (10). A person can then see the benefit of the orthonormal and vectorial methods by working with them. I praise the contributions of Brill, Cohen, Guth, Thornton and others (alphabetical order!). Their childhood homes should become shrines visited by schoolchildren. While we wait for that to happen, we can honor them by using the orthonormal opponent color matching functions, a simple embodiment of their hard work and clever ideas. The vectorial schema will be so logical that people will feel as if they always knew it, that it must have a simple history. When nobody can imagine that hard research was needed, that will honor their thorough work.

Conclusion

So-called color-matching data, the foundation of colorimetry, derive from experiments in which lights are added. Conceptually, three primaries are adjusted so that their sum matches a test light. Color technologies are much like the basic experiments: primary colors vary in power, but not in the shape of their spectral distributions. Even the viewing of colored objects under a white light fits this model: the primaries are the narrow bands within the spectrum (Figure 13) which add to give the white light (Figure 14). Colors add linearly, which should aid analysis and permit vector diagrams. Unfortunately, in the 20th century colorimetry became tied to the XYZ scheme. In the CIE's method, tristimulus vectors in the form [X Y Z] put arbitrary non-orthogonal quantities on the axes and squeeze together the cone of color vectors, discouraging any thought of vector diagrams. In the transition to chromaticity (x, y), color vectors are subject to a central projection which preserves direction but loses amplitude. Most color discussions then involve (x, y).

As the XYZ schema is used, it is slanted toward quality control. A test piece, such as a colored signal light, is measured with a spectroradiometer. In one step then, a computer converts the spectrum to three numbers such as (Y, x, y). The numbers emerge as from a black box, because the intermediate values X, Y and Z lack intuitive meaning. (Y has meaning, but not about color.) In the mid-twentieth century, the numbers were a boon and gave objectivity to users of signal lights and paints. In the pre-computer era, it was a serious task to "grind out" such a calculation, and the pencil-worker was thankful for the lack of negative numbers. But also in the mid-twentieth century, fluorescent lights and color television were invented, and other color technologies evolved. Considering color television, for example, it is not enough that an artist chooses lavender, lilac and aqua phosphors, and then the engineer does quality control. The engineer must be a jump ahead of the artist and pick three colors that are strong and dissimilar so that their mixtures will be diverse. As Thornton discovered, information about the strong and dissimilar colors of human vision is near the surface in the data of a color matching experiment, Fig. 1a. It is obscured by the XYZ system, but easily applied using the orthonormal basis and Cohen's space.

For lack of a good vector representation, the XYZ schema engenders low expectations. Signal lights are discussed rationally, but when a white light shines on some objects no clear insight is offered. Color rendering is treated as mysterious, yet dull. In fact, color rendering is about how white lights work, a basic topic. In Fig. 14, consider the chains of vectors for the two lights. The component arrows in each chain are calculated directly by colorimetry, based on the partitioning into wavelength bands, Fig. 13. No approximation or hidden assumption is needed, just vector methods which embody the trichromatic theory.

Appendix A, How Different Sets of Cmf's Can Predict the Same Matches

If one set of color matching functions (cmf's) is known, other sets can be derived from it that predict the same matches. Suppose that the starting set of cmf's form the 3 columns of a matrix $\mathbf{Q}_0$, perhaps representing the data from a color matching experiment:

$$\mathbf{Q}_0 = [|r\rangle\ |g\rangle\ |b\rangle]\ . \tag{A1}$$

Then two spectral power distributions L_1 and L_2 are expected to match if

$$\mathbf{Q}_0^{\mathrm{T}} L_1 = \mathbf{Q}_0^{\mathrm{T}} L_2\ . \tag{A2}$$

Superscript T means matrix transpose. Eq. (A2) is an equality of tristimulus vectors. Now define a transformed set of cmf's $\mathbf{Q}_1$:

$$\mathbf{Q}_1^{\mathrm{T}} = \mathbf{X}\, \mathbf{Q}_0^{\mathrm{T}}\ , \tag{A3}$$

where $\mathbf{X}$ is an invertible (non-singular) 3×3 matrix. ($\mathbf{X}$ is a transform, not related to the XYZ system.) Let Eq. (A2) be multiplied on both sides by $\mathbf{X}$,

$$\mathbf{X}\mathbf{Q}_0^{\mathrm{T}} L_1 = \mathbf{X}\mathbf{Q}_0^{\mathrm{T}} L_2, \tag{A4}$$

then apply Eq. (A3) to conclude that

$$\mathbf{Q}_1^{\mathrm{T}} L_1 = \mathbf{Q}_1^{\mathrm{T}} L_2\ . \tag{A5}$$

Since $\mathbf{X}$ is invertible, multiplying Eq. (A5) by $\mathbf{X}^{-1}$ gives back Eq. (A2), all steps are reversible and Eqs. (A2) and (A5) are equivalent.

Appendix B, Relationship to Guth's 1980 Model.

If one were to poll Sherman Lee Guth and his former students, one might find that the concepts I take from his work differ from the ideas he had in mind in 1980. In fact, it has been a theme of my career to apply Guth's 1980 model as a mathematical tool, in ways that he did not intend. Opponent concepts prove valuable for understanding the linear stage of color vision, the stage of transduction and of color matching, prior to the stage that Guth intended to study. In this article, I stray especially far from his vision, adding and subtracting his functions to make orthonormal ones.

Emphasizing the evolution of thought, rather than formulas, the notion of orthonormalizing Guth's functions arises like this:

1. Guth conformed his opponent functions to the familiar facts of color matching by starting with 2° observer functions, $\bar{x}$, $\bar{y}$, $\bar{z}$. He used a "slightly modified" 2° observer, but I use the official CIE 1931 functions with his formulas. In two steps, he gave formulas for cone sensitivities (from Smith and Pokorny), then subtracted the cone signals to give opponent signals. When we recognize that receptor-sensitivity overlap is a central theme of colorimetry, then Guth's 1980 model is charming and simple because:

a. The achromatic function, a multiple of $\bar{y}$, is a sum of the red and green signals, with certain coefficients. It has no blue input.

b. The red-green signal is red minus green only, with specified coefficients, no blue input. Thus, the red-green signal involves only the most-overlapping pair of cones.

c. The blue-yellow signal is blue minus red only, with coefficients, no green input. This signal involves only the least-overlapping pair of cones.

2. Guth normally discussed his functions in the sequence achromatic, red-green, and blue-yellow (symbols $\bar{a}$, $\bar{t}$, $\bar{d}$). One may notice that:

a. Opponent functions such as Guth's look like they might be orthogonal functions.

b. They are sequenced in a way that would be natural if they were orthogonal, namely the all-positive function first, then two functions that cross the abscissa and go negative over part of the wavelength domain.

c. The functions are the two highly overlapping cones (red and green) added; then the same cones subtracted; then the two least overlapping cones (blue and red) subtracted.

3. So, we are perhaps disappointed that Guth's functions are not orthogonal.

a. However, it is logical to apply the Gram-Schmidt procedure, keeping the functions in the order achromatic, red-green, blue-yellow.

b. When Gram-Schmidt is applied in this way, the achromatic function is renormalized, but otherwise unchanged. The achromatic function was a re-scaled $\bar{y}$, and it is still a re-scaled $\bar{y}$ in our orthonormal set.

c. The red-green function is changed, to make it orthogonal to $\bar{y}$, but because of the step-by-step nature of the Gram-Schmidt procedure, the orthogonalized red-green function still has the Guth-like property of no blue input. It is still red minus green, but the exact coefficients are new.

d. Finally, the blue-yellow function in the orthonormal set is the least Guth-like, with contributions from all 3 cones.

In the main text, the orthonormal set was found, in effect, by starting with the set $\{\bar{y}$, red cones, blue cones$\}$ and applying Gram-Schmidt. The resulting functions are the same, with the same connection to Guth's 1980 model.

Appendix C, Smooth Cone Sensitivity Functions and Other Calculations.

In the body of this article, any "achromatic sensitivity" is a multiple of the CIE's $\bar{y}$ for the 1931 2° Observer, and any "cone sensitivity" is a linear combination of $\bar{x}$, $\bar{y}$, and $\bar{z}$ for the 2° Observer. One reason for using the CIE functions in this way is that they are (unofficially) available as smoothed functions, tabulated at 1-nm wavelength steps. One may ask, how much difference would it make to have a different starting point, such as the more recent Stockman and Sharpe receptor sensitivities? In fact, it makes little difference, but it will be instructive to work out the comparison.

The derivation of the orthonormal cmf's required, as a starting point, a set of cone sensitivities and the relationship between cone functions (Figure 1b) and the whiteness function $\bar{y}$ (Fig. 1c or 1d). Equations (C1) and (C2) serve to concoct the red, green and blue cone functions:

$$\begin{bmatrix} \langle \bar{r} | \\ \langle \bar{g} | \\ \langle \bar{b} | \end{bmatrix} = \mathbf{M}_1 \begin{bmatrix} \langle \bar{x} | \\ \langle \bar{y} | \\ \langle \bar{z} | \end{bmatrix} \tag{C1}$$

where

$$\mathbf{M}_1 = \begin{bmatrix} 0.2435 & 0.8524 & -0.0516 \\ -0.3954 & 1.1642 & 0.0837 \\ 0 & 0 & 0.6225 \end{bmatrix} . \tag{C2}$$

The bra notation, such as $\langle \bar{r} |$, says that the functions are written as row vectors; they become the rows of the larger matrices. Eqs. (C1) and (C2) are taken from Guth, but slightly misapplied, since he intended that they be used with Judd's slightly modified $\bar{x}$, $\bar{y}$, $\bar{z}$, deemed more accurate.

Putting aside the Judd functions, suppose that the Stockman and Sharpe cone primaries are to be used. Let them be written as column vectors, $|r_{SS}\rangle$, $|g_{SS}\rangle$, $|b_{SS}\rangle$. Now, using the assumption that achromatic sensitivity involves only red and green cones, let $\mathbf{A} = [|r_{SS}\rangle\ |g_{SS}\rangle]$ and construct a projector matrix $\mathbf{R}_{rg}$ for the 2-dimensional space of these receptors:

$$\mathbf{R}_{rg} = \mathbf{A}(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T \ . \tag{C3}$$

If $|\bar{y}\rangle$ is the usual achromatic function, and $|\bar{y}_{rg}\rangle$ is to be the best fit to $|\bar{y}\rangle$ in the 2-dimensional space, then

$$|\bar{y}_{rg}\rangle = \mathbf{R}_{rg}\,|\bar{y}\rangle \ . \tag{C4}$$

Now recall the scheme of the smooth orthonormal functions used earlier. That is, the first function was proportional to $|\bar{y}\rangle$, the second function was a different linear combination of red and green cones, orthogonal to $|\bar{y}\rangle$, and the third function involved all 3 cones. If we form a matrix $[|\bar{y}_{rg}\rangle\ |r_{SS}\rangle\ |b_{SS}\rangle]$ and perform Gram-Schmidt orthonormalization on the columns of this matrix in sequence, the result will be a new orthonormal basis comprising linear combinations of the Stockman and Sharpe primaries:

$$[|\bar{y}_{rg}\rangle\ |r_{SS}\rangle\ |b_{SS}\rangle] \rightarrow \text{by Gram Schmidt} \rightarrow \mathbf{\Omega}_{SS} \ , \tag{C5}$$

where $\mathbf{\Omega}_{SS}$ is the new orthonormal set. Figure C1 shows the fit of the less smooth basis to the smooth one. At worst, one can say that the smooth functions differ little from more exacting modern ones. At best, the smoother functions convey the "big story" better. In Fig. C1, the smoother functions based on the interpolated 2° Observer are shown as gray dashed lines, while those based on the Stockman and Sharpe primaries are the thinner black lines.

The calculation just done combines two methods—Matrix **R**, and Gram-Schmidt. Suppose now that a spectral power distribution is given, and one seeks its projection into the vector space of color matching functions, its fundamental metamer. There are (at least) two paths:

1. One path uses the logic of Fourier series. Orthonormal color matching functions are available, or can be created by the Gram-Schmidt procedure. Coefficients are calculated by the proper simple formulas, and the original SPD is approximated by a 3-term series, as in the beginning of Appendix D.
2. By the alternate path, a projector matrix **R** is calculated, then the approximation to the SPD is calculated by a single final matrix multiplication, similar to Eqs (C3) and (C4), but with 3 column vectors in matrix A.

The resulting vectors—projections of the SPD—will be the same, within some roundoff error. In both cases, the projection into the vector space of color matching functions *is* the least-squares best fit to the SPD by a linear combination of cmf's. For the one task, path 2 may be more convenient, but if the numerical coefficients are needed, then path 1 generates them as a byproduct. A decade or two ago, a practical person might have noticed that **R** can be an extremely large matrix (up to 1.8 MB as double precision) and the path via Gram-Schmidt is the more efficient. With today's computer hardware, the occasional inefficient algorithm does no harm. Any logical method may be used, perhaps followed by other steps for checking.

Suppose that $\mathbf{\Omega}$ has just been calculated, a matrix whose columns are intended to be orthonormal functions. By the orthonormal property, one should find $\mathbf{\Omega}^{\mathrm{T}}\mathbf{\Omega} = \mathbf{I}$, where $\mathbf{I}$ is a small identity matrix, 3×3 for example. This property, that $\mathbf{\Omega}^{\mathrm{T}}\mathbf{\Omega} = \mathbf{I}$, simplifies the formula for projector matrix $\mathbf{R}$, such as Eq. (C3). That is,

$$\mathbf{R} = \mathbf{\Omega}\mathbf{\Omega}^{\mathrm{T}} \text{, only when columns of } \mathbf{\Omega} \text{ are orthonormal.} \tag{C6}$$

If the functions $\mathbf{\Omega}$ are linear combinations of some other functions, such as a particular set of cone sensitivities, then the Matrix $\mathbf{R}$ in Eq. (C6) should be numerically the same as found by Cohen's formula, $\mathbf{R} = \mathbf{A}(\mathbf{A}^{\mathrm{T}}\mathbf{A})^{-1}\mathbf{A}^{\mathrm{T}}$, based on the original functions. The computer can subtract the two versions of $\mathbf{R}$, then search the difference matrix for the greatest absolute discrepancy. Applying this checking calculation to the two versions of $\mathbf{R}$ based on Stockman and Sharpe cones gives a maximum difference of 3×10^{-16}. For comparison, the mean absolute value of an element of $\mathbf{R}$ in this case was 1.4×10^{-3}.

Appendix D, Fun with Orthonormal Functions

This appendix is really about notation and convenient calculation. The ideas are the well-known facts of generalized Fourier series. Suppose that $\{|\omega_1\rangle, |\omega_2\rangle, |\omega_3\rangle\}$ are a set of functions that are orthonormal:

$$\langle \omega_i | \omega_j \rangle = \delta_{ij} \quad . \tag{8}$$

To be concrete we envision a set of 3 functions of wavelength, but there could be any number of functions over any domain. Now consider a function $|L\rangle$, which could be the spectral power distribution of a light. We want to approximate $|L\rangle$ by a linear combination of the functions $|\omega_i\rangle$:

$$|L\rangle \approx c_1|\omega_1\rangle + c_2|\omega_2\rangle + c_3|\omega_3\rangle \ , \tag{D1}$$

where the c_j are constant coefficients. We seek a formula for c_1. Multiply Eq. (D1) on the left by $\langle \omega_1 |$:

$$\langle \omega_1 | L\rangle \approx c_1 \langle \omega_1|\omega_1\rangle + c_2 \langle \omega_1 | \omega_2\rangle + c_3 \langle \omega_1 | \omega_3\rangle \ . \tag{D2}$$

By orthonormality, Eq. (8), $\langle \omega_1 | \omega_2\rangle = \langle \omega_1 | \omega_3\rangle = 0$, and $\langle \omega_1|\omega_1\rangle = 1$. Then $c_1 = \langle \omega_1 | L\rangle$. In general,

$$c_j = \langle \omega_j | L\rangle \ , \tag{D3}$$

and then substituting Eq. (D3) into Eq. (D1) leads to

$$|L\rangle \approx |\omega_1\rangle\langle \omega_1 | L\rangle + |\omega_2\rangle\langle \omega_2 | L\rangle + |\omega_3\rangle\langle \omega_3 | L\rangle \ . \tag{D4}$$

By reasoning not reviewed here, the sum on the right in Eq. (D4) is the linear combination of $|\omega_1\rangle$, $|\omega_2\rangle$, $|\omega_3\rangle$ that is the least-squares best fit to $|L\rangle$. That is also a description of the fundamental metamer of $|L\rangle$, denoted by $|L^*\rangle$. Therefore,

$$|L^*\rangle = |\omega_1\rangle\langle \omega_1 | L\rangle + |\omega_2\rangle\langle \omega_2 | L\rangle + |\omega_3\rangle\langle \omega_3 | L\rangle \ . \tag{D5}$$

(Why is '≈' gone in Eq. (D5)? Because the fundamental metamer *is* the approximation.) Factoring the RHS of Eq. (D5) yields

$$|L^*\rangle = (|\omega_1\rangle\langle\omega_1| + |\omega_2\rangle\langle\omega_2| + |\omega_3\rangle\langle\omega_3|)|L\rangle \ , \tag{D6}$$

or,

$$|L^*\rangle = \left(\sum_{j=1}^{3} \left|\omega_j\right\rangle\left\langle\omega_j\right|\right) |L\rangle \ . \tag{D7}$$

The sum in parentheses is called a unity operator, $\mathbb{1}$, and could have N terms:

$$\mathbb{1} = \sum_{j=1}^{N} \left|\omega_j\right\rangle\left\langle\omega_j\right| \quad . \tag{D8}$$

Comparing Eq. (D7) to Eq. (10) shows that the unity operator performs the same function as projector matrix **R**, and suggests that it *is* Matrix **R**. Recall that $|\omega_i\rangle$ is a column matrix and $\langle\omega_j|$ is a row matrix, so $|\omega_i\rangle\langle\omega_j|$ is a large square matrix. The sum within the matrix multiplication is implicit; the sum from j=1 to N is separate and indicates the sum of N large square matrices. So $\mathbb{1}$ is a large square matrix like **R**, it is not yet proved that they are equal.

Postponing that proof, why do we need a different symbol and formula for **R**, Eq. (D7)? We need the unity operator not as a formula for the projector matrix, but as a shorthand way to derive equations like Eq. (D4) or (D5), which include explicit formulas for the coefficients, as in Eq. (D3). Now letting N=3, we notice an alternate way of writing Eq. (D8):

$$\mathbb{1} = \left[\left|\omega_1\right\rangle \quad \left|\omega_2\right\rangle \quad \left|\omega_3\right\rangle\right] \begin{bmatrix} \left\langle\omega_1\right| \\ \left\langle\omega_2\right| \\ \left\langle\omega_3\right| \end{bmatrix} \quad . \tag{D9}$$

In this case, the summation implicit in the matrix product is the one written explicitly in Eq. (D8). It is natural to write the orthonormal set as the columns of a matrix $\mathbf{\Omega}$, that is

$$\mathbf{\Omega} = [|\omega_1\rangle \ |\omega_2\rangle \ |\omega_3\rangle] \quad . \tag{D10}$$

Therefore

$$\mathbb{1} = \mathbf{\Omega}\,\mathbf{\Omega}^{\mathrm{T}}. \tag{D11}$$

Now the formula for **R** holds for any transformed set of color matching functions, so substitute $\mathbf{A} = \mathbf{\Omega}$ in Eq. (9):

$$\mathbf{R} = \mathbf{\Omega}[\mathbf{\Omega}^{\mathrm{T}}\mathbf{\Omega}]^{-1}\mathbf{\Omega}^{\mathrm{T}} \quad . \tag{D12}$$

But the grouping $\mathbf{\Omega}^{\mathrm{T}}\mathbf{\Omega}$, because of orthonormality, is the 3×3 identity matrix:

$$\mathbf{\Omega}^{\mathrm{T}}\mathbf{\Omega} = \mathbf{I}_{3\times 3} \quad , \tag{D13}$$

whose inverse is also the identity matrix, therefore

$$\mathbf{R} = \mathbf{\Omega}\mathbf{\Omega}^{\mathrm{T}} \quad . \tag{D14}$$

Comparing Eq. (D11) to Eq. (D14) confirms that $\mathbf{R} = \mathbb{1}$.

The key idea of this appendix is contained in Eq. (D4) or (D5), or in Eq. (D8) or (D9), which are tools for deriving those equations. Eq. (D3) is the explicit formula for the coefficients. If the summation notation of Eq. (D8) seems awkward, Eq. (D9) can be used to derive formulas. For example, $|L^*\rangle = \mathbb{1}|L\rangle$, then

$$|L^*\rangle = \begin{bmatrix} |\omega_1\rangle & |\omega_2\rangle & |\omega_3\rangle \end{bmatrix} \begin{bmatrix} \langle\omega_1| \\ \langle\omega_2| \\ \langle\omega_3| \end{bmatrix} |L\rangle \quad . \tag{D15}$$

On the RHS, three matrices are multiplied. Formally multiplying the second and third matrices gives

$$|L^*\rangle = \begin{bmatrix} |\omega_1\rangle & |\omega_2\rangle & |\omega_3\rangle \end{bmatrix} \begin{bmatrix} \langle\omega_1|L\rangle \\ \langle\omega_2|L\rangle \\ \langle\omega_3|L\rangle \end{bmatrix} \quad . \tag{D16}$$

Formal multiplication in Eq. (D16) then gives Eq. (D5), the desired result. A further succinct insight is that $\mathbf{\Omega}^{\mathrm{T}}\mathbf{\Omega} = \mathbf{I}_{3\times 3}$, Eq. (D13), but if the order of multiplication is reversed, $\mathbf{\Omega}\mathbf{\Omega}^{\mathrm{T}} = \mathbf{R}$, Eq. (D14).

Application. Suppose that we seek the relationship between the orthonormal vectors, $\mathbf{\Omega}$, and Guth's opponent functions (renormalized as in Fig. 1d). Call the array of Guth's vectors $\mathbf{G}$. Then

$$\mathbf{G} = \mathbb{1}\mathbf{G} \quad . \tag{D17}$$

In Eq. (D17), there is equality and not approximate equality because we know that the Guth color matching functions are linear combinations of the columns of $\mathbf{\Omega}$. Apply Eq. (D11):

$$\mathbf{G} = \mathbf{\Omega}\, \mathbf{\Omega}^{\mathrm{T}}\, \mathbf{G} \quad . \tag{D18}$$

A realistic situation is assumed: that $\mathbf{G}$ and $\mathbf{\Omega}$ exist on a computer as arrays of numbers. It might be that $\mathbf{\Omega}$ was just found from $\mathbf{G}$ by the Gram-Schmidt algorithm. We now seek a 3×3 matrix that is the transform from one to the other. All that we need to do is group the terms. Define $\mathbf{X} = \mathbf{\Omega}^{\mathrm{T}}\, \mathbf{G}$. Then $\mathbf{G} = \mathbf{\Omega}\, \mathbf{X}$, and $\mathbf{\Omega} = \mathbf{G}\, \mathbf{X}^{-1}$. The inverse may be the more interesting. Numerically,

$$\mathbf{X}^{-1} = \begin{bmatrix} 1 & 0.3433 & 0.6442 \\ 0 & 1.0573 & 0.2706 \\ 0 & 0 & 1.1726 \end{bmatrix} \quad . \tag{D19}$$

We can then see that the first vector of $\boldsymbol{\Omega}$ is the same as the first vector of **G**. The second vector of $\boldsymbol{\Omega}$ is a combination of the first 2 vectors in **G**, the ones that depend only on red and green cones, and $\boldsymbol{\Omega}$'s third vector is a combination of all the Guth vectors. The same approach, beginning with Eq. (D17) could be used to find other relationships, such as $\boldsymbol{\Omega}$ in terms of $\bar{x}$, $\bar{y}$, $\bar{z}$. To emphasize individual functions, Eq. (D9) can be used for the unity operator.