Tutorial

Vectorial Color

Instructor:

JAMES A. WORTHEY, PE, PhD

Tuesday, 2008 November 11, 13:30-15:30

Color Imaging Conference 16
Portland, Oregon

Legacy Understanding

Fictitious but realistic color-matching data.

Cones, red, green and blue.

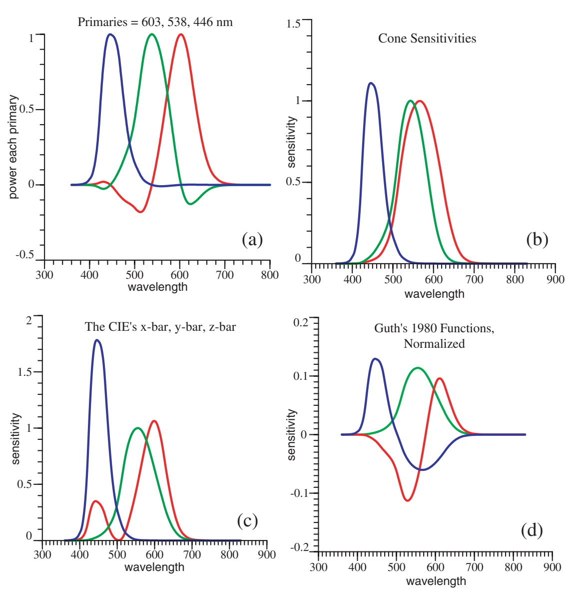

Traditional $\bar{x}$, $\bar{y}$, $\bar{z}$.

GUTH's 1980 opponent model, normalized. Achromatic function is the same as $\bar{y}$.

Linear transformations of color-matching data predict the same matches!

Figure 1. <The only numbered figure!>

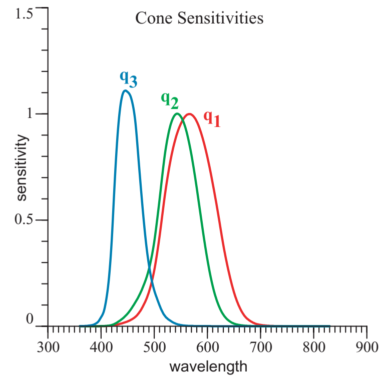

How does a set of 3 functions predict color matches?
A set of color-matching functions, CMFs, can be thought of as column vectors, which become the columns of a matrix A.

$$A = \begin{bmatrix} q_1 & q_2 & q_3 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \\ \vdots & \vdots & \vdots \end{bmatrix} \qquad (1)$$

A can contain any of the sets of color-matching functions in Fig. 1.

The spectral power distribution of light L_1 can be written as a column vector. It is then summarized by a tristimulus vector V,

$$V = A^{\mathrm{T}} L_1 \quad . \qquad (2)$$

Light L_1 is a color match for Light L_2 if

$$A^{\mathrm{T}} L_1 = A^{\mathrm{T}} L_2 \quad . \qquad (3)$$

If Eq. (3) holds for one set of CMFs shown above, then it will hold for the others. That much is in traditional sources such as Wyszecki and Stiles's book.

Moment of Reflection

Now look back to the 4 graphs of Fig. 1. They are equivalent for color matching, but play other roles:

- The cone functions relate to adaptation and constancy.
- The color matching data give clues about choosing primaries for additive mixing (emphasized by Bill Thornton).
- The opponent model fits the intuitive ideas that red and green are opposites, as are blue and yellow.

Other Issues at the Stage of Linear Additivity

What's missing is a good scheme for **vectorial addition of colors**. In his 1970 Book, Tom Cornsweet emphasized vector addition to account for overlapping receptor sensitivities:

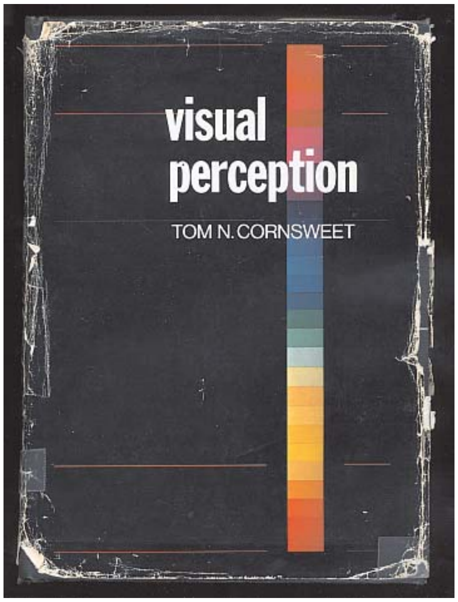

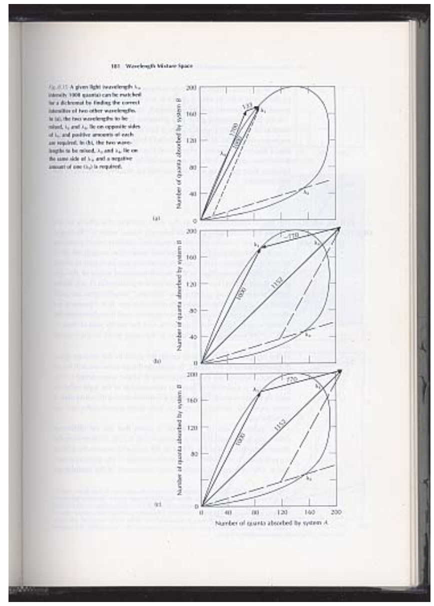

My method is inspired by Cornsweet, but also by Jozef Cohen.

Jozef Cohen's Innovations

Spectral reflectance can be measured at hundreds of points across the visible domain, but for most objects it graphs as a smooth curve:

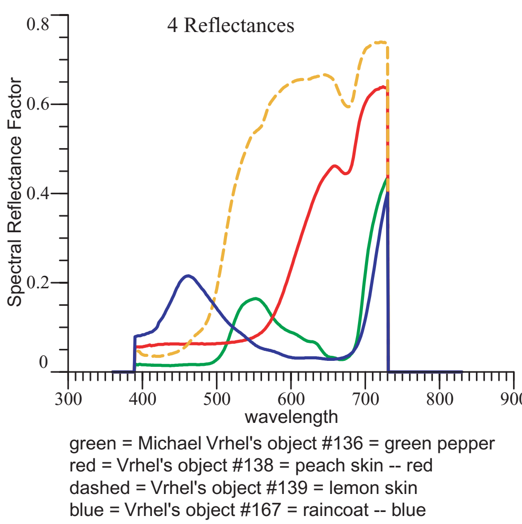

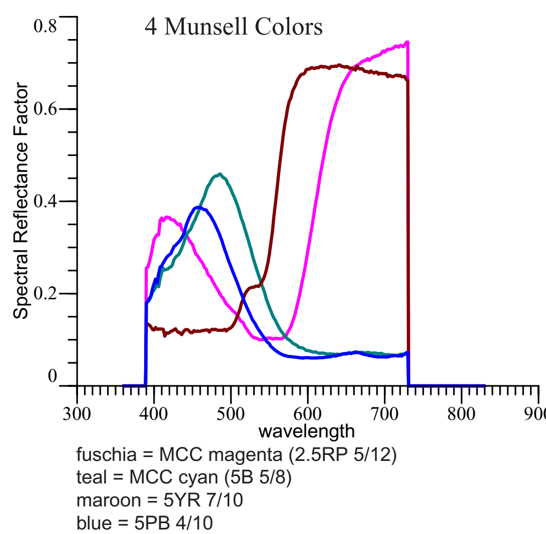

The smoothness of object colors is an underlying fact for *color constancy*, *object metamerism*, and *color rendering*. In 1964, at the dawn of the computer age, Jozef Cohen used principal components analysis to find that a population of Munsell Colors could be modeled as a sum of 3 or 4 basis functions. [Cohen, Jozef, "Dependency of the spectral reflectance curves of the Munsell color chips," *Psychonom. Sci*. 1, 369-370 (1964).]

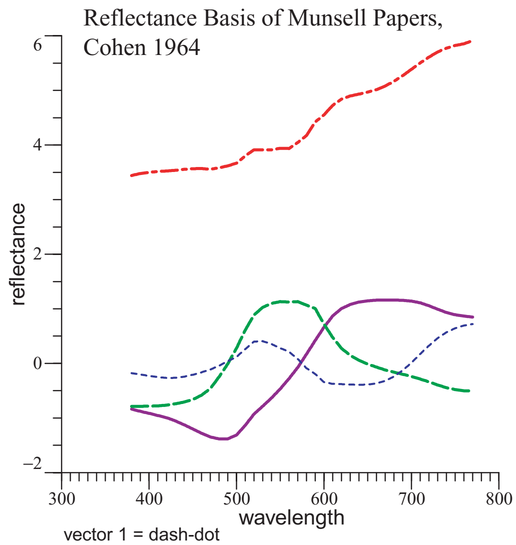

vector 1 = dash-dot
vector 2 = solid
vector 3 = long dashes
vector 4 = short dashes

The "Dependency" paper was little noted for 20 years, but is now widely cited. The idea is referred to as a "linear model," and is the basis for many studies with a statistical flavor. Having used some linear algebra ideas in a curve-fitting task, COHEN then asked, "Could linear algebra be applied directly in thinking about color-matching data?"

In the traditional presentation, color-matching data are transformed and one version is as good as another. COHEN asked a daring question: "Is there an invariant description for a colored light?"
We may well ask, "What would it even mean, that a function or vector is invariant?" COHEN started with WYSZECKI's idea of "metameric blacks." If 2 lights have different SPDs, but are colorimetrically the same, they are metamers. Subtracting one of them from the other gives a metameric black, a non-zero function with tristimulus vector = 0. Consider an example of D65 and a mixture of 3 LEDs adjusted to match:

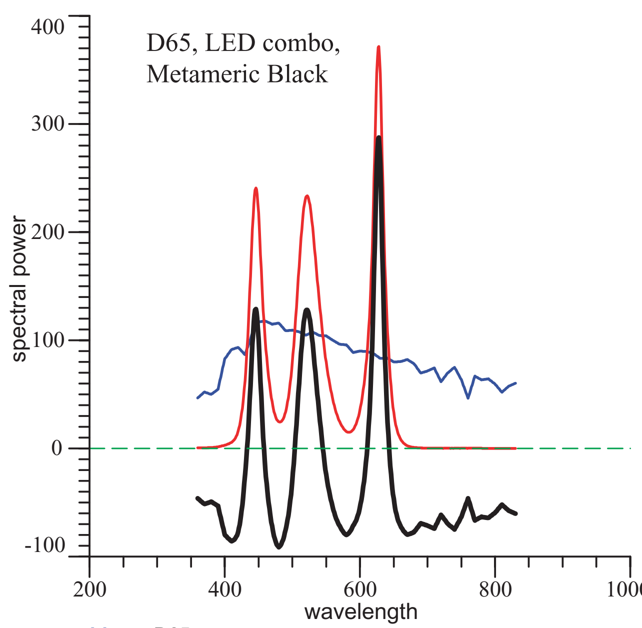

blue = D65
red = Combo of LEDs 86, 56, 28, meaning
HLMP-EH25-SV000 Red-orange
HLMP-CM15-S0000
Roitchner Lasertechnik LED 450-01U
Yxy(D65) = [1.057e+004, 03127, 0.3290]
Yxy(LED combo) = [1.057e+004, 03127, 0.3290]
black = metameric black, [X, Y, Z] = [0, 0, 0]

The figure at left shows the concept of metameric black. A mixture of 3 LEDs matches D65. Subtracting D65 from the LED combo gives the black graph, a function that crosses zero and is colorimetrically black. The metameric black is a component that you cannot see, as noticed by WYSZECKI.

COHEN then said, "Let's take a light (such as D65) and separate it into the component that you definitely can see, and a metameric black." (My wording.) The component that you definitely can see is the light projected into the vector space of the color matching functions. COHEN called that component...

... ta-dah ...

The Fundamental Metamer.

COHEN needed a formula to compute the fundamental metamer. As a practical matter, the "projection into the space of the CMFs" means a linear combination of CMFs that is a least-squares best fit to the initial light. COHEN was clever in what he did. As on p. 3, let A be a matrix whose columns are a set of CMFs,

$$A = \begin{bmatrix} q_1 & q_2 & q_3 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \\ \vdots & \vdots & \vdots \end{bmatrix} \quad . \qquad (1)$$

For a light $L(\lambda)$, we want to find the 3-vector of coefficients, $C_{3\times1}$, such that

$$L \approx AC \ . \qquad (4)$$

Where '≈' symbolizes that least-squares fit. Today, you might do a web search and you would learn to solve for C using the

$$\text{MOORE-PENROSE pseudo-inverse} = A^+ = (A^\mathrm{T}A)^{-1}A^\mathrm{T} \text{ , so that} \qquad (5)$$

$$C = A^+L = (A^\mathrm{T}A)^{-1}A^\mathrm{T} L \ . \qquad (6)$$

That would give you C as a numeric 3-vector, and if we define ...

$$\text{Fundamental Metamer of } L = L^* = \text{The least-squares fit to } L, \qquad (7)$$

then,

$$L^* = AC \ . \qquad (8)$$

A numerical value for 3-vector C gives a numerical vector for L^*, but no new insight. The value of C—the coefficients for a linear combination of CMFs—depends on which CMFs you are using. Nothing is invariant.

Deriving everything from a blank page, COHEN in effect combined Eq. (6) and Eq. (8) to find

$$L^* = [A(A^\mathrm{T}A)^{-1}A^\mathrm{T}]\, L \ . \qquad (9)$$

Then L is the original light, L^* is its fundamental metamer, and expression in square brackets is projection matrix **R**:

$$\mathbf{R} = A(A^\mathrm{T}A)^{-1}A^\mathrm{T} \ . \qquad (10)$$

COHEN noticed that **R** is invariant, and that led to other interesting ideas.

My algebra looks different from COHEN's. But the invariant projection matrix **R** remains important. **R** is invariant in the strongest possible sense. If A is any set of equivalent CMFs, like any of the sets in Fig. 1, then **R** is *the same large array of numbers*, not scaled or transformed in any way. For proof, see Q&A.

Fundamental Metamer Example

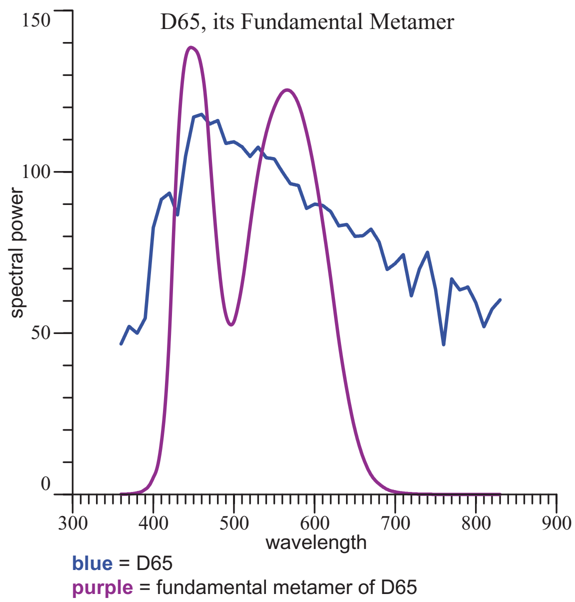

At left for example, the blue curve shows L = D65. The purple curve is L^*, the fundamental metamer of D65. The two curves are metamers in the ordinary sense. L^*, a linear combination of CMFs, is found by:

$$L^* = \mathbf{R}L \quad . \qquad (11)$$

Eqs (8) and (11) have '=' and not '≈' because L^* by definition *is* the least-squares approximation.

So, that's JOZEF COHEN's **Highly Original Contribution.**

Now on p. 5, we were talking about **vectors...**

COHEN References:

COHEN, JOZEF, "Dependency of the spectral reflectance curves of the Munsell color chips," *Psychonom. Sci.* **1**, 369-370 (1964).

COHEN, JOZEF, *Visual Color and Color Mixture: The Fundamental Color Space*, University of Illinois Press, Champaign, Illinois, 2001, 248 pp.

So, CORNSWEET used vectors to deal with overlapping CMFs (p. 5, above). Now, to be consistent with JOZEF COHEN's discoveries, and enjoy other benefits, we can make color vectors by starting with

Orthonormal Opponent Color Matching Functions.

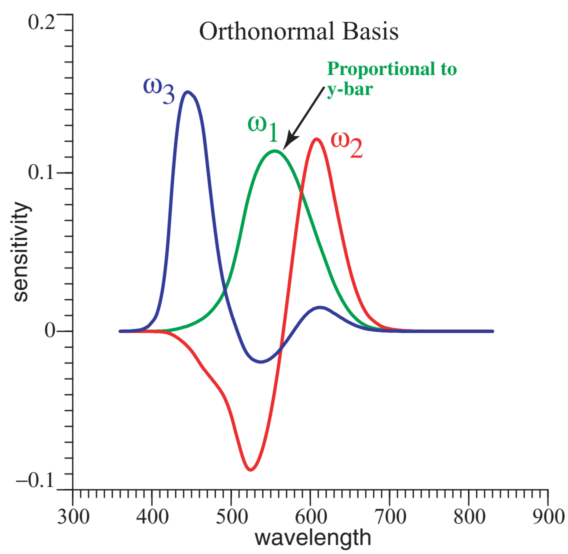

Functions are easy to generate, or can be found at http://www.jimworthey.com/orthobasis.txt .

Orthonormal Opponent Functions

$$\Omega = [\omega_1\ \omega_2\ \omega_3] \qquad (12)$$

1. ω_1 is achromatic sensitivity (whiteness). The function ω_1 is proportional to the usual y-bar, but normalized. That is $\langle\omega_1|\omega_1\rangle = 1$. ω_1 is a sum of (red cones) + (green cones), with appropriate coefficients. There is no blue input to ω_1 .
2. There is also no blue input to ω_2 . ω_2 is a difference, (red cones) – (green cones), such that it is orthogonal to ω_1 .
3. The third function, ω_3 , is the most messy. It has inputs from blue, red, and green cones.

Orthonormality

$$\langle\omega_i|\omega_j\rangle = \delta_{ij} \qquad (13)$$

meaning:

$$\langle\omega_i|\omega_j\rangle = 1 \text{ if } i = j \ ,$$
$$\langle\omega_i|\omega_j\rangle = 0 \text{ if } i \neq j \ .$$

Notation of bra, $\langle f|$, and ket, $|g\rangle$

A ket, $|g\rangle$ indicates a vector, usually a function of λ, written as a column matrix.

A bra, $\langle f|$ indicates a vector written as a row matrix.

Therefore, the one is the transpose of the other: $\langle f| = |f\rangle^T$.

Then by the ordinary rules of matrices, a complete bracket represents an inner product:

$$\langle f|g\rangle = \Sigma f_\lambda g_\lambda \ , \qquad (14)$$

where the matrix product effectively sums over the wavelength domain. For instance, the λ domain might be [360, 361, 362, ... 830], then vectors would have 471 elements.

Working Class Summary

Although many clever ideas (from COHEN, GUTH, THORNTON, BUCHSBAUM, etc.) guided its development, in the end the orthonormal basis is not radically new:

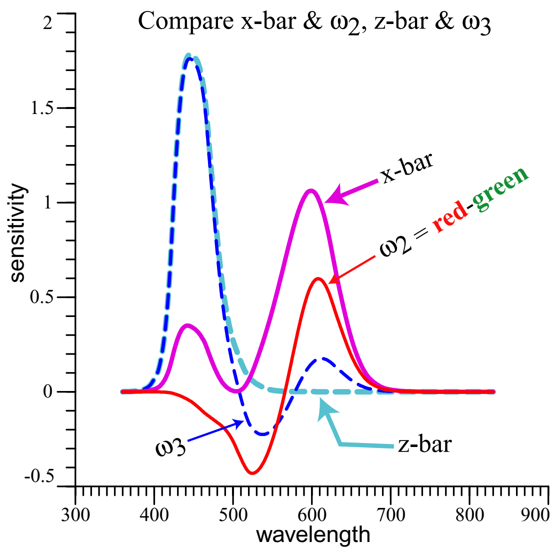

1. ω_1 is proportional to old-fashioned y-bar (not shown). So ω_1 is not new.

2. ω_3 is similar to the old z-bar, which is essentially a blue cone function.

3. In the XYZ system, a non-intuitive feature is x-bar, an arbitrary magenta primary. It is replaced in the new system by ω_2, a red-green opponent function. The opponent function confronts the overlap of cone sensitivities by finding a difference of red and green.

Let me say that again. An important feature of human vision is the overlapping sensitivities of the red and green cones. An opponent system confronts the issue of overlap. It remixes the red and green signals into ω_1 = red + green, and ω_2 = red – green.

Practicality

A new orthonormal basis could be created for any standard observer or camera. If it will suffice to have orthonormal opponent functions based on the 1931 2° Observer, then tabulated data are available at http://www.jimworthey.com/orthobasis.txt .

```
11
Orthonormal Color Matching Functions
James A. Worthey, Sat 2004 Oct 23, 20:57:46
1st column is wavelength in nanometers.
2nd column is omega-1 = achromatic
3rd column is omega-2 = red-green
4th column is omega-3 = blue-yellow
Put columns 2-4 in a big matrix called Omega.
If ‘ is matrix transpose, you can do this check:
Omega’*Omega = Identity Matrix (3x3) .
OTOH, Omega*Omega’ gives Matrix R, the projection matrix.
First line of this file is the number of these comment lines!
360       4.457966e-007     -2.2732411e-007     5.19311685e-005
361     5.00036628e-007    -2.69438911e-007     5.83360809e-005
362     5.61041793e-007     -3.1959058e-007     6.55531511e-005
363     6.29616395e-007    -3.78213858e-007     7.36777575e-005
364     7.06564849e-007    -4.45743615e-007     8.28052783e-005
365     7.92691682e-007     -5.2261486e-007      9.3031092e-005
366     8.89228099e-007    -6.10848068e-007      0.000104556575
367     9.97816844e-007    -7.12879535e-007      0.000117585251
368     1.11987975e-006    -8.29715601e-007      0.000132215511
369     1.25683798e-006    -9.62361805e-007      0.000148545746
370     1.41011485e-006    -1.11182622e-006      0.000166674347
371     1.58042235e-006    -1.27658478e-006      0.000186520506
372     1.77058528e-006     -1.4627818e-006      0.000208614278
373     1.98519588e-006    -1.68299053e-006      0.000233970598
374     2.22884074e-006    -1.94971689e-006      0.000263604238
375     2.50611211e-006    -2.27561136e-006      0.000298530207
376     2.82701851e-006    -2.68023553e-006      0.000340414724
377     3.19139606e-006    -3.15714084e-006      0.000388841422
378     3.58857044e-006     -3.6799146e-006      0.000441701398
379     4.00786849e-006    -4.22203069e-006      0.000496885676
380     4.43861818e-006    -4.75688934e-006      0.000552285297
```

etc, etc. Recall that graphs are on page 12.

Fun with Matrices

From p. 13,

$$\Omega = [\omega_1\ \omega_2\ \omega_3] \qquad (12)$$

$$\langle \omega_i | \omega_j \rangle = \delta_{ij} \quad . \qquad (13)$$

Orthonormality, Eq. (13) can be expressed in matrix form:

$$\Omega^T\Omega = \begin{bmatrix} \langle \omega_1 | \\ \langle \omega_2 | \\ \langle \omega_3 | \end{bmatrix} \begin{bmatrix} |\omega_1\rangle & |\omega_2\rangle & |\omega_3\rangle \end{bmatrix} = \begin{bmatrix} \langle \omega_1 | \omega_1 \rangle & \langle \omega_1 | \omega_2 \rangle & \langle \omega_1 | \omega_3 \rangle \\ \langle \omega_2 | \omega_1 \rangle & \langle \omega_2 | \omega_2 \rangle & \langle \omega_2 | \omega_3 \rangle \\ \langle \omega_3 | \omega_1 \rangle & \langle \omega_3 | \omega_2 \rangle & \langle \omega_3 | \omega_3 \rangle \end{bmatrix}$$

$$= \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} = I_{3\times 3}$$

$$\Omega^T\Omega = I_{3\times 3} \quad . \qquad (15)$$

If you multiply Ω and its transpose **in the reverse order**, you get a different interesting result:

$$\Omega\Omega^T = \mathbf{R} \quad . \qquad (16)$$

Eq. (16) is easily proved, but not right now. [Just substitute $A = \Omega$ in Eq. (10).]

Graphing a Vector

Now recall Eq. (2) above, and let $A = \Omega$. That is, let the set of CMFs be the orthonormal set. Then

$$V = \Omega^T |L_1\rangle. \quad (17)$$

The ket notation is a reminder that $|L_1\rangle$ is a column vector. Three-vector V is a tristimulus vector and can be graphed in 3 dimensions. Below is the vector for a 603 nm narrow-band light at unit power.

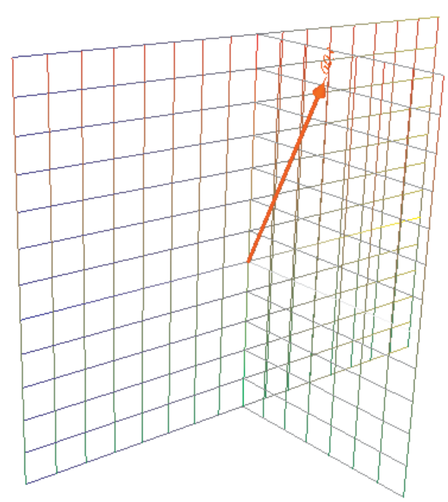

The vector is drawn from the origin. Its length depends on the optical power of the light and the eye's sensitivity at that wavelength.

The achromatic (= white) axis is the line from the origin that fades to white. Up-down is the red-green axis with red at the top. To the left from the origin is the V_3 or blue axis. The opposite of blue is yellow. The graph paper is colorized according to the color directions. The intuitive meaning of the opponent-color functions ω_1, ω_2, ω_3, leads to an intuitive graph.

5 Vectors

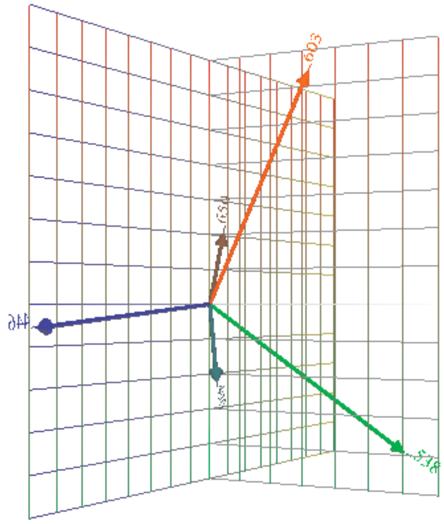

At left 5 colors are graphed:
THORNTON's Prime Colors, 603 nm, 538 nm, 446 nm, plotted at unit power, give long vectors. An extreme red, 650 nm, and a blue-green, 490 nm, also at unit power, give shorter vectors.

What is the significance of vector length?

The algebra of color vectors is specifically about additive color mixing. The mixture of lights is predicted by the sum of their vectors.

5 Vectors Added

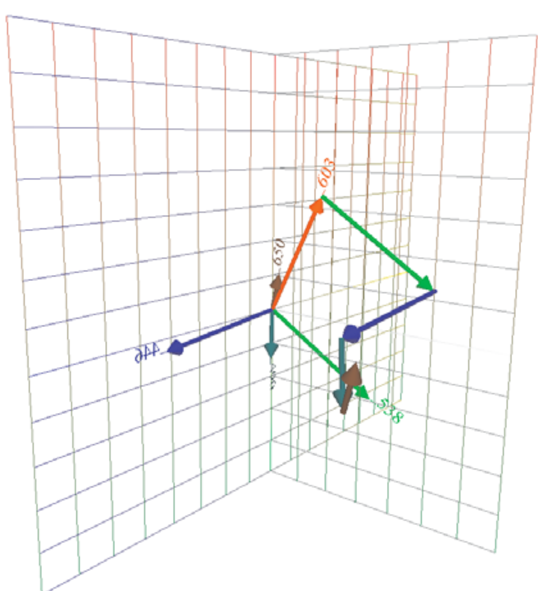

The same 5 vectors point from the origin, again for the 5 unit-power lights, at 603, 538, 446, 650, and 490 nm. Then the 5 vectors are added, tail to head. The resultant is a white vector.

So, vector length has a non-mysterious meaning:

According to the eye's overlapping sensitivities, and a light's wavelength and amplitude, a light maps to a vector with length and direction. Then vector addition predicts color matches.

Long Vectors

Vectors $[X\,Y\,Z]^T$ are traditionally not graphed. Now suppose that it is 1950 and we are trying to invent color television. It will be helpful if the phosphor colors correspond to **long vectors** in **well-separated directions**. Indeed, the NTSC phosphors are approximately at the longest-vector λs, which ≈ THORNTON's Prime Colors.

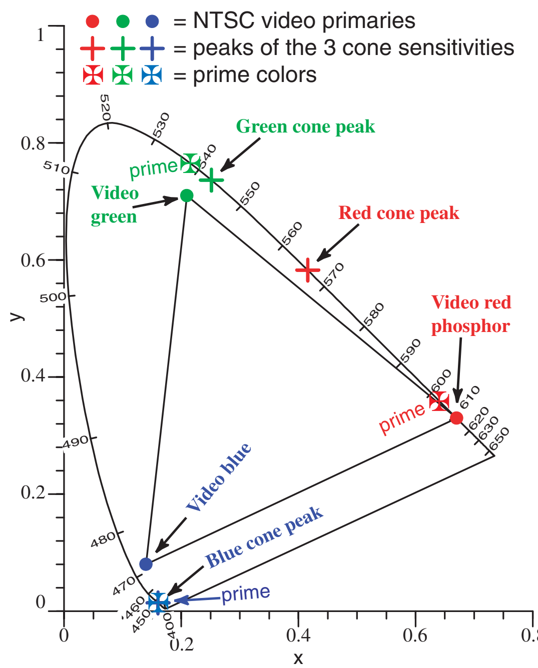

Wavelengths of Strong Action in Mixtures, nm.

2º Observer			
Longest vectors:	**445**	**536**	**604**
Prime colors =	**446**	**538**	**603**
10º Observer			
Longest vectors:	**445**	**535**	**600**
Prime colors =	**445**	**536**	**600**

The vectorial approach explains why the red phosphor is so far from the red cone peak. In fact, a 3-D vector diagram will lay out the facts for any choice-of-primaries problem.

Locus of Unit Monochromats

If λ of a narrow-band light is stepped through the spectrum and then a unit-power vector is plotted at each step, the vectors trace out a locus:

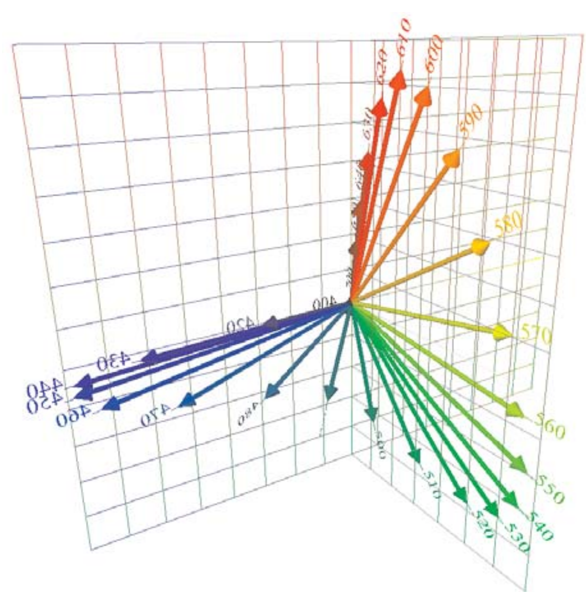

JOZEF COHEN referred to this curve in 3 dimensions as "**The Locus of Unit Monochromats**."

The curve is a spectrum **locus**. By **monochromat**, Cohen meant a narrow-band light.

The word **unit** means unit-power. Each monochromat has the same optical power, which gives meaning to the vector lengths. The "unit power" idea was one of COHEN's great leaps forward, because in the legacy system, a vector $[X\,Y\,Z]^T$ is found, and then we rush to plot it in the (x, y) diagram, losing the vector length. "Unit power" leads to a picture in which vector length has meaning.

Locus of Unit Monochromats (continued)

The Locus of Unit Monochromats (LUM) can be shown as a surface; then the locus is really the curve along the edge.

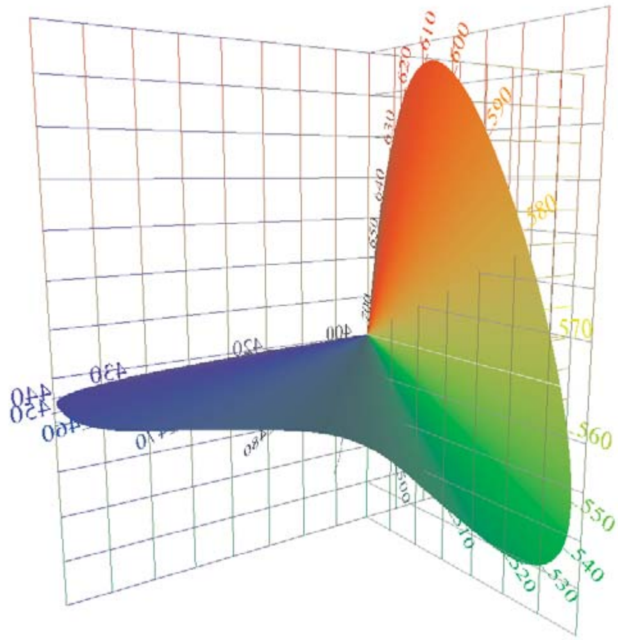

Jozef Cohen usually did not identify the LUM as a locus of *tristimulus* vectors. His vectors have more than 3 elements, and no preferred axes.

The LUM as Cohen described it has an invariant shape. The LUM based on the orthonormal opponent functions is the same invariant locus. Algebraic tinkering could rotate the LUM with respect to the axes, but its shape would not change (except for a potential mirror-image transformation).

To generate the LUM from the orthonormal basis, simply graph the rows of Ω.

5 Vectors Added (revisited)

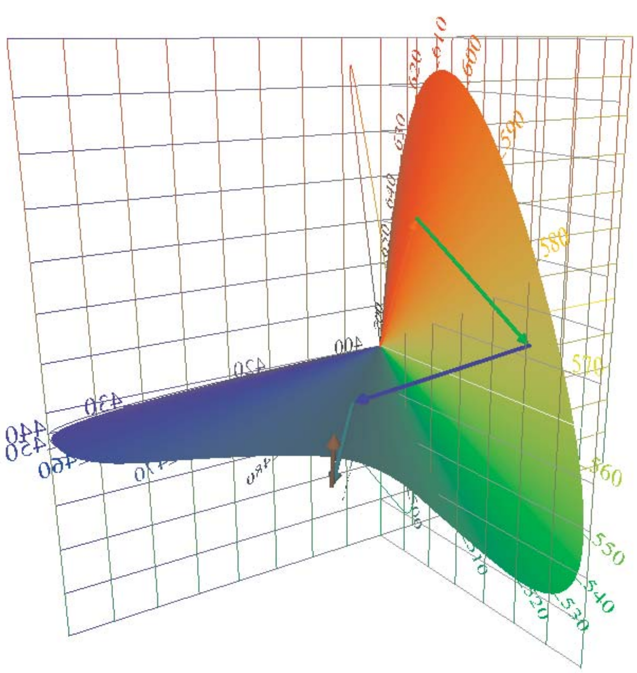

The Locus of Unit Monochromats can be included to give a frame of reference to a vector sum or other interesting diagram.

Composition of a White Light

The so-called “equal energy light” is one that has *constant power per unit wavelength* across the spectrum, indicated by a solid black line in the figure at the right.

The straight-line spectrum makes a simple discussion and is similar to a more realistic light, 5453 K blackbody (or 5500 if you like).

For the next step, let’s assume an equal-energy light that packs all of its power at the 10-nm points, 400, 410, and so forth, as indicated by the green dots.

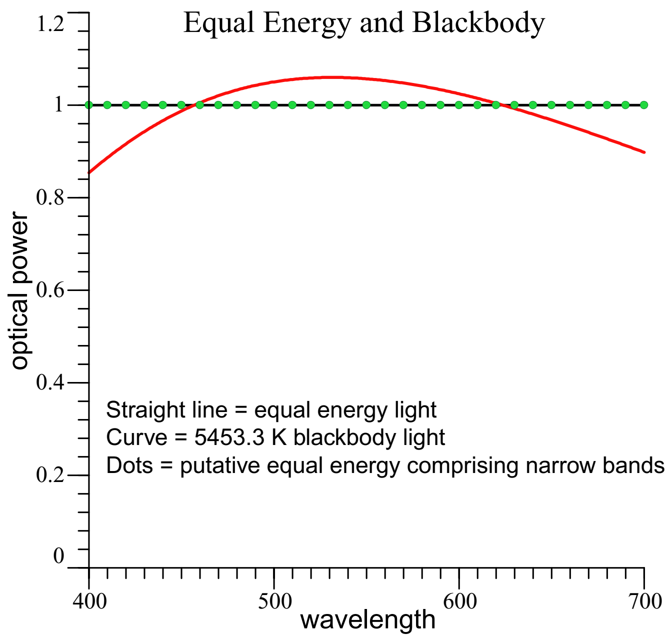

Composition of a White Light (continued)

At the right, the 31 narrow bands add up to make Equal Energy White. Each band contributes an arrow: blue, blue ..., blue-green, blue green ... green ... green-yellow ... orange ... red. Added tail to head, the arrows sum to a particular white.

The resultant of all that vector addition can be graphed as a single arrow from the origin. Conceptually, that arrow has 3 components, also shown as white, red, and blue arrows. The one arrow, or the set of 3, those symbolize the traditional approach, except it would be done in XYZ space.

The chain of 31 arrows forms a pattern that will repeat for most white lights. It progresses towards blue, then swings in the green direction, then back towards red.

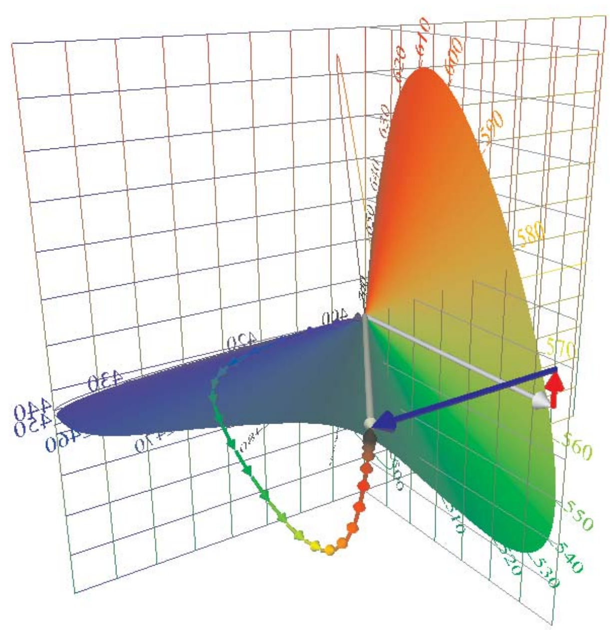

Composition of a White Light (continued)

The swing toward red almost cancels the swing towards green. The greenish and reddish vectors are all present in the light, but red cancels green, approximately. I have explained the same idea in the past without vectors.

Now let this white light be shined on colored objects, and think of what an object color does. A red object reflects a range of the reddest vectors, while absorbing green and blue. The light looks white, but it must have red in it to reveal the red object as different from gray or green.

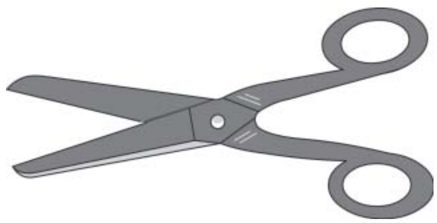

The most saturated colors will be those that reflect one or two segments of the spectrum and absorb the rest. You can think of those colors as “snipping out” a segment of the chain of vectors that compose the white light.

Is it a new idea to snip out part of the spectrum? No! The idea is called limit colors and is associated with MacAdam, Schrödinger and Brill, among others. On the next page is an illustration from MacAdam’s book, showing the possible ways to snip out parts of the spectrum. I won’t discuss limit colors at length, but they are a well-known idea.

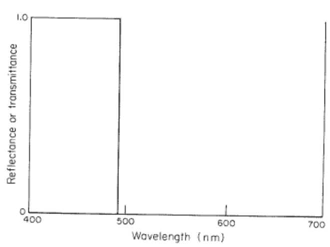

Fig. 7.17 Spectrophotometric curve of short-end type of optimal color

7.5 Maximum Visual Efficiency of Colored Materials 123

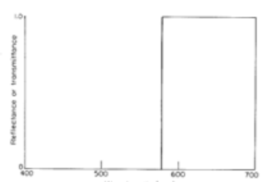

Fig. 7.18 Spectrophotometric curve of long-end type of optimal color

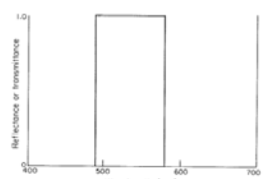

Fig. 7.19 Spectrophotometric curve of high-middle type of optimal color

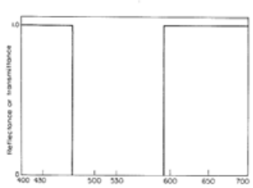

Fig. 7.20 Spectrophotometric curve of low-middle type of optimal color

Figures from D. L. MacAdam, *Color Measurement: Theme and Variations*, Springer, Berlin, 1981.

Observation: Many discussions of limit colors focus on a gamut of possible colors. Here we look at a more basic step: how a single limit color interacts with the components of a white light.

One Example: A Short-λ End “Optimal Color”

A hypothetical “short-end” color reflects all wavelengths up to 540 nm. Longer wavelengths are absorbed. The light source is the equal-energy light. Then the reflected light yields the vector chain seen at right. The entire equal-energy spectrum, as above, would sum to equal the one arrow, or the chain of three.

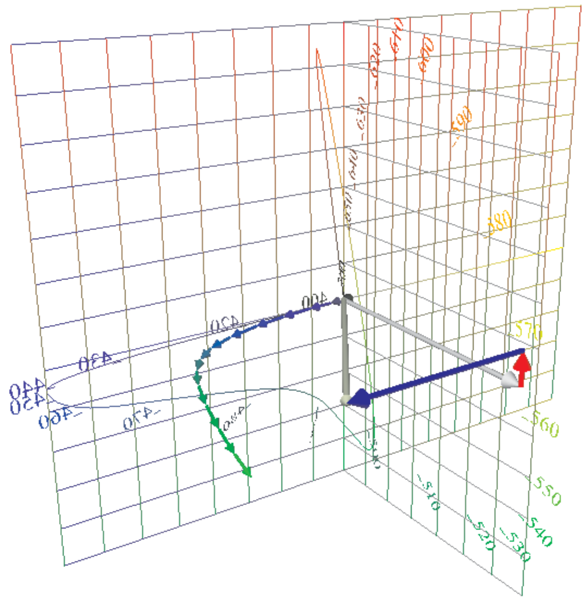

When the full chain of vectors for the equal-energy light is shown, you can picture pieces of the chain being snipped out according to MacAdam’s 4 types of optimal color:

- **short-end**
- **long-end**
- **middle**
- **low-middle**—meaning that short and long λ’s are reflected, but a band in the middle is not reflected.

We’ll revisit Composition of White Lights after a digression to more general issues.

Why So Many 3-Dimensional Drawings?

Why do I use so darned many of these 3-dimensional drawings? For detailed analysis and design, one might see a way to capture key information in a flat 2-D drawing. But the 3-D drawings have a

Special Charm.

If a person views some objects under a white light, then the eye and the light operate together. The eye and the light are each constants over the scene.

- The chain of vectors in 3-D preserves all the information that colorimetry can give.
- No assumption enters the chain of vectors in 3D. There are no preferred paint chips or anything like that.
- Almost no simplification enters the drawings. If we stick to 10-nm steps, some tiny wiggles are lost, but conceptually—and in computer calculations—the smallest details can be used.
- In short, the 3D picture makes no pre-judgment about what is important.
- We can find even more uses for Cohen’s space.

Here is another use of the orthonormal opponent functions ...

Minimizing Correlation

One benefit of the orthonormal basis is that v_1, v_2, v_3 are less correlated than other sets of tristimulus values.

- In metrology, it is easier to predict propagation of errors if the starting variables are less correlated.
- In information theory, if variable p is known, and q can be estimated from p, then q carries less information. The NTSC television standard uses opponent colors as a step in image compression, so the idea is not new.

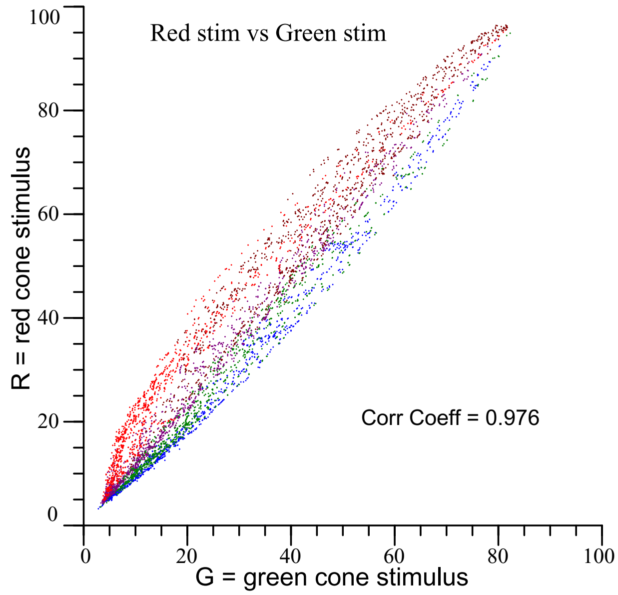

Spectral reflectances are available for a set of 5572 paint chips: Antonio Garcia-Beltran, Juan L. Nieves, Javier Hernandez-Andres, Javier Romero, “Linear Bases for Spectral Reflectance Functions of Acrylic Paints,” *Color Res. Appl.* **23**(1):39-45, February 1998. Professor Garcia kindly sent me the raw data.

To make realistic visual stimuli, let the paint chips be illuminated by D65. Then one stimulus can be plotted against another to see if they are correlated or independent. Garcia et al. put the chips in color groups, so I’ve colored the dots accordingly. The graph above shows that red and green cone stimuli are highly correlated, correlation coefficient = 0.976. X vs Y in the legacy system is a little better, while the opponent stimuli v_1 and v_2 are most independent, Correlation Coefficient = 0.180 .

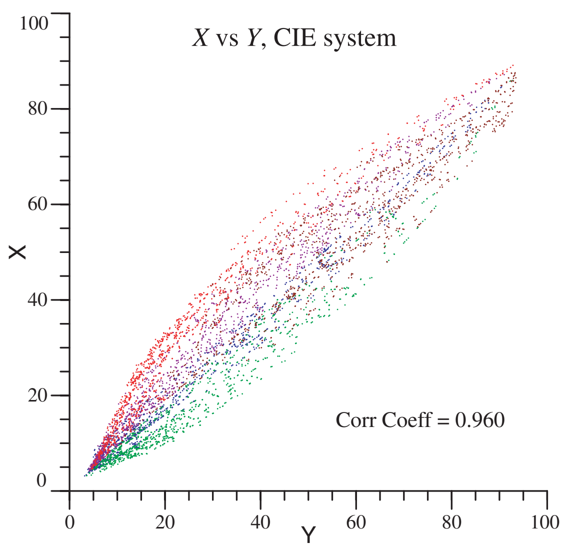

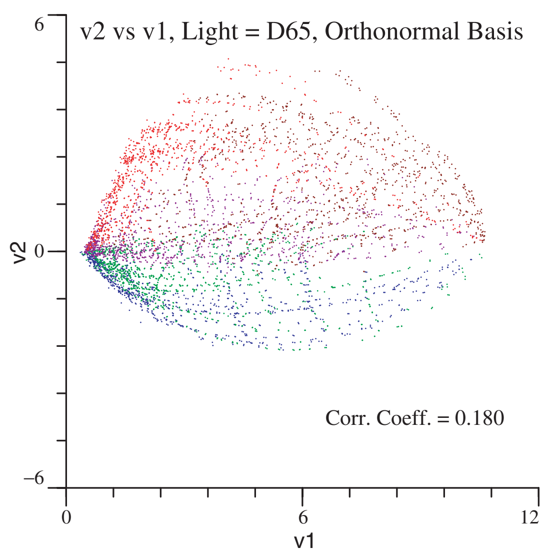

Relationships Among Equivalent Sets of CMFs

The graphs just presented illustrate some advantages of the opponent functions. Equations were avoided, but below the surface was the idea that transformed CMFs are equivalent for color matching, but not equivalent for other purposes.

Now consider another application whose charm is more mathematical. Referring to Fig. 1, we recall that a transformed set of CMFs can be interesting and look different from another set that's equivalent-for-matching. Any transformation of CMFs can be represented as a 3×3 matrix. *If the starting functions are the orthonormal ones*, then the transformation matrix has a simple form. Let C be any set of CMFs known to be linear combinations of Ω: $C = [|c_1\rangle\, |c_2\rangle\, |c_3\rangle]$. As always, $\Omega = [|\omega_1\rangle\, |\omega_2\rangle\, |\omega_3\rangle]$. Then the sets are related by $C = \Omega\, T$, where

$$T = \begin{bmatrix} \langle \omega_1 | c_1 \rangle & \langle \omega_1 | c_2 \rangle & \langle \omega_1 | c_3 \rangle \\ \langle \omega_2 | c_1 \rangle & \langle \omega_2 | c_2 \rangle & \langle \omega_2 | c_3 \rangle \\ \langle \omega_3 | c_1 \rangle & \langle \omega_3 | c_2 \rangle & \langle \omega_3 | c_3 \rangle \end{bmatrix} \quad . \qquad (18)$$

Notice that each matrix element is a tristimulus value for a function $|c_j\rangle$, and in fact each *column* is the tristimulus *vector* for one of the $|c_j\rangle$. The 3 vectors can be plotted in COHEN's space.

CMFs Graphed as if they were Lights

Thanks to orthonormality, the functions $|\omega_j\rangle$ plot to the v_1, v_2, v_3 axes. (**Not** emphasized in the drawing!) Legacy function y-bar also plots to the $+v_1$ axis. Red and green cones plot near to each other and to $+v_1$. Also, *r* cones and *g* cones lie in the v_1–v_2 plane. Blue cones and z-bar are the same direction and quite close to $+v_3$. x-bar is **not** in the v_1–v_2 plane, or on the spectrum locus. In short, the graph shows similarities and differences among functions.

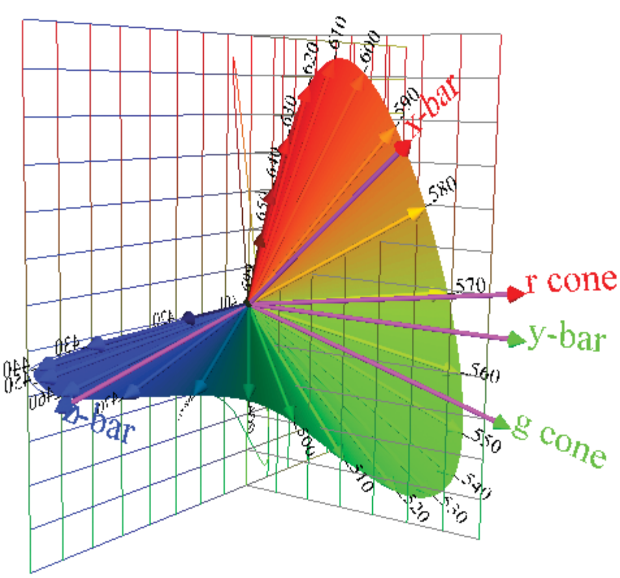

Simplicity would be lost if a similar plot were made with axes *X*, *Y*, *Z*. Then, for example the y-bar function would not plot to the *Y* axis, since y-bar has non-zero inner products with x-bar and z-bar.

Table 1. Correlation coefficient and direction cosine for various pairings of functions. The correlation coefficients are based on D65 and the paint samples as in Figs. 1-3. Direction cosine is a measure of overlap between the sensitivity functions.		
Functions compared	**Correlation Coefficient**	**Direction Cosine**
R cones, G cones	0.976	0.918
ω_1, ω_2	0.180	0
$\bar{x}$, $\bar{y}$	0.960	0.760
R cones, B cones	0.520	0.058
G cones, B cones	0.619	0.121
$\bar{x}$, $\bar{z}$	0.548	0.255
$\bar{y}$, $\bar{z}$	0.557	0.082
ω_1, ω_3	0.522	0
ω_2, ω_3	−0.303	0

The orthonormal pairs, with cosine = 0, give the lowest correlation of the paint chips, but not zero correlation. With other populations of object colors, correlations may vary.

Opponent Colors and Information Transmission

Above I've shown the benefit of an opponent system for information transmission (image compression) in an empirical way, by starting with thousands of paint chips. BUCHSBAUM and GOTTSCHALK proceeded differently, deriving an opponent-color system to optimize information transmission. They started with cone functions that are similar to the ones I use, and got a result like my opponent basis. That is, they got an achromatic function (not shown) similar to y-bar, and the red-green and blue-yellow functions shown at right. The achromatic and red-green functions have minimal blue input, and the blue-yellow function crosses the abscissa twice and looks very much like the ω_3 function that I have been using.

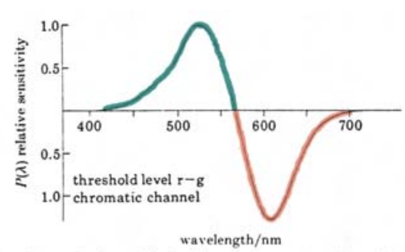

FIGURE 4. First chromatic channel derived by eigenvector transformation of threshold level Vos–Walraven primaries, corresponding to a r–g opponent channel.

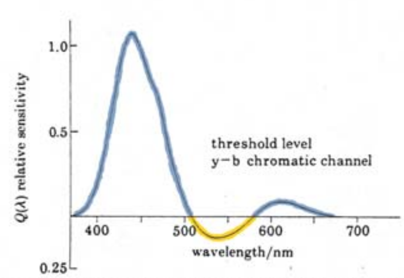

FIGURE 5. Second chromatic channel derived by eigenvector transformation of threshold level Vos–Walraven primaries. The spectral shape of this channel is discussed in the text.

This equation, scanned right out of the original article shows how their first two functions have little blue input, etc:

$$\begin{bmatrix} A \\ P \\ Q \end{bmatrix} = \begin{bmatrix} 0.887 & 0.461 & 0.0009 \\ -0.46 & 0.88 & 0.01 \\ 0.004 & -0.01 & 0.99 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}. \qquad (12a) \qquad (19)$$

Their functions *A*, *P*, *Q* are orthogonal, but not normalized.

Long story short, Buchsbaum and Gottschalk started with cone functions and derived opponent functions. They had a specific goal of optimum information transmission. I derived opponent functions in a simple way and discovered their connection to Cohen's work. In the end, the sets of functions are extremely similar, confirming that the opponent basis is appropriate for image compression and propagation-of-errors. All orthonormal sets based on the same cones lead to the same Locus of Unit Monochromats, so in that sense Buchsbaum's set cannot be wrong.

[Buchsbaum, G. and A. Gottschalk, "Trichromacy, opponent colours coding and optimum colour information transmission in the retina," *Proc. R. Soc. Lond.* **B 220**, 89-113 (1983).]

Value of Cohen's Space and the Orthonormal Basis

Benefits of the orthonormal basis can be listed:

- Makes nice spread-out vector diagrams. [Greatest separations possible.]
- Opponent feature confronts the fact of red-green overlap.
- Opponent feature gives meaning to the axes: whiteness, red-green, blue-yellow.
- Gives decorrelated stimuli, good for image compression, or propagation of errors.
- Orthonormal functions lead to simple formulas, such as Eq. (15) - (16). [Also see Q&A item on "Fun with Orthonormal Functions."]
- Using the orthonormal basis, and treating CMFs as lights, ω_1 plots to the v_1 axis, and so forth. By comparison, a non-orthogonal basis would lead to a crazy graph where, e. g. $\bar{x}$ would **not** plot to the *X* axis.

Cohen's Insight; the Fundamental Metamer

The last list item is the most important and echos Jozef Cohen's ideas. If you have two lights (2 SPDs) as stimuli to vision, **they have an intrinsic vector relationship**. Each light has an invariant fundamental metamer, which is its projection into the vector space of the CMFs. Tristimulus vectors *V* obtained using Ω have the same amplitudes and vector relationships as fundamental metamers. Functions $\bar{x}$ and $\bar{y}$ are intrinsically 40.5° apart, but ω_1, ω_2, ω_3 are 90° apart, so v_1, v_2, v_3 are more appropriate axes than *X*, *Y*, *Z*.

End Digression, Return to

Vector Composition of Lights

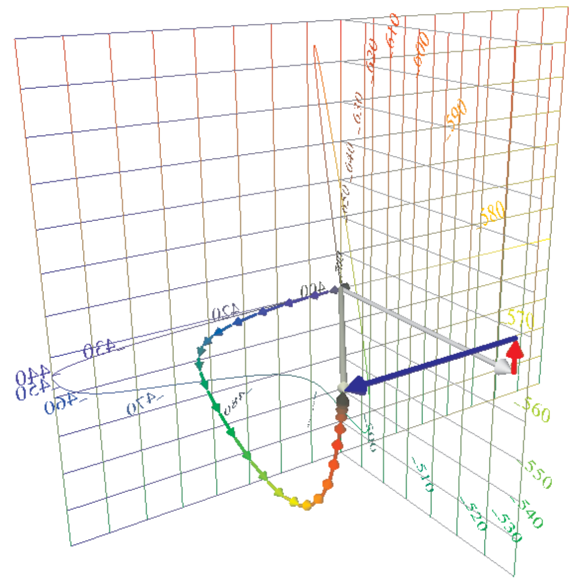

The figure at the right once again shows the equal energy light as a vector sum. Power is assumed to be in narrow bands at 10 nm steps,

400, 410, ... , 700 nm .

Although the power is the same in all bands, they map to vectors of varying amplitude and direction, according to the facts of color matching. As a practical matter, the little vectors are based on certain rows of matrix Ω.

A key issue for lighting and color is that the chain of arrows swings toward green and back towards red. If this light shines on a white paper, the green and red almost cancel, but if it shines on a red pepper,  , the red pigment selects a group of red vectors, by absorbing green ones. Then it matters how much red was in the light.

Some Lights Have Less Red and Less Green

A white light has net redness or greenness that is small or zero. The same white point can be reached by different lights in different ways.

SPDs of two lights are plotted at right. The black line is for a High Pressure Mercury Vapor light, while the blue is JMW Daylight, adjusted to have the same tristimulus vector. (Yes, that means they are matched for illuminance and chromaticity.) The wavelength domain is chopped according to the dashed vertical lines. The wavelength bands are 10 nm, except at the ends of the spectrum, with most bands centered at multiples of 10 nm. If one light then multiplies the columns of Ω, those products could be graphed as a distorted LUM, but we skip that step. Instead, the color composition of each light will be graphed as a chain of vectors.

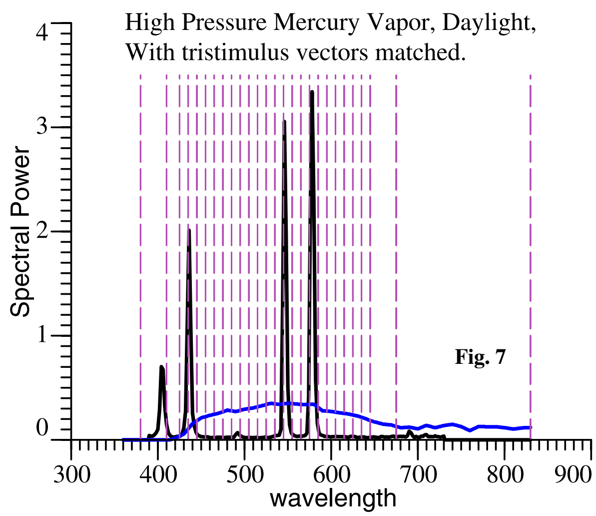

Fig. 7

Vectors will be drawn according to indicated wavelength bins.

Comparing Color Composition of Lights

Now the same two lights are compared in their vector composition. The smooth chain of thin arrows shows the composition of daylight. Slightly thicker arrows show the mercury light.

The mercury light radiates most of its power in a few narrow bands, leading to a few long arrows that leap toward the final white point. Compared to “natural daylight,” the mercury light makes a smaller swing towards green, and a smaller swing back towards red. Such a light would leave the red pepper starved for red light with which to express its redness.

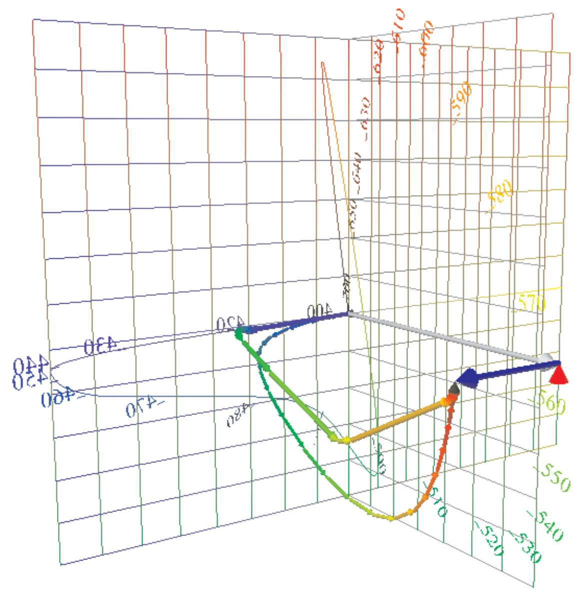

Comparing Lights (continued)

Above, two lights are compared by the narrow-band components of their tristimulus vectors. At right the same comparison is shown, but projected into the *v*1-*v*2 plane. The loss of red-green contrast is the main issue with lights of “poor color rendering,” and that shows up in this flat graph. If you were really designing lights, you might use the *v*1-*v*2 projection as a main tool. You might want to add wavelength labels to the vectors. (In the VRML picture, λ info is indicated by coloration. In the picture above, vectors that are approximately parallel pertain to the same band.)

You are not limited to a certain projection, or to any step that loses information.

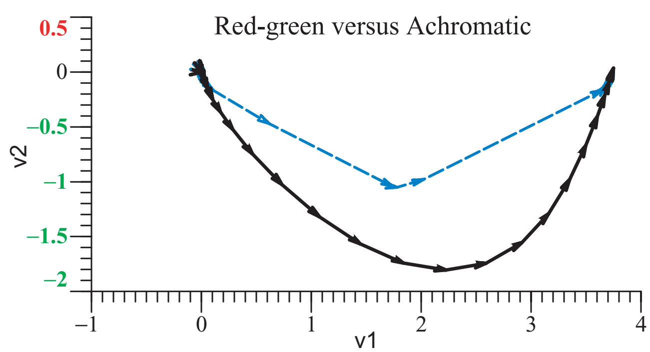

dashed = high pressure mercury vapor
solid = JMW daylight
Tristimulus vectors of the lights are equal.

Comparing Lights (continued)

Other interesting data can be plotted in COHEN's space. Suppose that the 64 Munsell chips from VRHEL et al. are illuminated first by daylight and then by the mercury light. Since the mercury light lacks red and green, we expect it to create a general loss of red-green contrast among the 64 chips.

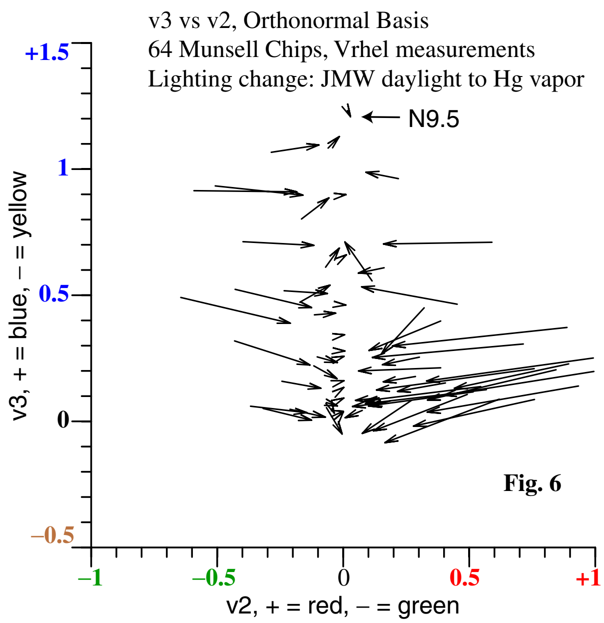

Fig. 6

The graph at right is a projection into the v_2-v_3 plane. Each arrow tail is the tristimulus vector of a paint chip under daylight. The arrowhead is the same chip under the mercury light. The lightest neutral paper is N9.5, and is a proxy for the lights. Notice that 3-vectors projected into a plane still have the properties of vectors in the 2D plane. As expected, red and green paint chips suffer a tremendous crash towards neutral. Actual neutral papers appear as arrows of zero length. Each point is calculated from detailed spectral data.

Comparing Lights (continued)

For more details on this and other comparisons of lights, please see graphic materials for "How White Light Works," at:
http://www.jimworthey.com/jimtalk2006feb.html
. Or, for various links, see the CIC 16 Tutorial Page:
http://www.jimworthey.com/qna/tutorial_cic16.html
.

Cameras and the "Luther Criterion"

The **Luther Criterion**, also known as the Maxwell-Ives Criterion:

For color fidelity a camera's spectral sensitivities must be linear combinations of those for the eye.

- For instance, in Fig. 1 on p. 2, if one of the graphs represented a camera, it would meet the criterion.
- On the other hand, it is less obvious how to describe a camera that departs from the ideal.
- Some find a figure of merit, which disposes of the issue, but teaches little.
- At CIC 14, Brill and I related the Luther Criterion to the Locus of Unit Monochromats, the LUM. We also brought Thornton's prime colors into the picture and developed the ideas in many steps.
- However, the poster itself—available on my web site—presents a concise method for the actual calculation, "The Fit First Method."
- Now, assuming that you appreciate the Locus of Unit Monochromats, I'll explain the Fit First Method in **a few bold steps**.

Bold Steps For Camera Analysis

- The eventual goal is to simulate the original scene. The Luther Criterion is a main issue, not incidental.
- Color vision depends on the 3 cone sensitivities (or other color matching functions) *taken together as a system*.
- The Locus of Unit Monochromats is an invariant 3D graph that summarizes color mixing by humans. It accounts for the 3 cone types as a system. (Any observer, but I use 2° Standard Observer.)
- The camera has its own LUM. If the Luther Criterion is met, it will be the same as the human LUM.
- By comparing the human and camera LUMs, we can relate the camera's color-mixing properties to those of human, in detail.
- There is one stumbling block. The LUMs may be similar or identical as shapes floating in space. To compare them, we need to rotate them into good alignment. The Fit First Method creates good alignment.

Compare Camera LUM to that of Human

At right is a screen grab from the virtual reality comparison of a Dalsa 575 sensor's LUM to the 2° human observer LUM. The spheres represent the LUM of the camera. The arrowheads trace the best fit to human LUM by the camera functions.

Please try to see the big picture for a moment, letting go of the details. The human LUM is an invariant representation of trichromatic color vision. The camera has its own LUM. We want to compare them, but must somehow position them for a clear comparison. The **Fit First Method** finds the camera LUM *in a good alignment*.

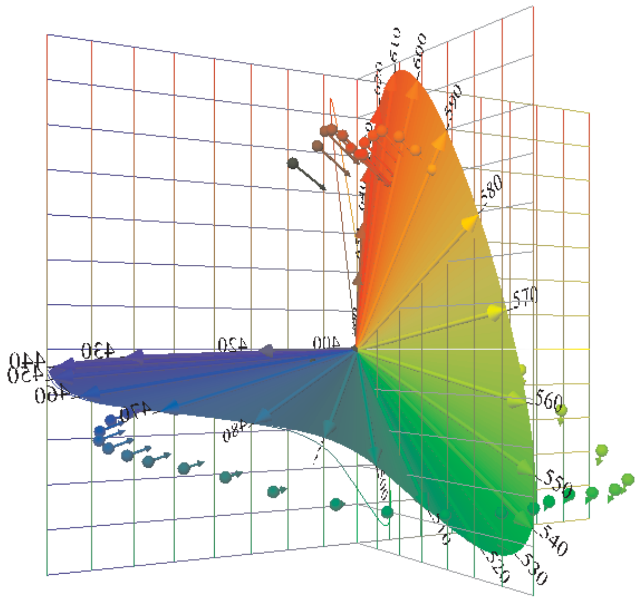

The Fit First Method

Conceptually, the camera's LUM (spheres) is more fundamental than the fit to the human LUM (arrowheads). The trick of the Fit First Method is to find the best fit first, then find the LUM from that.

Here is the computer code:

```
Rcam = RCohen(rgbSens) # 1
CamTemp = Rcam*OrthoBasis # 2
GramSchmidt(CamTemp, CamOmega) # 3
```

The camera's 3 λ sensitivities are stored as the columns of array rgbSens . Because of the invariance of projection matrix **R**, it doesn't matter how the functions are normalized, or whether they are actually in sequence r, g, b. Statement **1** finds Rcam, the projection matrix **R** for the camera. `RCohen()` is a small function, but conceptually,

```
RCohen(A) = A*inv(A'*A)*A'
```
(20)

.

In other words, step **1** applies Cohen's formula for the projection matrix. Then `Rcam` is the projection matrix for the camera. In step **2**, the columns of `OrthoBasis` are the human orthonormal basis, Ω . The matrix product `Rcam*OrthoBasis` finds the projections of the human basis into the vector space of the camera. **But**, the wording about projection is another way of saying that step **2** finds the best fit to each ω_i by a linear combination of the camera functions. So, step **2** is the “fit” step.

Step **2** does 3 fits at once, but let’s look at just one. At right, the dashed curve is ω_1, the human achromatic function. The camera in question happens to be a Nikon D1. The solid curve is a linear combination of that camera’s *r*, *g*, and *b* functions that is the least-squares best fit to ω_1. There would be other ways to solve the curve-fitting problem, but projection matrix **R** is convenient. A best fit is found for each ω_j separately. The resulting re-mixed camera functions are not an orthonormal set.

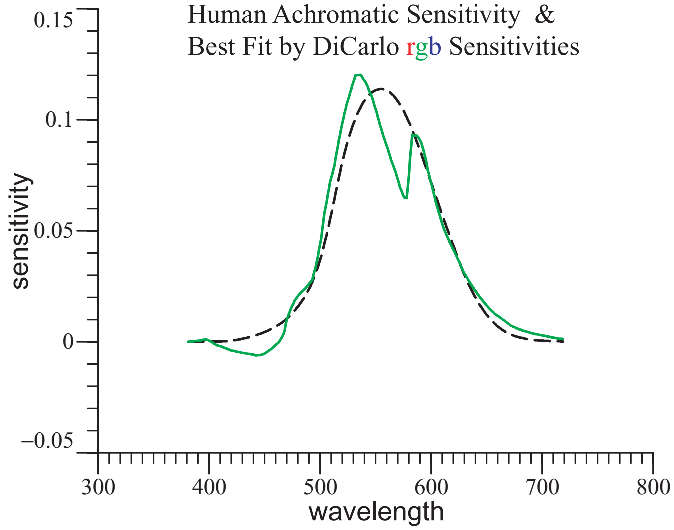

Ref: DiCarlo, Montgomery & Trovinger, CIC 12.
dashed = human achromatic = ω1.
solid = approx using r, g, & b sensors of Nikon D1.

Step 3, Orthonormalize the Re-mixed Camera Functions

Since the re-mixed camera functions are computed separately, they are not orthonormal and would not combine to map out a true Locus of Unit Monochromats. But they mimic Ω and are in the right sequence. We need to make them orthonormal, which is what the Gram-Schmidt method does, Step **3**.

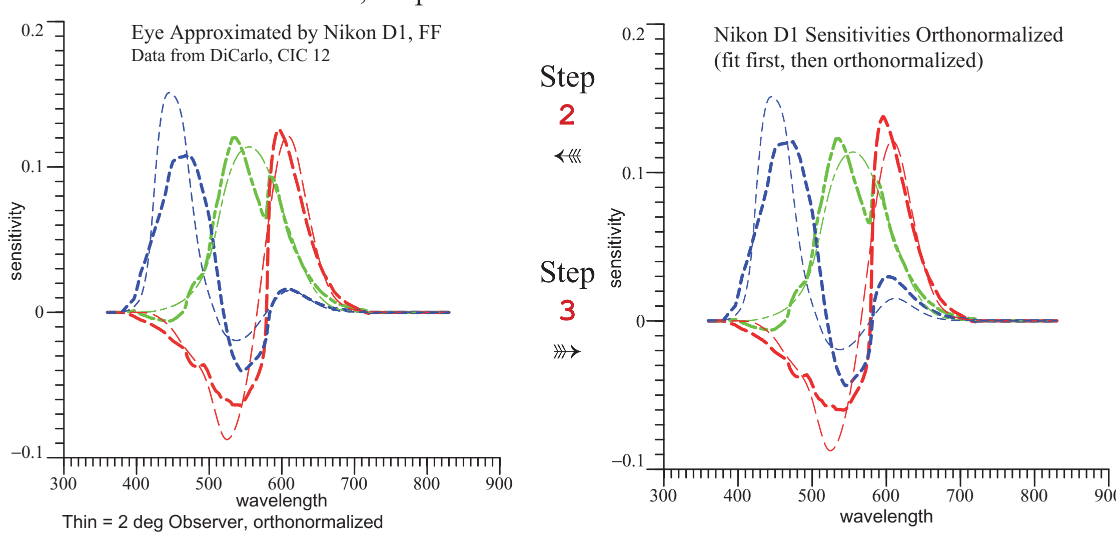

Thin = 2 deg Observer, orthonormalized
Thicker = 2 deg Observer approximated by Nikon D1
long-short = ω1, long dashes = ω2, short dashes = ω3
Fit first, then orthonormalized, then fit again.
The second fit should be the same as the first, but
for orderly software, the fit is repeated.

thicker curves = Nikon D1 sensitivities, FF
thinner = 2 degree observer
long-short = ω1, long dashes = ω2, short dashes = ω3
Camera sens. fit first, then orthonormalized.
Ref:DiCarlo, CIC 12

The two sets of graphs above look similar. But the one on the left shows the set of "fit" functions. The one on the right shows the orthonormal basis of the Nikon D1. The thinner curves pertain to the 2° observer, the thicker ones to the camera.

Why Does it Matter?

When you have the orthonormal basis, for the eye or for a camera, **you can do many things with it**. Combining the 3 functions generates the Locus of Unit Monochromats. The orthonormal property leads to some simple derivations. On the Q&A page, see "Can we have fun with orthonormal functions?"

The Noise Issue

To extract color information, some subtraction must be done. The signals subtract but the noise adds (in quadrature). The noise discussion becomes more concrete when one can say exactly what subtraction will be done. The camera's orthonormal basis is a natural to be the canonical transformed sensitivity. Orthogonality means "no redundancy" and normalization standardizes the amplitude. There's a numerical noise example worked out in the CIC 14 article. For now, the point is that expressing camera sensitivities as an orthonormal basis is a giant step towards dealing with noise.

Camera Example, Nikon D1

The graphs on p. 50 pertained to the Nikon D1, based on data from CIC 12. Here are the camera's red, green, and blue sensitivities:

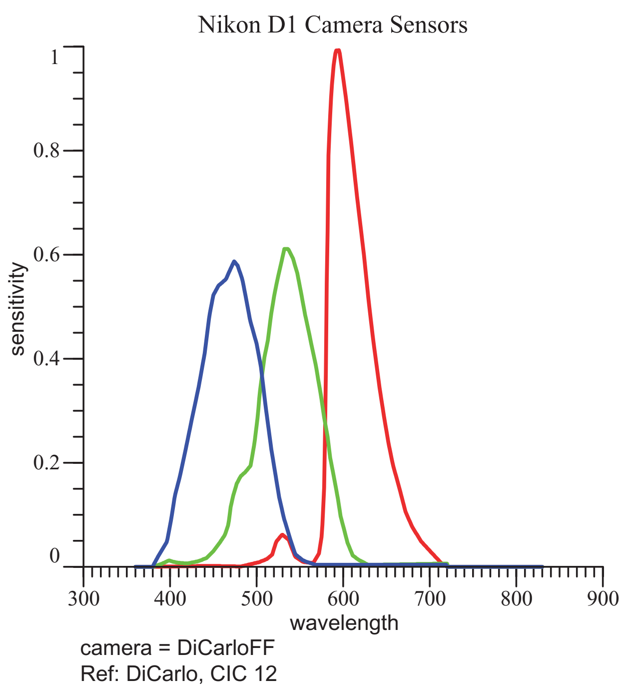

The camera's LUM can be compared to the eye's. Rather than another perspective picture, we now view the LUMs in orthographic projection. The dashed curves are the human locus. The solid curves are the camera's locus, while the tips of the small green arrows are points on the best-fit sensitivity function.

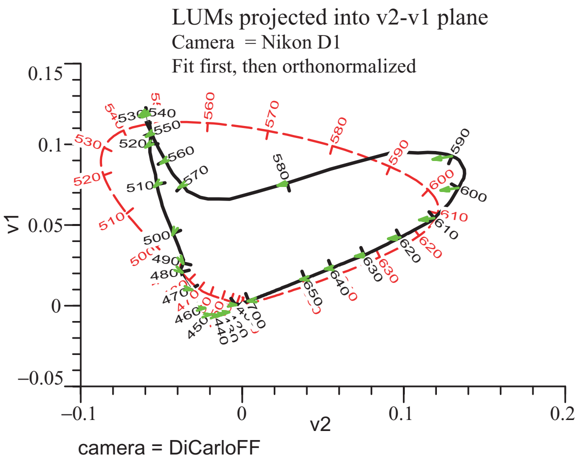

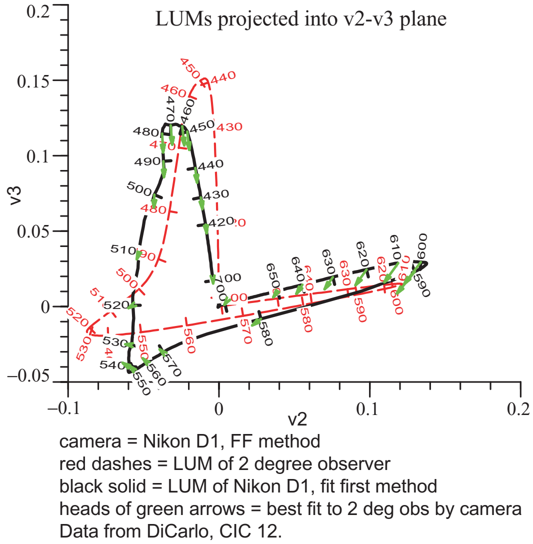

Now you may say "These curves mean nothing to me!" That may be at first, but the graphs contain a complete description of the camera sensor, with no hidden assumptions, and few details lost.

Finding Some Meaning in the Camera's LUM

Consider the left-hand graph, "LUMs projected into v_2-v_1 plane." Only the red and green receptors contribute to the human LUM in this view, and v_1 is the achromatic dimension for human, based on good old y-bar. In this plane at least, the particular camera tends to confuse wavlengths in the interval 510 to 560 nm, which are nicely spread out as stimuli for human. Yellows, say 560 to 580 nm, have lower whiteness than they would for human. The camera has other differences from human that may be harder to verbalize. To the extent that finished photos look wrong, one could revisit these graphs for insight.

More Examples

Five detailed examples were prepared in 2006, and they are linked from the further examples page:
http://www.jimworthey.com/furtherCamDesignLUM.html . For instance, QUAN's optimal sensor set indeed looks good in any of the graphical comparisons to 2º observer. The smooth and overlapping sensitivities of "Foveon X3 without prefilter" allow it to discriminate wavelengths all through the spectrum.

Something Completely Different: a 4-band Array

Sony publishes a specification for a 4-band sensor array, the ICX429AKL. I'm not sure of the intended application, but it could potentially be applied in a normal trichromatic camera. The Fit First Method readily fits the 4 sensors to the 3-function orthonormal basis. Some further calculations call for re-examining the matrix methods, but in any case, the projection matrix handles the initial curve-fitting step. The four sensitivities are seen at right. The key steps look the same:

Rcam = RCohen(rgbSens)
CamTemp = Rcam*OrthoBasis
GramSchmidt(CamTemp, CamOmega)

. Recall that the projection matrix Rcam is a big square matrix of dimension N×N, where N is the number of wavelengths, which might be 471. The 4th sensor adds a column to the array rgbSens, but does not change the dimensions of the result Rcam.

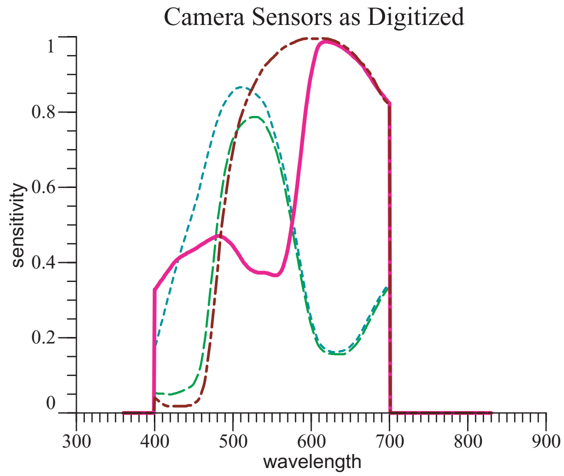

camera = Sony ICX429AKL **4-color array**
Fit first, then orthonormalized.
short dashes = cyan, longer dashes = green
dash-dot = yellow, solid = magenta

For an auxiliary step, I had to revise the program. The program output explains the algorithm as follows:

Similar to Eqs. 15-18 in CIC 14 paper, transform from sensors to CamOmega:
We want to solve CamOmega = rgbSens * Y , where Y is coefficients for 3 linear combinations.
MPP = inv(rgbSens'*rgbSens) * rgbSens'
Y = MPP*CamOmega =

0.24845	0.13737	0.35971
-0.34663	-0.43708	-0.45585
-0.22433	-0.088434	-0.033503
0.26839	0.22522	0.055984

Column amplitudes = vector length of each column =

0.55158	0.51812	0.58433

The columns of rgbSens actually contain the 4 sensitivities, cyan, green, yellow, magenta. MPP is the Moore-Penrose Pseudoinverse. (See Wikipedia and pp. 9-10 above.) Keeping in mind that the sensitivities are all ≥0, matrix Y gives some idea how much subtraction is done to produce the sensor chip's orthonormal basis. That's a step toward thinking about noise.

Below are the 3 orthonormal functions, and also the 3 best-fit functions made from the 4 camera sensitivities. The only source of “noise” is the errors that I introduced while converting graphs to numbers. It becomes more visible here, after subtractions.

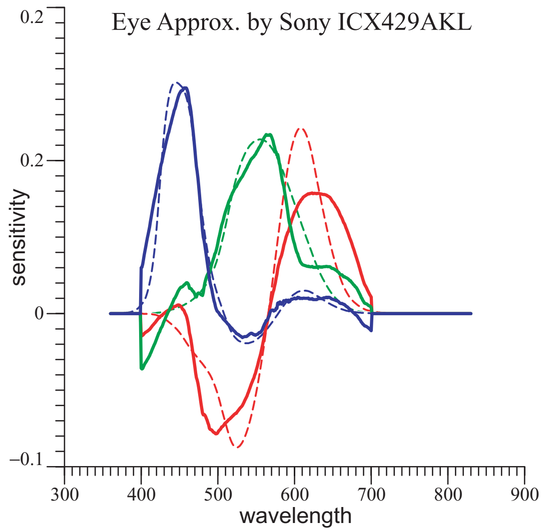

camera = Sony ICX429AKL **4-color array**
dashed curves = 2-deg observer, orthonormal
solid = best fit of camera functions to 2-deg observer

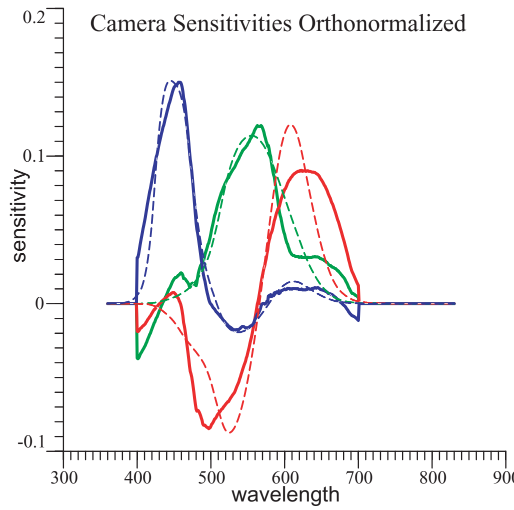

camera = Sony ICX429AKL **4-color array**
Fit first, then orthonormalized.
dashed curves = 2-deg observer
solid = orthonormal combo of camera functions

Some noise also shows up in the projections of the LUM, below. It would appear that color fidelity is not good; reds and oranges may lose some redness.

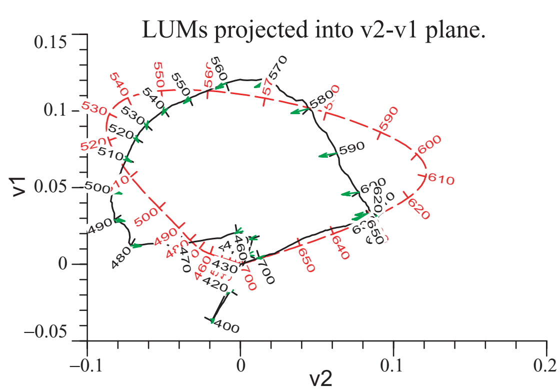

camera = Sony ICX429AKL, 4-band array
Fit first, then orthonormalized.
dashed = 2-deg observer
solid = LUM of the 4-band array
arrowheads = best fit of camera to human

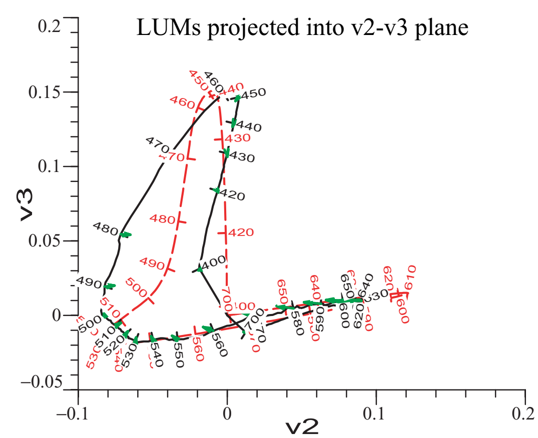

camera = Sony ICX429AKL, 4-band array
Fit first, then orthonormalized.
dashed = 2-deg observer
solid = LUM of the 4-band array
arrowheads = best fit of camera to human

Combining LEDs to Make a White Light

In the 1970s, WILLIAM A. THORNTON asked an interesting question: If you would make a white light from 3 narrow bands, how would the choice of wavelengths affect vision of object colors under the light? His research led to the Prime Colors, a set of wavelengths that reveal colors well. From that start, he invented 3-band lamps and was named Inventor of the Year in 1979. He continued his research and made the definition of prime colors more precise.

For the 2º Observer, THORNTON's Prime Colors are **603, 538, 446** nm. [See **CIC 6**, and MICHAEL H. BRILL and JAMES A. WORTHEY, "Color Matching Functions When One Primary Wavelength is Changed," *Color Research and Application*, **32**(1):22-24 (2007).] A light with 3 narrow bands at those wavelengths, will tend to enhance red-green contrasts, making some colors more vivid, though it would do a bad job with saturated red objects. You might think that combining 3 LEDs whose SPDs **peak** at the prime colors would enhance reds and greens. This idea falls short because LEDs are not narrow-band lights. We'll now see what happens when real LEDs are combined.

At each step below, detailed information is graphed, not lost. Trial-and-error moves along quickly.

Let LED Peaks ≈ the Prime Colors

The reference white is 5500 Kelvin blackbody. From 119 types measured by Irena Fryc, 3 LEDs are chosen by their peak wavelengths, as shown at left. Then we see at right that the blackbody makes a bigger swing to green and back. This LED combo will dull most reds and greens.

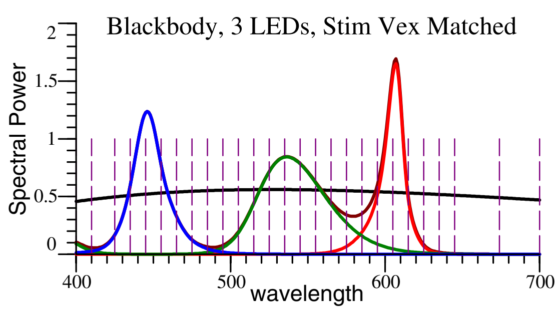

Black = blackbody 5500K, maroon = total of LEDs
Red, green, blue = individual LEDs
Indices of the LEDs = 80, 57, 28, and their names are:
Roitchner Lasertechnik LED600-03V
CREE NIST 0003, and
Roitchner Lasertechnik LED 450-01U
The peaks occur at **607**, **536**, and **446** nm
Compare to prime colors **603**, **538**, **446** nm

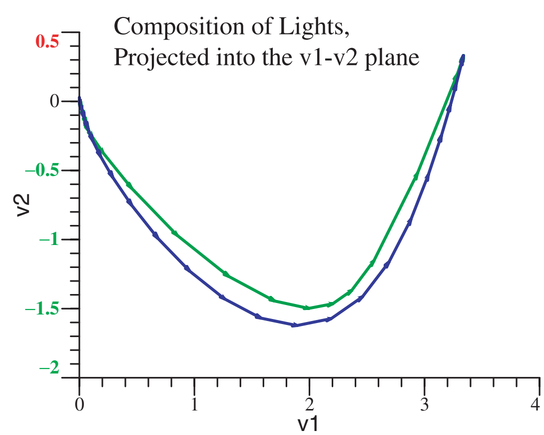

Blue = blackbody reference light, **Green** = 3 LEDs
Mon 2006 Jan 30, 11:45:30
Indices of the LEDs = 80, 57, 28 .
Peaks are at 607, 536, 446 nm.

For the **next trial**, we let the green be greener (shorter λ) and the red be redder (longer λ). In all cases, the LED amplitudes are adjusted for a color match with the blackbody. Again, SPDs are at left and the resulting composition of the lights is at right.

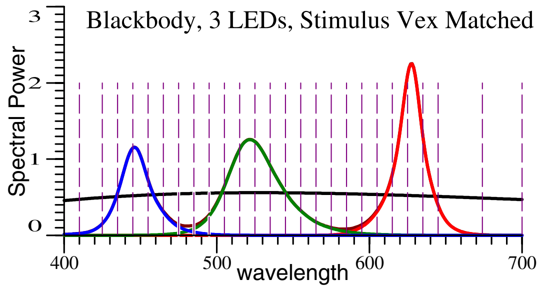

Black = blackbody 5500 K, maroon = total of 3 LEDs.
Red, green, blue = individual LEDs
Indices of the LEDs = 86, 56, 28 and their names are:
HLMP-EH25-SV000 Red-Orange
HLMP-CM15-S0000, and
Roitchner Lasertechnik LED 450-01U
The peaks occur at **628**, **522**, and **446** nm.
Compare prime λ's at **603**, **538**, and **446** nm.

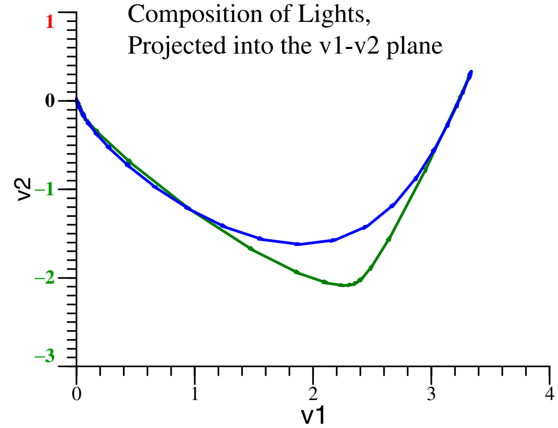

Blue = blackbody 5500 K reference, **Green** = 3 LEDs.
Indices of the LEDs = 86, 56, 28 and their names are:
HLMP-EH25-SV000 Red-Orange
HLMP-CM15-S0000, and
Roitchner Lasertechnik LED 450-01U
The peaks occur at 628, 522, and 446 nm

This light shows that indeed the LED combo can exaggerate reds and greens. A small exaggeration might be good, but this looks clumsy, with certain colors especially distorted. More details can be seen on the web page.

At this point, we guess that a broader band is needed in the red. The red prime color is at 603 nm in the orange, but longer-wavelength reds plot to distinct vector directions, as seen in the locus of unit monochromats. The problem is known to lighting experts, but you can see it yourself in the 3D locus of unit monochromats. Our fix is to let the "red LED" be in fact two reds, used in equal proportions. Think of those green and red limit colors, reflecting bands at the ends of the spectrum. Some of those reds and greens will still be dulled, but we are tracking the blackbody pretty close. More tweaking is possible.

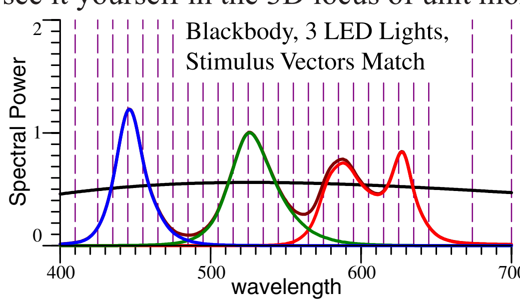

Black = blackbody 5500 K, maroon = total of LEDs
Red, green, blue = individual LEDs
Indices of the LEDs = 86, 78, 53, 28 and their names are
HLMP-EH25-SV000 Red-Orange
LEDTRONICS INC. L200CY5B, and
LEDTRONICS INC. BP280-OAG-050T-S
Roitchner Lasertechnik LED 450-01U
The peaks occur at **628**, **588**, **526**, **446** nm
Compare prime λ's at **603**, **538**, and **446** nm.

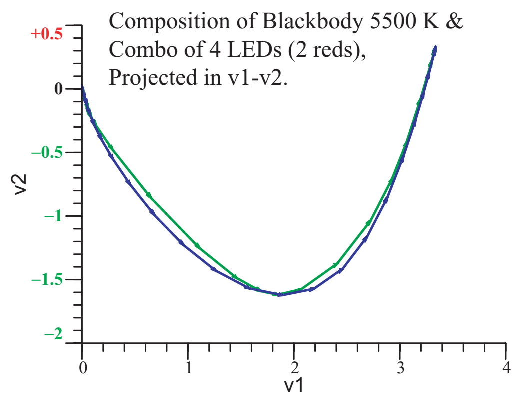

Blue = blackbody 5500 K reference, **green** = 4 LEDs
Indices of the LEDs = 86, 78, 53, 28 and their names are
HLMP-EH25-SV000 Red-Orange
LEDTRONICS INC. L200CY5B, and
LEDTRONICS INC. BP280-OAG-050T-S
Roitchner Lasertechnik LED 450-01U
The peaks occur at 628, 588, 526, 446 nm

Vectorial Color - Summary

- Sets of functions can be equivalent for color matching, even though their graphs look different, and they have other meanings, such as cone sensitivities or opponent functions, Fig. 1. This great insight is part of our legacy from the 20th century and earlier.
- But, which color matching functions are best for graphing color vectors? There is no 20th-century answer, but you can take GUTH's 1980 functions, and orthonormalize them, and that will work. That's the orthonormal basis.
- JOZEF COHEN noticed that colors have intrinsic vector relationships. His work led to the invariant Locus of Unit Monochromats.
- In its original use, COHEN's projection matrix **R** takes any spectral power distribution *L* and finds a metamer *L** that is a linear combination of the CMFs used to make **R**.
- More generally, **R** can be used to fit any function by a linear combination of functions. Using **R** in this way leads to the 3-step Fit First Method for finding a camera's Locus of Unit Monochromats.
- We have analyzed lights and cameras. The graphical methods tend to preserve details, not lose them.
- Matrix **R** and the orthonormal basis lead to interesting algebraic methods. See the Q&A page.
- In making vector diagrams, you'll find that you must adjust a vector length, or adjust two vectors to match, and so forth. Those steps may feel unfamiliar.

Invariants

My numbers and graphs are based on the 2º observer. If a different standard observer were used, details could change, but certain basic ideas would not change:

- The formula for **R**:

$$\mathbf{R} = A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}} \quad . \qquad (10)$$

- The fact that **R** is always the same numerical array, for a given observer.
- The relationship of **R** to the orthonormal basis:

$$\Omega^{\mathrm{T}}\Omega = \mathrm{I}_{3\times 3} \quad . \qquad (15)$$

$$\Omega\Omega^{\mathrm{T}} = \mathbf{R} \quad . \qquad (16)$$

- The LUM for the 2º observer, including the wavelengths of the longest vectors. Plotting the rows of Ω draws the LUM. A different orthonormal basis for the 2º observer will draw *the same LUM*.
- The fact of red-green overlap, and the implication that a light can be deficient in red *and* deficient in green.
- The idea that a camera has an LUM, and if its LUM is the same as the eye's, then the camera meets the LUTHER criterion.

If vectorial color is forgotten and then re-invented in 50 years, these ideas will come out the same. They transcend the personalities of those who discovered them.

Food for Thought...

Vectorial Color Is About Color Mixing.

It Is Not About Any Neural Processing, Only Transduction.

COHEN's ideas relate to *making best use of color mixing data*. The orthonormal CMFs are a convenient way to map stimuli into COHEN's space. They are not a hypothesis about retinal wiring.

The End

Thank you very much for signing up and for your attention.

Please feel free to contact me at any time. I am always

eager to discuss lighting, color, cameras etc.

Jim Worthey

www.jimworthey.com

Lighting for a Copy Machine??

In 2007, I promised to combine the camera analysis and the lighting analysis, to analyze something like a copy machine, where the sensors and the lighting are under engineering control. In 2008, I did not make that promise. **This is now a bonus section.** It gets confusing because there are 2 independent variables, the light and the sensors. It will be assumed that the copy machine uses a combination of LEDs for a light. Then the copier discussion is a continuation of the LED discussion above, pp. 58-61.

So, We Know How to Design a Light,
Now Let the Camera also be a Variable

Let the lights be the last pair demonstrated above. The reference white is 5500 K blackbody. The 2 red LEDs are in fixed proportion, but the R, G, and B LEDs are adjusted to match the blackbody. The adjustment of R, G, B is done for human, and then separately for the camera. Using the Fit First method, it is possible to graph the human and camera sensitivities together, and then the compositions of the lights as seen by human and by camera.

Analysis of (very hypothetical) Color Copier

The same LEDs are used as in the final light design above. The camera is the Nikon D1, whose sensitivities are given on page 48. The data were swiped from DiCarlo's paper at CIC 12.

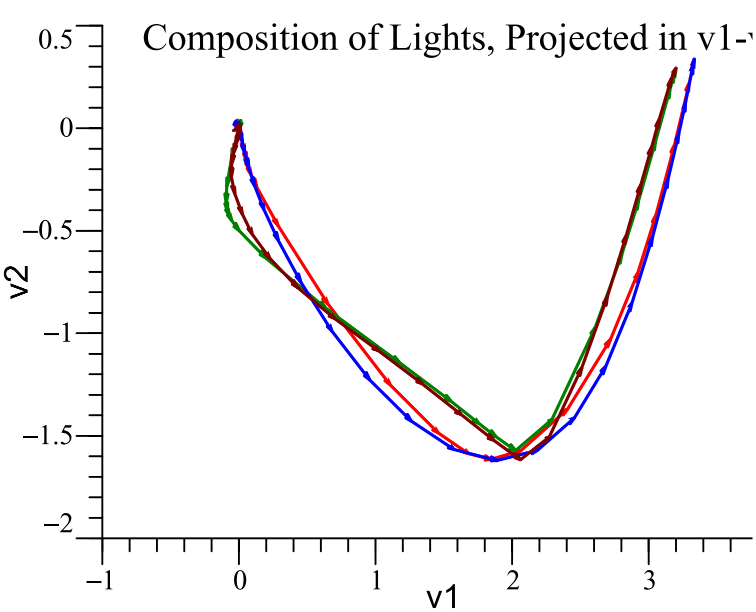

Red = 3 LEDs, human observer; Blue = blackbody, human observ
Green = 3 LEDs, camera; maroon = blackbody, camera.
Indices of the LEDs = [86, 78, 53, 28] .

A Different Hypothetical Copier

Now the light is the same, but the camera is the “optimal” sensor, taken from Quan’s dissertation.

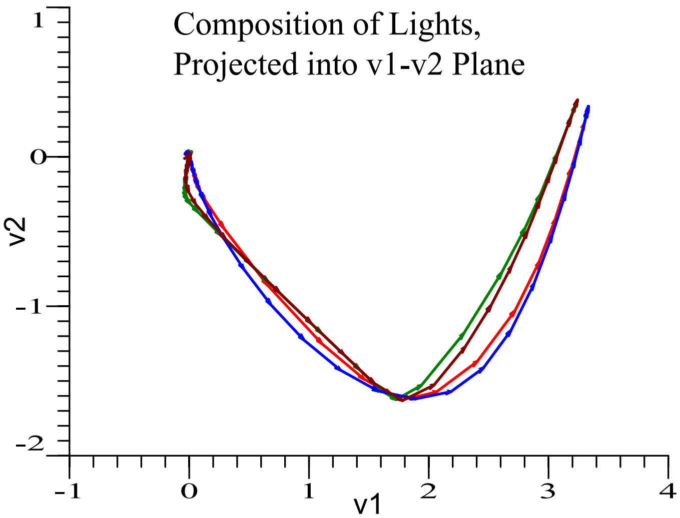

Red = 3 LEDs, human observer; Blue = 5500 K blackbody, human.
Green = 3 LEDs, camera; maroon = 5500 K blackbody, camera.
Indices of the LEDs = [86, 78, 53, 28] .
Peaks are at [628, 588, 526, 446] nm.
Camera is Quan's Optimal Rgb, represented by orthonormal basis.

The camera’s total “reach” into the green and back towards red is like the eye’s. But, many reds and greens will be dulled. You can see it by thinking of them as pass-band colors. (MacAdam limit colors as above)

The camera falls short because parts of the chain of vectors are too straight. Successive bands point in the same direction in color space, at least in this projection. We can see the same thing by looking at the Quan sensor’s Locus of Unit Monochromats.

Quan Sensor Locus of Unit Monochromats

At right, the LUM of the Quan Sensor is compared to that of human. This sensor is better than many; It tracks the human LUM in a general way. We can see shortcomings in two areas.

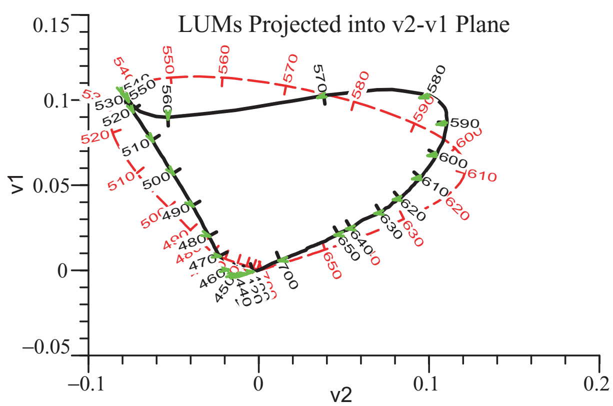

camera = Quan Optimal *rgb*,FF method
Fit first, then orthonormalized
red dashes = LUM of 2 degree observer
black solid = LUM of Quan Optimal *rgb*, fit first method
heads of green arrows = best fit to 2 deg obs by *rgb*

Picture the “unit monochromats,” the vectors that define the curves. The human vectors have steadily changing direction but long amplitude, from about 540 nm to 605 nm. Then the human LUM goes smoothly round the bend, but there is still some change of direction at 610, 620, 630 nm. By contrast, the Quan sensor goes round the bend too soon, at 585. In a second quirk, the Quan sensor tends to lump together 520, 530, 540, and 550 nm, wavelengths that would be well discriminated by human. These features of the Quan LUM translate into the too-straight regions when lights are composed on p. 61.

The lumping-together of wavelengths around 540 is something that many cameras do, and it happens because the red sensitivity is negligible there. The QUAN sensor does discriminate those wavelengths along the v_3 or blue-yellow dimension. There is no way to fix the lumping-together by changing the light.

At the long-wavelength end, it might be possible to get some marginal improvement by boosting the light at long wavelengths, since the issue is a kind of absolute sensitivity dropoff.