

Investigation Of The Relationship Between Spectral Smoothness And Color Constancy

John Seymour

Keywords: color inconstancy, metamers, illuminants, spectral smoothness

Abstract

When a spot color ink is initially formulated, the ink technician needs to decide between many potential combinations of base inks that could yield the desired color under the desired illuminant. One of the primary criteria for selection is to restrict the number of base inks that are used. A second criterion which is commonly used is the *metameric index*. This is the color difference as computed between the spectrum of the target color and predicted spectrum of the recipe under consideration. A third potential criterion is color constancy – the formulation should, as much as possible, maintain its appearance under different illuminants. This criterion is never used, at least in part because we don't know how to measure it!

This paper considers a few possible metrics for color constancy, specifically at whether a smooth spectrum is an indicator of a spectrum that will prove to be more color constant.

Introduction

A color formulator often has the spectrum of the target color from a previously printed product already on the store shelves. Using this as a target under various illuminants is a very reasonable choice. One wants the new product to match the old under whatever illuminant happens to be in the stores.

On the other hand, it is common for the spectrum of the target color to be based on measurements from a color guide, such as Pantone. Matching the printed product against a Pantone guide in a viewing booth in the pressroom is a useful step in assuring the proper color has been printed. A properly formulated ink must provide an acceptable match under the agreed upon illuminant. But there is no practical benefit for a product on a shelf to perfectly match the appropriate color in the

John the Math Guy, LLC & Clemson University

Pantone guide. There is no guarantee that the existing product will be a spectral match to a patch in the Pantone guide.

So, how can a color formulator pick from among the multiple possible base ink combinations?

A rule-of-thumb method for picking from a list of potential recipes is a subjective judgement about the smoothness of the spectrum. This is summarized in the statement that smooth spectra are less likely to have problems with metamerism [1]. The statement itself is not rigorous, since it does not include a mathematical quantification of “smooth” and it presumes that there is a defined spectrum to assess metamerism against. But it is a reasonable goal. After all, a gray ink with a perfectly flat spectrum will not change color under different illuminants; pure gray will always look gray.

A careful reading of the previous paragraph will reveal a very subtle change of topic from *metamerism* to *color constancy*. The former is defined as the property that two objects with dissimilar spectral reflectance curves may match under one illuminant, but not under another. The latter is defined as the degree to which an individual object changes color between two illuminants. (For further explanation of the difference, see [6].) Color constancy is the more appropriate measure when there is not an established spectrum for the target color – the brand owner wants the color to “look the same” under any illumination.

This paper investigates the premise that a suitably defined measure of spectral smoothness can be used as a proxy for color constancy. This paper will define one measure of color constancy based on previous work by this author [2] and will define two metrics for spectral smoothness. The first experiment will use a previously developed database of metamers [3] to test the correlation between the measures. The second experiment will use a new collection of metamers to resolve a remaining question.

How to use CIELAB (and why not to)

It would seem obvious to use CIELAB to determine the amount of color change when an illuminant is changed. You would compute the CIELAB values under one illuminant, compute them again under the second illuminant, and then compute the color difference between these two CIELAB values using your favorite color difference formula. That last step is problematic.

Brill [4] has provided a warning in the very title of a recent paper: Is CIELAB one space or many? His thesis is that CIELAB values computed under the D50 illuminant with the 2° observer, for example, should be treated as a different color space than CIELAB values computed under F11, for example.

This is a general statement. This author has directly addressed the question of whether CIELAB can be used to determine color constancy [5]. Quoting from this paper:

The conclusion is that cross-illuminant color difference calculations, where the two CIELAB values to be compared have been computed under different illuminants, will provide misleading results. Cross-illuminant color differences tell more about the failings of CIELAB than about the color inconstancy of the sample.

The problem is inherent to CIELAB. The equations used to compute $L^*a^*b^*$ include a normalization step that is performed on XYZ values. This step is performed to ensure that pure white is always (100, 0, 0). The human visual system includes a normalization step to rebalance the channels when the illuminant changes in spectral distribution or intensity. This renormalization is akin to the auto-gain feature which is built into digital cameras. However, in the human visual system, this normalization is not performed on the XYZ values. The tristimulus functions which beget the XYZ values do not exist in the human visual system, aside from in the cognitive brains of certain individuals who have studied color science!

CIELAB should not be used to assess color constancy. At the very least, any assessment of color constancy should be done with a color space that it built on the cone functions. Examples of this include IPT [7], CIECAM02 [8], and CIECAM16 [9].

How can we assess color inconstancy?

If “smoothness is next to godliness” as Smiley’s dictum claims, then perhaps the smoothness of the spectrum is simple proxy for color constancy. It seems a reasonable hypothesis that a surface with a smooth spectrum might change color less when the illuminant changes.

Two experiments were performed to test this hypothesis. The first was broad but shallow. A database of over 5,000 spectra of real objects was interrogated to see if there was a correlation between two measures of smoothness and color inconstancy. The second experiment was narrow, but deeper. A single spectrum was perturbed to create 1,000 metamers under D50/2 metamers. The same potential correlation was investigated.

Measure of color inconstancy

These experiments will use ConeLab2 to assess color inconstancy. This color space was developed for a previous paper [5] to assess the magnitude of false color inconstancy that is caused by using the tristimulus functions instead of the spectral responses of the cones. ConeLab2 was designed to provide $L_c a_c b_c$ values that as

similar as possible to $L^*a^*b^*$, but which are based on LMS rather than XYZ. The equations are provided below, copied from [5] with minor editing.

The *LMS* values are computed from spectra in the same way as the *XYZ* tristimulus values, except with estimations of the *L*, *M*, and *S* cone functions. But which estimates of the cone functions should be used? A valid argument could be made for any of the variety that are available, but one consideration puts the decision in perspective: the differences between the various LMS approximations are insignificant compared to the difference between any of them and the Standard Observer.

A second consideration is that it would be useful for metamers under CIELAB to still be metamers in the new color space. This allows for the metameric data base to be used again. Ultimately, it was decided to go with an estimate of the LMS functions which is based on the Bradford Chromatic adaptation transform. The fact that this transform is a linear combination of the XYZ values allows for CIELAB metamers to continue to be metamers.

The cone functions $l(\lambda)$, $m(\lambda)$, and $s(\lambda)$ are defined in terms of the standard observer in Equations 1 – 3, where λ is the wavelength.

$$l(\lambda) = 0.8951\bar{x}(\lambda) + 0.2664\bar{y}(\lambda) - 0.1614\bar{z}(\lambda) \quad (1)$$

$$m(\lambda) = -0.7502\bar{x}(\lambda) + 1.7135\bar{y}(\lambda) + 0.0367\bar{z}(\lambda) \quad (2)$$

$$s(\lambda) = 0.0389\bar{x}(\lambda) - 0.0685\bar{y}(\lambda) + 1.0296\bar{z}(\lambda) \quad (3)$$

The *LMS* values are computed analogous to *XYZ* from the reflectance spectrum of the sample, $r(\lambda)$ under the designated illuminant, $I(\lambda)$.

$$L = \sum_{\lambda} r(\lambda) I(\lambda) l(\lambda) \quad (4)$$

$$M = \sum_{\lambda} r(\lambda) I(\lambda) m(\lambda) \quad (5)$$

$$S = \sum_{\lambda} r(\lambda) I(\lambda) s(\lambda) \quad (6)$$

Next, we need to replicate the functionality of the equations that define CIELAB. The analog of L^* , which is called L_c , is worth a little discussion. The *Y* tristimulus function corresponds well with our perception of lightness, so this functionality needs to be preserved. L_c must be based on something like *Y*. To be consistent with ConeLab2, L_c should be computed from *LMS*. But *Y* is not one of the cone functions, but rather a weighted sum of the *LMS* functions.

The human visual system incorporates a nonlinearity which must be reflected in L_c . The question comes, do we compute a weighted sum of LMS and then apply to nonlinearity to that? This would mean that $L_c=L^*$. Or do we perform the weighting function separately on each of L , M , and S , and then compute a weighted sum of these values?

If the nonlinearity occurs in the cones or the encoding of the signal from the cones, then individually applying the nonlinearity to the L , M , and S values more closely emulates the human visual system. This option was chosen, since it is consistent with the premise of the earlier paper.

The CIELAB equations for L^* , a^* , and b^* contain the scaling values 116, 16, 500, and 200. These need to be modified to account for the fact that X , Y , and Z are functions of L , M , and S . The values for the constants k_n need to be determined so that the new space is as close to the same size as CIELAB as possible. These equations incorporate nonlinearity correction of LMS as well as normalization according to LMS by dividing by L_n , M_n , and S_n .

$$L_c+16 = k_1 f\left(\frac{L}{L_n}\right) + k_2 f\left(\frac{M}{M_n}\right) + k_3 f\left(\frac{S}{S_n}\right) \quad (7)$$

$$a_c = k_4 f\left(\frac{L}{L_n}\right) + k_5 f\left(\frac{M}{M_n}\right) + k_6 f\left(\frac{S}{S_n}\right) \quad (8)$$

$$b_c = k_7 f\left(\frac{L}{L_n}\right) + k_8 f\left(\frac{M}{M_n}\right) + k_9 f\left(\frac{S}{S_n}\right) \quad (9)$$

Linear regression was used to find the k_n scaling factors in Equations 7, 8, and 9 that provide the best fit to a large collection of small ΔE_{ab} color differences based on the printing metamers database. Equation 10 is the equation for ConeLab2, in matrix form, based on the regression.

$$\begin{bmatrix} L_c+16 \\ a_c \\ b_c \end{bmatrix} = \begin{bmatrix} 55.34 & 56.19 & 5.78 \\ 271.29 & -311.89 & 39.07 \\ 97.25 & 83.01 & -178.79 \end{bmatrix} \begin{bmatrix} f\left(\frac{L}{L_n}\right) \\ f\left(\frac{M}{M_n}\right) \\ f\left(\frac{S}{S_n}\right) \end{bmatrix} \quad (10)$$

Caveats

It was previously demonstrated that CIELAB does not accurately predict color inconstancy due to a reliance on the artificial Standard Observer. ConeLab2 corrects this error, but there has not been any experimental verification of this statement. One known weakness is that ConeLab2 includes the assumption of complete adaptation, so it is expected that changes in color may be overestimated. It is not known how accurately ConeLab2 predicts color inconstancy, but it is almost certainly better than CIELAB. It also benefits from the simplicity of CIELAB and IPT.

ConeLab2 is in many ways an improvement to CIELAB, simply because it is based on LMS rather than XYZ. Still, it is not recommended as a replacement for CIELAB. In particular, ConeLab2 is scaled to match CIELAB. As such, it will suffer the same distortions in color difference equations. When CIELAB is relaxed, it should be replaced with a color space that does better at predicting color difference data.

The first experiment

A database of 5,402 spectra of real objects was created. These were actual measurements of Crayola crayons (97), expanded gamut characterization data (876), GOE guide data (2060), Munsell color checker (24), Olympic paint (1225), and Pantone guide (1126).

For each of the 5,402 spectra, the ConeLab2 values were computed under two illuminants and the Pythagorean distance was computed between these. (This is the same formula used to compute the ΔE_{ab} color difference, but since ConeLab2 is a different color space, it cannot rightly be called ΔE_{ab} . I did not use ΔE_{00} since this color difference equation was been created to correct for the nonuniformity inherent to CIELAB. Since some of the nonuniformity may have come from the chose of XYZ, this color difference formula may not be prudent.)

In addition, the first and second derivative spectra were computed for each of the 5,402 spectra. The RMS was computed of the derivatives to reduce these two measures of smoothness down to single numbers. Finally, the chroma values (in ConeLab2) were computed. Each of these three values correlated well with the ΔE_c color difference due to illumination change.

Table 1 shows the results of using the various parameters to predict the color inconstancy when going from FL 1 to D65. The first line shows that there is, on average, just over 6 ΔE_c of color change due to this change in illuminants. The next two lines show that either the first or the second derivative account for about 40% of the variation. Either of these would be a simple and moderately successful predictor of how drastically a given sample would change color.

Illum #1	Illum #2	Model	Residual (ΔE_c)
FL 1	D65	None	6.26
FL 1	D65	1st deriv	3.86
FL 1	D65	2nd deriv	3.99
FL 1	D65	C*	3.21
FL 1	D65	1st deriv, C*	3.04

Table 1 – Color inconstancy from FL 1 to D65

These positive results could be somewhat misleading, however. The first and second derivatives are also moderately good predictors of the chroma of a sample. Colors with high chroma tend to have abrupt changes in at least one part of their spectrum. The next line in Table 1 table looked at whether the chroma (computed from the ConeLab2 values), might also be a good predictor. As can be seen, C* is a slightly better predictor.

It could be that C* is providing additional information to predict the color difference, or it could be that derivatives and chroma are just highly correlated but do not provide additional information. The final line in Table 1 shows a tiny reduction in residual when regression is done with both C* and the first derivative. The tiny improvement shows that the derivatives do not add much to the prediction.

Tables 2 and 3 show the same story with different pairs of illuminants.

Illum #1	Illum #2	Model	Residual (ΔE_c)
D50	A	None	4.17
D50	A	1st deriv	3.86
D50	A	2nd deriv	2.50
D50	A	C*	1.568
D50	A	1st deriv, C*	1.566

Table 2 – Color inconstancy from D50 to A

Illum #1	Illum #2	Model	Residual (ΔE_c)
D50	LED 4500K	None	4.64
D50	LED 4500K	1st deriv	2.67
D50	LED 4500K	2nd deriv	2.78
D50	LED 4500K	C*	1.99
D50	LED 4500K	1st deriv, C*	1.95

Table 3 – Color inconstancy from D50 to white LED

One disappointing conclusion from this analysis is that, if one wishes to have a brand color with minimal color inconstancy, then one should avoid high chroma colors.

The difference between “predicting color inconstancy via derivatives” and “predicting color inconstancy via chroma” at first may seem like a moot point. Either will do a reasonable job. But in the practical sense, the distinction is important. One of the practical use cases of a color inconstancy metric is during formulation of a recipe. If the color inconstancy for a specific recipe is closely tied to the chroma, then finding a recipe with a smoother spectrum is fruitless.

The second experiment

Is color inconstancy more a function of the chosen color, or of the way that the color is formulated? We know that higher chroma colors will tend to have more color inconstancy, but once a brand owner has decided on a brand color, is the color inconstancy locked in, or is there opportunity to reduce this undesirable trait through prudent choice of formulation?

To test this, Pantone 4665 was chosen because it is relatively smooth, but still “interesting”. This color is a desaturated brown with ConeLab2 values of (72.38, 11.03, 16.41) under D50/2. The spectrum is the black dashed line in Figure 1. One thousand metamers for this spectrum were created. Figure 1 shows the spectrum of ten metamers for this color. All the spectra have the same ConeLab2 values, since they were designed to be metamers.

The smoothness of the original spectrum was important for this analysis, since this would guarantee a range of smoothness. Those metamers with small departures from the original will tend to be smoother than those which differ more from the original.




Figure 1 – Ten metamers to the straight-line function

Details about how the metamers were created can be found in the Appendix.

Part 1 of second experiment

All of these 1,000 spectra were converted to ConeLab2 values under a variety of illuminants:

1. The official D50, D55, D65, and D75 from CIE 15.2.
2. A D50 simulation lighting booth from one manufacturer.
3. A D50 simulation lighting booth from another manufacturer.
4. 4500K LED lighting as is common in retail outlets.
5. 2800K LED lighting as is common in living room lighting.

The color values for the first lighting condition will, of course, all be the same. The spectra were computed to be metamers in CIELAB D50/2, so they will be metamers in ConeLab2 as well. The first set of illuminants looks at the variation in perceived color through the normal progression of daylight.

The second and third illuminants were chosen to investigate the effects of lighting booth simulation of D50. How widely to lighting booths disagree from with the D50 spectral calculations when it comes to apparent color?

The fourth and fifth illuminants look at two types of lighting that are found in common life.

Daylight simulators

Figure 2 shows the acbc coordinates of 1,000 metamers under the four CIE daylight simulators. It is seen that as the color temperature increases (the illuminant gets bluer), the perceived color moves closer to gray. T

he centroids of the metamers have shifted by $0.21 \Delta E_c$, $0.56 \Delta E_c$, and $0.80 \Delta E_c$. It could be that higher color temperatures de-emphasize the blue end of the spectrum, which is what primarily distinguishes tan from gray. As the color temperature departs from D50, the metameric spread also increases. This makes sense, since the difference in the spectra is larger.




Figure 2 - Comparison of perceived color under daylight simulator illuminants

Another takeaway from this plot concerns the likelihood of color constancy. The green dot represents the color of 1,000 metamers under D50. When converted to color under D55 (the red dots) almost all of them have a shift in color – not a large one. The maximum color change is less than $1 \Delta E_c$. But there are metamers where

there is virtually no color shift. In this collection of 1,000 metamers, the closest to being color constant changes color by $0.22 \Delta E_c$. These few spectra are, for all practical purposes, color constant.

Further, since the distribution of colors for the metamers increases with the difference in color temperature, it appears that the same metameric colors are color constant under the different daylight simulators. In fact, one single metamer has the best color constancy under D55, D65, and D75.

This simulation was repeated with 50,000 metamers. The closest metamer to being color constant was virtually unchanged. The conclusion is that There are certain metamers of this spectrum that are nearly perfectly color constant, but none are exactly color constant under the various phases of daylight.

That said, the perceived variations in color are minimal. All would be considered acceptable color matches if viewed side-by-side. It is unlikely that such a change would be detectable under normal conditions, that is, when the viewing is separated in time.

Viewing booths

Emission spectra were collected from a variety of viewing booths with fluorescent bulbs. The spectra of the booths were found to be distinct between two manufacturers. Spectra were averaged according to the manufacturer to arrive at two illuminants which will be called Booth 1 and Booth 2. The centroid for Booth 2 has moved a negligible $0.22 \Delta E_c$, whereas, for Booth 1, the shift is $1.04 \Delta E_c$. This again should be undetectable under normal conditions.




Figure 3 – Comparison of perceived color under fluorescent bulb viewing booths

The size of the scatter for the plots is somewhat larger for these illuminants than for the daylight simulators. This is to be expected because of the spikiness of the fluorescent tubes. But here again, with a spread of around ΔE_c , it is not expected to be detectable.

Everyday illuminants

The final comparison in Figure 4 is among illuminants that are likely to be found in everyday life. Note that the scale on the plot has been increased from 6 units square to 20 units square. This was necessary to accommodate the larger variation.

D65 (in yellow) is shown again as one representation of daylight. The red points which are labelled “Store” represent the range of metamers when viewed under 4550K LED lighting which is common in retail outlets today. The blue points, labelled “Home” are the perceived colors of the metamers under 2800K LED, which is a common warm white LED light. Both of these were measured by the author.




Figure 4 – Comparison of perceived color under common illuminants

The store lighting is only somewhat worse, on average, than the D65 outdoor lighting. This shows that color temperature is not a good indicator of color inconstancy, since according to color temperature, the 4550K LED should be just a bit closer to D50 than D55.

The much larger dispersion of the perceived colors of the metamers under store lighting is unexplained. As can be seen, different metamers of this color may vary by up to 5 units in b_c . This is conceivably noticeable.

The big surprise, however, is the drastic shift of perceived color under warm white LED lighting which is commonly found in a living room. The centroid of

these colors has shifted by $9 \Delta E_c$. This is expected to be noticeable under normal conditions. Since the variation is comparatively small, this change I perceived color is inevitable.

Part 2 of second experiment

Do the measures of smoothness correlate with the amount of color change under a different illuminant? Would either of these be a reasonable proxy for color constancy?

In the extreme case of color change from D50 to common LED lighting (blue points in Figure 4), the answer is clearly “no”. There is a larger color change for all the metamers. There may be a slight difference among them, but the overwhelming color inconstancy is very similar.

The color inconstancy from D50 to D75 (yellow points in Figure 2) has some potential, since there are at least some metamers under $0.5 \Delta E_c$, and some close to $3.0 \Delta E_c$. Figures 5 and 6 show the color inconstancy from D50 to D75 as a function of the RMS first and second derivatives.




Figure 5 – Color inconstancy of metamers as a function of average first derivative




Figure 6 - Color inconstancy of metamers as a function of average first derivative

Neither of these potential metrics show any utility for predicting color inconstancy. The smoothness of the spectrum has little or no correlation to the color constancy among this set of metamers.

Conclusions

This paper investigated a variety of techniques for picking, from a group of metamers, those that will have the best color constancy, that is, which will change color the least under different illuminants. There are several findings.

- 1) CIELAB cannot be used to assess color constancy. The difference in CIELAB values between one illuminant and another is more a function of the failure of CIELAB than anything inherent in the object being measured.
- 2) A color space called ConeLab2 was used to assess color constancy. This color space is an improvement over CIELAB, but needs further refinement before it replaces CIELAB. The key difference is that ConeLab2 is built on the cone functions rather than the Standard Observer.
- 3) The color inconstancy of 5,402 spectra was measured between various pairs of illuminants using ConeLab2. It was found that three metrics, the first derivative smoothness, the second derivative smoothness, and the CIELAB chroma, all had fair correlations to the color inconstancy.
- 4) It was noted that the two measures of smoothness also correlate to C*, suggesting that smoothness is only secondarily correlated with color inconstancy.

- 5) A second experiment looked at 10,000 metamers of a single color. All metamers had the same C^* value, but differed in smoothness. The color inconstancy was evaluated for these metamers. It was found that there was no correlation between either measure of spectral smoothness and color inconstancy.

Overall, the conclusion is that color inconstancy is more a property of the color than of the specific spectrum that creates the color.

There were two additional findings that were indirectly related to the thesis of the paper.

- 6) A means for creating metamers has been described.
- 7) For the color that was investigated, one pair of lighting conditions showed a very large color inconstancy. The change in perceived color from D50 to warm LED lighting was dramatic.

Bibliography

- [1] Personal communication with Steve Smiley
- [2] John Seymour, Color inconstancy in CIELAB: A red herring? Color Research and Application, Feb 2022
- [3] Seymour, John, Metameric Proclivity, Technical Association of the Graphic Arts, 2020
- [4] Brill, M. Is CIELAB one space or many?. Coloration Technology. 2021;137:83–85
- [5] Seymour, John, Color inconstancy in CIELAB: A red herring? Color Research & Application 47(4), February 2022
- [6] Seymour, John, Gosh, I love the word metamerism, YouTube <https://youtu.be/8qT9ZJ74C7E>, Feb 9, 2023
- [7] Fritz Ebner and Mark Fairchild, Development and Testing of a Color Space (IPT) with Improved Hue Uniformity, The Sixth Color Imaging Conference: Color Science, Systems, and Applications, Jan 1998
- [8] N. Moroney, M. D. Fairchild, R.W.G. Hunt, C Li, M. R. Luo and T. Newman, The CIECAM02 Color Appearance Model, CIC2002, Scottsdale, Arizona, pp23 27.
- [9] C. J. Li, Z. Li, Z. Wang, Y. Xu, M. R. Luo, G. H. Cui, M. Melgosa, M. H. Brill, and M. R. Pointer, Comprehensive color solutions, CAM16, CAT16 and CAM16 UCS, Col., Res. and Appl., (2017), pp 703 718.

Appendix – Creating the metamers

Creation of metamers for a given spectrum is thankfully a routine operation in linear algebra, based on the Gram-Schmidt process. Metamers (under D50/2) for any spectrum (a vector in 31 or 36 space) can be found by adding any vector which is orthonormal to the three spectral products $\bar{x}(\lambda)D50(\lambda)$, $\bar{y}(\lambda)D50(\lambda)$, and $\bar{z}(\lambda)D50(\lambda)$.

One starts with the three spectral product vectors. Singular value decomposition is used to create an orthonormal basis for these three functions. These are three basis vectors in 36 space which:

1. Span the space covered by the three spectral product vectors, that is, each of the spectral products can be expressed as the weighted sum of the three basis vectors.
2. Each have a norm of 1, that is, the dot product of any of the vectors with itself is 1.
3. Each are normal (orthogonal) to the other two, that is, the dot product of one of the basis vectors with any of the other two is zero.

In the Gram-Schmidt process, vectors are added to the set, one-by-one, with each of the new vectors being orthonormal to all preceding vectors. To do this, a seed vector is generated which is not in the space spanned by all previous vectors. For this exercise, I started with the vector which was 1 at all wavelengths. The next seed was a cosine function with frequency of one cycle over the range of wavelengths. The next was a sine wave with that same frequency. The fourth and fifth seeds were cosine and sine waves with frequency of two. The rest of the 33 seeds were generated with higher and higher frequencies.

To turn a seed into the net orthonormal vector, the projection of that seed onto all previous vectors is subtracted from the seed. This is then normalized.

The result is a set of 36 orthonormal vectors. The first three are a linear combination of the spectral product functions and are not used. Vectors 4 through 36 are metameric vectors. Any linear combination of these can be added to a spectrum to find a metamer under D50/2.

The first three of the metameric vectors are shown in Figure XX, and the next three in the following figure.




Figure XX




Figure XX + 1

A spectrum which is metameric to a starting spectrum is created by adding a randomly weighted sum of the metameric spectra to the starting spectrum.

The figures above show that, because of the order of sine/cosine seeds, the smoothest vectors in the orthonormal set tend to occur earlier in the sequence. Since the reflectance spectra of solid objects tend to be smooth, the earlier metameric spectra are more apt to provide physically reasonable spectra. For this work, I chose to work with only the first five metameric spectra, and chose weightings randomly distributed between -0.2 and +0.2.