Using Optimal Surface Reflectance Functions to Study Color Space Properties[1]

Johan M. Lammens

Scuola Superiore S. Anna
Advanced Robotics Technology and Systems (ARTS) Laboratory
via Carducci 40, 57127 Pisa, Italy
e-mail lammens@arts.sssup.it
URL http://www.cs.buffalo.edu/~lammens

Submitted to IEEE-workshop on physics-based modeling 95

Abstract

We present a computational definition of the Optimal Color Stimuli surface, which corresponds to the limit of physically realizable surface spectral reflectance functions. We can use this surface to quantify and visualize color space properties, as we demonstrate for the CIE XYZ and L*a*b* spaces, as well as our own neurophysiologically-based NPP space. We then extend the technique to compute a representative set of surface spectral reflectance functions, all contained within the volume bounded by the OCS surface, and their corresponding representations in a color space. This set of spectra can be used to generate training sets for a backpropagation algorithm computing color space transforms, for example.

[1]The research reported in this paper was done while the author was a member of the SNePS Research Group at the department of Computer Science, State University of New York at Buffalo, and was supported in part by Equipment Grant No. EDUD-US-932022 from SUN Microsystems Computer Corporation, and in part by NASA under contract NAS 9-19004.

1 Introduction

In order to study color categorization (color naming), we have derived a color space from neurophysiological data, which we call the NPP color space [5]. To study the properties of this color space and to compare it to other ones, we need a way to visualize and to quantify its "shape". Since color is a perceptual phenomenon that arises in response to certain physical stimuli, there are two important aspects to the problem:

1. How do different physical stimuli translate into different locations in the color space?
2. How do different perceptual categories translate into different locations (or regions) of the color space?

This paper deals primarily with the first problem; for the second we refer to [5] and [4]. In what follows, we will present our computational derivation of the so-called Optimal Color Stimuli (OCS) surface [11], which we will use as a reference to compare the properties of different color spaces. We also present a further refinement of the same technique, that allows us to model not only the OCS surface, but also a set of points enclosed by it, corresponding to a representative set of surface spectral reflectance functions.

2 Some Definitions

To simplify our discussion, we briefly recapitulate some basic definitions relating to color perception. The physical stimulus involved in color perception is light, or more precisely: electro-magnetic radiation in the visible wavelength range, approximately 380–770 nm. We will represent the spectral energy distribution of the stimulus (the radiation striking the retina) as a function of wavelength $E(\lambda)$, after [1]. The stimulus energy distribution $E(\lambda)$ of the radiation that is reflected off (or transmitted by) an object is determined by (1) the spectral distribution of the radiation incident on the object (the light source) $S(\lambda)$, and (2) the spectral reflectance characteristics of the object $R(\lambda)$, so we can write

$$E(\lambda) = S(\lambda)R(\lambda) \tag{1}$$

in a simplified form. This simplification does not take into account additional factors like incident and reflection angles, transmission absorption and refraction, etc. It is easily seen from this equation that changing the light source or changing the reflectance characteristics of the object changes the stimulus $E(\lambda)$. Were we informally to think of *object color* as $R(\lambda)$, then the task of the color vision system is to recover $R(\lambda)$ from $E(\lambda)$ (for some examples of this approach, see [8, 7]). Or as [12] puts it:

> The brain strives to acquire a knowledge about the permanent, invariant and unchanging properties of objects and surfaces in our visual world. But the acquisition of that knowledge is no easy matter because the visual world is in a continual state of change. Thus, the brain can only acquire knowledge about the invariant properties of objects and surfaces if it is able to discard the continually changing information reaching it from the visual environment. [p. 355]

The human visual system is to some extent able to identify the same "object color" (surface reflectance, under the view just mentioned) under widely varying lighting conditions, a phenomenon known as *color constancy.*

The response of the visual system to $E(\lambda)$ brings us into the domain of neurophysiology. The human visual system is sensitive to differences in $E(\lambda)$ due to the presence of three photopigments with different spectral sensitivities in three types of photoreceptor cells (called cones) in the retina [6]. The cone types are sometimes called *red*, *green*, and *blue*, for the spectral color corresponding approximately to their wavelength of maximum sensitivity, but more appropriate designations are *long*, *medium*, and *short* wavelength sensitive cones, respectively. If we represent the cone action spectra (or linear transforms thereof, see [3]) as $\overline{x}(\lambda), \overline{y}(\lambda), \overline{z}(\lambda)$ and integrate them with the spectral energy distribution $E(\lambda)$ of the stimulus [1, 3], we obtain the so-called CIE tristimulus values X, Y, Z:

$$X = \int E(\lambda)\overline{x}(\lambda)d\lambda \quad (2)$$

$$Y = \int E(\lambda)\overline{y}(\lambda)d\lambda \quad (3)$$

$$Z = \int E(\lambda)\overline{z}(\lambda)d\lambda \quad (4)$$

Any two stimuli that result in identical X, Y, Z values cannot be distinguished from each other; i.e. they are metamers.

3 Quantifying and Visualizing Color Space Properties

3.1 Computing the Optimal Color Stimuli Surface

To visualize the general "shape" of a color space, we have computed the shape of the Optimal Color Stimuli (OCS) Surface in various spaces. We can represent all physically possible surface-spectral reflectance functions[2] in a solid known as the Object Color Solid. The surface of this solid represents the limit of physically realizable surface colors, known as Optimal Color Stimuli, and can be generated by computing the response of a given set of sensors to a continuum of two kinds of spectra:

> The spectral reflectance $\sigma(\lambda)$ is either zero or unity, and in moving through the visible spectrum, there are generally not more than two transitions between these values. Optimal color stimuli are imaginary stimuli in the sense that no actual object surfaces have reflectance curves with abrupt transitions of this kind. However, they are of considerable interest because they represent limiting cases of all (non-fluorescent) object-color stimuli. [...] Two types of curves must be distinguished; the first has zero reflectance (or transmittance) at wavelengths $\lambda < \lambda_1$ and $\lambda > \lambda_2$, the second at wavelengths $\lambda_1 < \lambda < \lambda_2$. [11, p.181 ff.]

We have used the following differentiable approximations to these two types of reflectance functions:

$$\sigma_0(\lambda, \gamma, \theta) = \begin{cases} \frac{0.5}{(1-\lambda+\theta)^2}, & \lambda \le \theta \\ 1 - \frac{0.5}{(1+\lambda-\theta)^2}, & \theta < \lambda \le \theta + \frac{\gamma}{2} \\ 1 - \frac{0.5}{(1-\lambda+\theta+\gamma)^2}, & \theta + \frac{\gamma}{2} < \lambda \le \theta + \gamma \\ \frac{0.5}{(1+\lambda-\theta-\gamma)^2}, & \theta + \gamma < \lambda \end{cases} \quad (5)$$

$$\sigma_1(\lambda, \gamma, \theta) = \begin{cases} 1 - \frac{0.5}{(1-\lambda+\theta)^2}, & \lambda \le \theta \\ \frac{0.5}{(1+\lambda-\theta)^2}, & \theta < \lambda \le \theta + \frac{\gamma}{2} \\ \frac{0.5}{(1-\lambda+\theta+\gamma)^2}, & \theta + \frac{\gamma}{2} < \lambda \le \theta + \gamma \\ 1 - \frac{0.5}{(1+\lambda-\theta-\gamma)^2}, & \theta + \gamma < \lambda \end{cases} \quad (6)$$

[2] Or transmittance functions, which makes no difference for our purpose.

where λ is wavelength in nm, γ is the width of the "gap" in nm, and θ is the start of the "gap" in nm. Some typical examples of reflectance spectra generated by these functions are shown in Figure 1. If we assume a flat-spectrum (white) light source, defined by $\sigma(\lambda) = 1$, the light

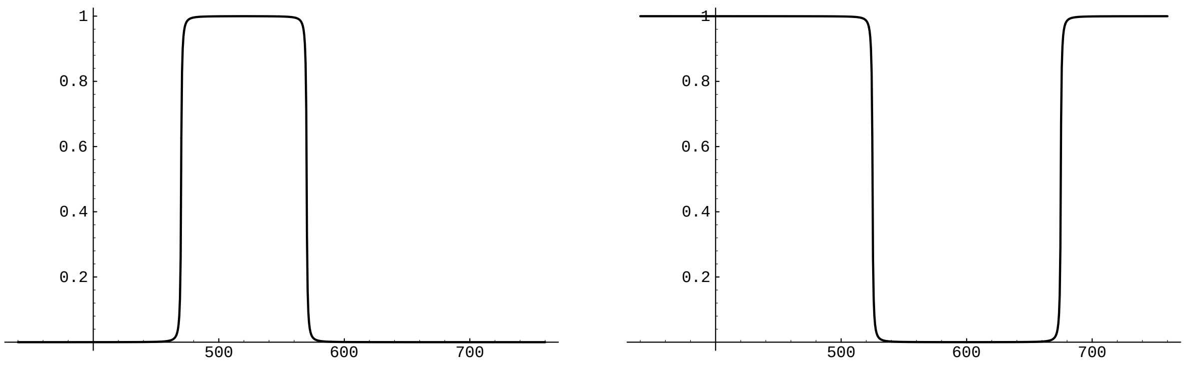

Figure 1: Some examples of the two kinds of spectra needed to generate Optimal Color Stimuli. Left: $\sigma_0(\lambda, 100, 470)$, right: $\sigma_1(\lambda, 150, 525)$. On the X axes: wavelength (nm), on the Y axes: relative reflectance. (After [11, p. 181].)

reaching the sensors has a spectrum identical to the reflectance function,[3] and we can compute Optimal Color Stimulus coordinates for *linearly* responding sensors as follows:

$$p_0 = List_{i=1}^{N} \left[\int_{\lambda_l}^{\lambda_u} \sigma_0(\lambda, \gamma, \theta)\, \overline{s_i}(\lambda)\, d\lambda \right] \tag{7}$$

$$p_1 = List_{i=1}^{N} \left[\int_{\lambda_l}^{\lambda_u} \sigma_1(\lambda, \gamma, \theta)\, \overline{s_i}(\lambda)\, d\lambda \right] \tag{8}$$

where $List_{i=1}^{N}$ is a list of expressions with index variable i ranging from 1 to N, N is the number of dimensions (basis functions) of the color space, $\overline{s_i}$ is the spectral sensitivity of each of the basis functions, and λ_l and λ_u represent the lower and upper limit of sensitivity for the sensors used, typically in the neighborhood of 300 and 800 for the human visual system, respectively. By varying γ and θ over the visible wavelength range, and plotting the resulting points p_0 and p_1, we can compute the shape of the OCS surface. It is made up of two "halves" that fit together like clam shells, corresponding to the sets of points p_0 and p_1.

If we use different sets of basis functions $\overline{s_i}$, the shape of the OCS surface will change accordingly. If we choose the standard CIE XYZ functions as defined in Section 2, we get the result shown in Figure 2. This is the typical "torpedo-like shape" that [11] refer to. For the actual computations involved in creating figures 2ff, we used a computationally more efficient technique than suggested by equations 7ff, making use of the special properties of the functions σ_0 and σ_1 and using a list of partial integrals as a kind of cache. The surface color[4] in Figure 2

[3] Ignoring other factors such as atmospheric absorption and scattering, viewing angle, etc., as we did in equation 1.

[4] For those having a color print available, or viewing this document on a color monitor. It is interesting to view this and related figures with a PostScript previewer like ghostview, where one can see the surface being built up in 3D. This gives a better appreciation of the shape, probably as good as possible short of an animated plot.

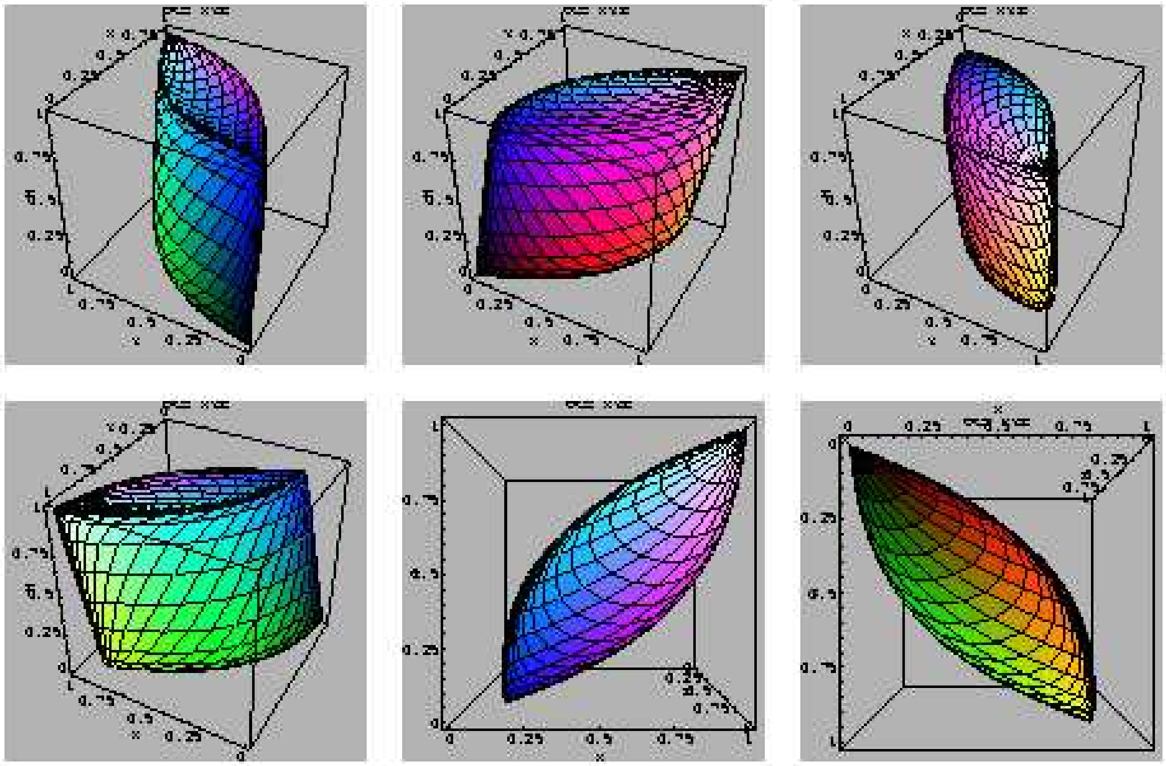

Figure 2: The Optimal Color Stimuli surface in the CIE XYZ color space (see text). The color rendering (if any) is only approximate, derived by using a transform from XYZ to RGB for typical computer monitors, and re-normalizing the result. In English reading order: four views rotating counter-clockwise around the surface, and a view from above and below.

is (necessarily) only an approximation, derived as follows:

$$\{r, g, b\} = Lt\left(\frac{\mathbf{M_{XR}} \cdot \{x, y, z\}}{\mathbf{C_X}}\right) \tag{9}$$

$$Lt(x) = \begin{cases} 0, & x < 0 \\ x, & otherwise \end{cases} \tag{10}$$

$$\mathbf{M_{XR}} = \left\{ \begin{array}{ccc} 1.26497 & -0.569038 & -0.198262 \\ -0.42012 & 0.797003 & 0.0196274 \\ 0.0275845 & -0.111426 & 0.46304 \end{array} \right\} \tag{11}$$

$$\mathbf{C_X} = \mathbf{M_{XR}} \cdot List_{i=1}^{3}\left[\int_{\lambda_l}^{\lambda_u} \sigma_1(\lambda, 0, 0)\, \overline{s_i}(\lambda)\, d\lambda\right] \tag{12}$$

where $\{r, g, b\}$ are $[0, 1]$ normalized RGB coordinates, Lt is a limiting function serving to limit RGB coordinates to the gamut of the display device, M_{XR} is a linear transform from XYZ to "typical computer monitor" RGB coordinates such as the ones given in [9] or [2], $\{x, y, z\}$ are the CIE XYZ coordinates computed with equations 7ff (p. 3), and $\mathbf{C_X}$ is the set of RGB coordinates corresponding to a maximum radiance flat spectrum (white). The latter is used as a normalization factor for display purposes, which is basically Von Kries adaptation [11]. It is clear from these equations that the displayed color has to be approximate, because of the limitations of the gamut of the display device, the "typical" transform used, and the inability to control for such things as gamma correction. Nevertheless, the rendered color is a reasonable indication of the "real" surface color corresponding to that particular point on the OCS surface.[5]

It is interesting to note that only a relatively small volume of the XYZ cube actually corresponds to physically realizable colors. The same is true of course if we compute the OCS surface in RGB coordinates, but the intersection of the volume enclosed by the OCS surface with the all-positive (i.e., displayable) RGB subspace is even smaller than in the XYZ case (Figure 3).

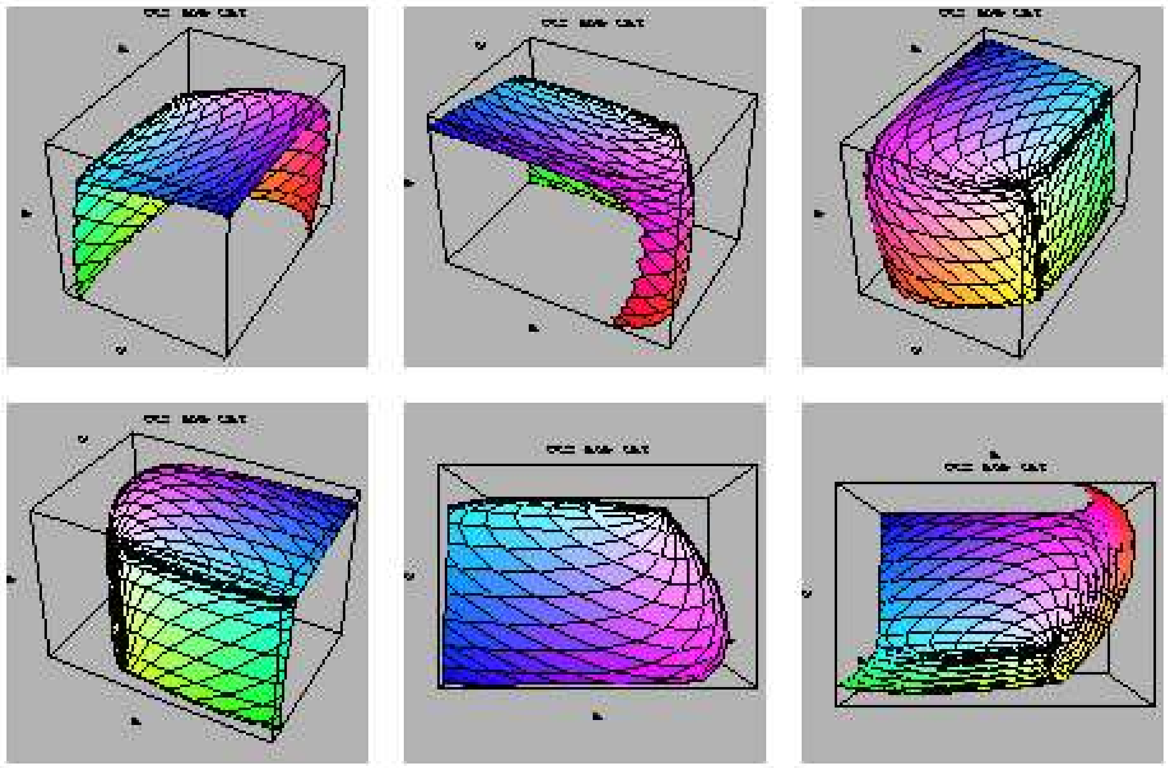

Figure 3: The OCS surface transformed into RGB space using the transform $\mathbf{M_{XR}}$ from equation 11 (p. 4). Viewpoints as in Figure 2. The parts of the surface that lie outside of the positive octant (i.e., outside of the gamut of a typical RGB computer monitor) have been clipped.

Choosing different basis functions results in different shapes of the OCS surface. The computation for the NPP basis functions is somewhat more involved then for the CIE XYZ functions, because of the nonlinearities involved in the former, which are defined as sigmoid-of-linear functions [5]. The equations for computing points on the surface are

$$p_0 = List_{i=1}^{N}\left[sig_i\left(c_i \int_{\lambda_l}^{\lambda_u} \sigma_0(\lambda,\gamma,\theta)\,\overline{lin_i}(\lambda)\,d\lambda\right)\right] \tag{13}$$

$$p_1 = List_{i=1}^{N}\left[sig_i\left(c_i \int_{\lambda_l}^{\lambda_u} \sigma_1(\lambda,\gamma,\theta)\,\overline{lin_i}(\lambda)\,d\lambda\right)\right] \tag{14}$$

[5] Were it not that there is no such thing as a real color, of course.

where symbols are as in equations 7ff, sig_i being the sigmoid component, and lin_i being the linear component of the basis functions. Since the value of the sigmoid function is a non-linear function of its input, there is an issue here with respect to scaling the linear responses that was not relevant for the linear CIE XYZ functions. We have to determine values for the constants c_i. The method we have used is essentially scaling with respect to the equal-energy or "white" response. The argument goes as follows. We want to scale the brightness (Br) dimension to range from perceived black to white, without affecting the adaptive state of the visual system. Once we have determined a scaling factor for the Br dimension we will use it to scale the other dimensions as well. It is justifiable to use the same scaling factor for all color space dimensions, since the NPP functions we derived in [5] preserve the ratios of responses among the cell types on whose responses the NPP space is based. Assuming (as usual) that the perception of "white" arises from the viewing of a stimulus with a flat equal-energy spectrum $\sigma_w(\lambda) = 1$, we can normalize responses with respect to this type of stimulus. Since "white" is at the top of the Br dimension (always within the adaptive range), a flat spectrum "white" (maximum relative radiance) stimulus must result in the maximum response for the Br function, which is the function value at the wavelength of maximum response and at maximum relative radiance:

$$Br_{lin}^{max} = Max_\lambda \left[Br_{lin}(1, \lambda) \right] \tag{15}$$

The response of the linear Br function to a flat-spectrum "white" stimulus is given by[6]

$$Br_{lin}^{w} = \int Br_{lin}(1, \lambda)\, \sigma_w(\lambda)\, d\lambda \tag{16}$$

And the desired scaling factor is then just the quotient of the two:

$$c = \frac{Br_{lin}^{max}}{Br_{lin}^{w}} \tag{17}$$

The normalized response of a linear function to a stimulus $\sigma(\lambda)$ is then simply

$$f_{lin}^{n}(\sigma) = c \int f_{lin}(1, \lambda)\, \sigma(\lambda)\, d\lambda \tag{18}$$

The final normalized function value is

$$f_{sig} \left[f_{lin}^{n}(\sigma) \right] \tag{19}$$

as discussed above. Applying this normalization to the NPP basis functions and computing the OCS surface results in the shape shown in Figure 4.

For comparative purposes we have also represented the OCS surface in CIE L*a*b* coordinates (Figure 5). The CIE L*a*b* space represents an attempt to create a perceptually equidistant color space for reflected light. The L*a*b* model is based on psychophysical principles only, however, not on neurophysiological data as the NPP space is. Although the order of hues around the NPP and L*a*b* spaces is the same, there are marked differences in the overall shapes (e.g. the sharp protrusion of the L*a*b* space in the blue region is absent in the NPP space), and the relative positions and areas that certain colors occupy on the respective surfaces.

In [11], only black and white hand-drawn approximations of the OCS shape (in different color spaces, but not NPP of course) are shown, and we are not aware of any attempts to

[6] For integrating the linear response functions any relative radiance will do, as long as it is the same for all functions, since the result is scaled before sigmoidification. We have chosen a relative radiance of 1 for convenience.

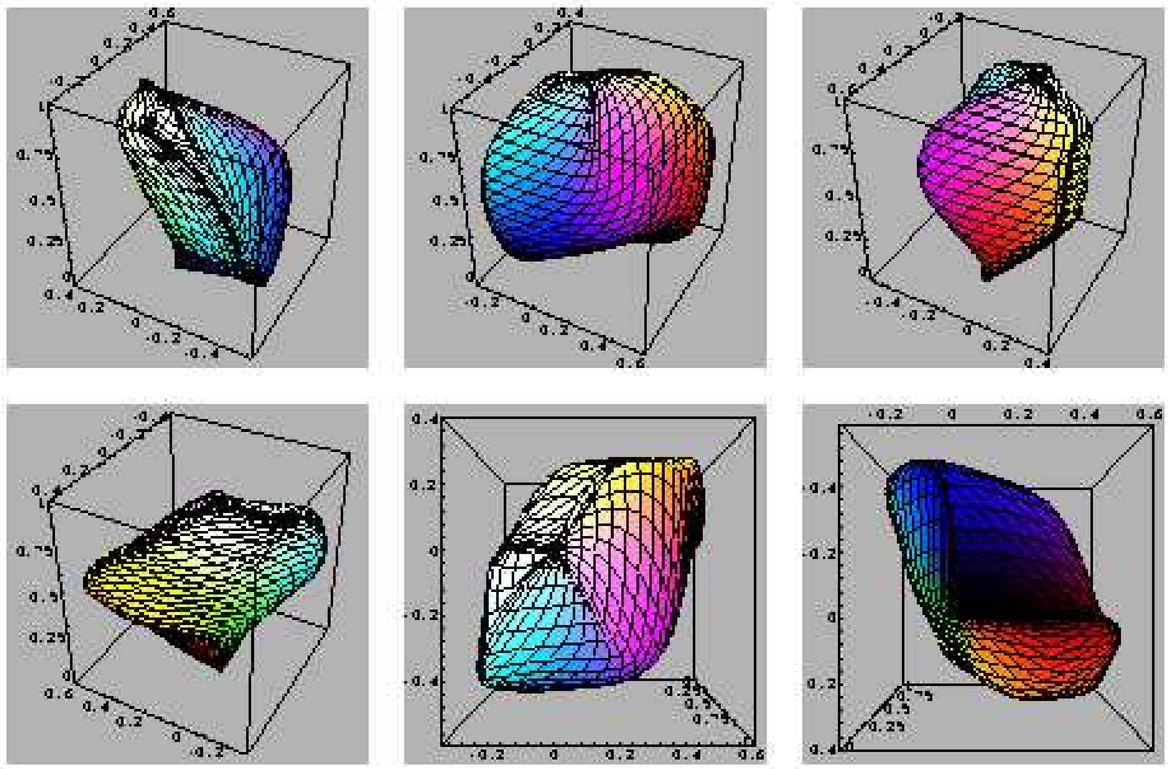

Figure 4: The OCS surface in NPP coordinates (see text).

define the complete shape and its surface color analytically as we have done. In [5] and [4] we have used the OCS surface as a frame of reference to investigate the distribution of perceptual color categories. It is well suited for that purpose, since it represents the limit of all physically possible surface reflectances (giving rise to color perception when viewed under an appropriate light source), and we can represent it in the neurophysiologically-based NPP space.

3.2 Computing Other Representative Stimuli

Having constructed the NPP space, we needed to be able to transform data coming from color sensors (typically an RGB color camera) into it, if we want the NPP space to be usable for robotic agents [5] or for general computer vision work. Since most (all?) color cameras in use today are based on the CIE XYZ standard, we need to transform XYZ coordinates to NPP coordinates. This will proceed in two steps: a transform from XYZ to a set of 6 linear basis functions, followed by sigmoidification and a final linear transform. Since it is not possible to determine a simple linear transform from XYZ to the 6 basis functions within a reasonable margin of error, we have used the error back-propagation algorithm to determine such a transform, as commonly used in artificial neural networks research [10]. In contrast with many applications of the backpropagation technique we did not use the network as a classifier, but rather to learn six simultaneous functions of three real variables (the CIE X, Y, and Z coordinates), implementing a transform from one space to another. Using the networks this way imposes much stricter requirements on the obtained transform than using them as a classifier does, since both the domain and range are continuous-valued, and there are no “bins” in the range within which differences in function values are not important, as is the case in classifier nets.

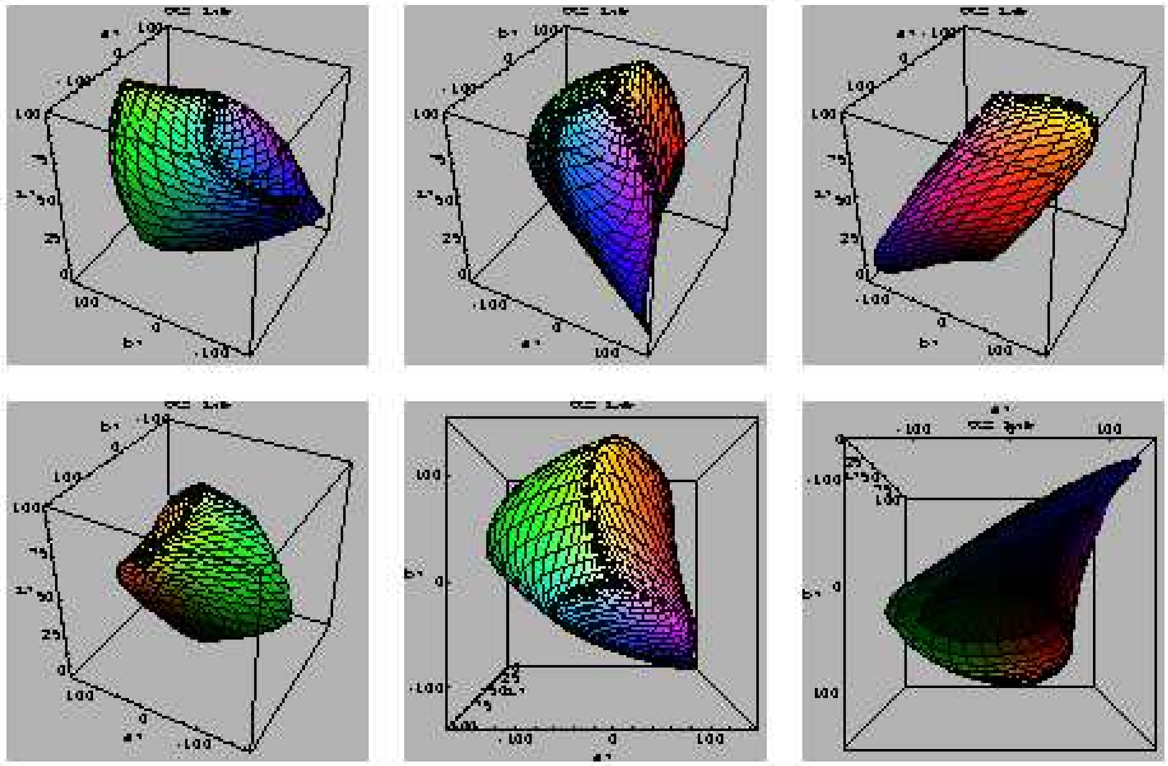

Figure 5: The OCS surface in CIE L*a*b* coordinates (see text).

One of the most critical problems in using the error backpropagation algorithm to learn a particular transform is the selection of a proper training set. Since this learning technique is a supervised one, one needs to construct a set of input-output pairs that covers the N-dimensional input and output subspaces (assuming N input and output nodes) sufficiently well to produce a smooth transform, without significant discontinuities and generalizing well to new input/output pairs. In our particular case, the natural thing to do might seem to be to construct a training set based on the pure frequency responses of the XYZ and NPP functions, by varying the frequency over the visible wavelength range:

$$
\begin{aligned}
in(\lambda) &= \langle \overline{x}(\lambda), \overline{y}(\lambda), \overline{z}(\lambda) \rangle && (20) \\
out(\lambda) &= List_{i=1}^{6} \left[\overline{lin_i}(\lambda) \right] && (21)
\end{aligned}
$$

This type of training set turns out not to work well, and some analysis reveals why: it does not contain any training points on the gray axis (since no pure frequency signal results in an equal response of all three XYZ functions) or any points in the purple region of the color space (since purple is a "mixture of red and blue" which cannot be obtained from a pure frequency signal), among other areas. What we need to do is to create a representative training set containing points corresponding to complex stimulus spectra, including grays, purples, and all other kinds. Since spectra can be thought of as real-valued functions of wavelength, the set of possible spectra is infinite. Fortunately we can exploit some constraints on the physically possible surface reflectance functions to generate a number of representative spectra, analogous to the way we constructed the OCS surface in Section 3. We know that all physically possible spectra must result in points enclosed by that surface (including points on the gray axis), so these are the spectra of interest for constructing a training set. Based on the technique described in Section 3, we have created an extended one to generate a set of spectra covering both the OCS surface and the space contained within it. The spectra are given by the following function:

$$\sigma(\lambda, \gamma, \theta, \rho_1, \rho_2) = \begin{cases} \rho_1, & \lambda \leq \theta \\ \rho_2, & \theta < \lambda \leq \theta + \gamma \\ \rho_1, & otherwise \end{cases} \tag{22}$$

where λ is wavelength in nm as usual, γ is the width of the "gap" in nm, θ is the start of the "gap" in nm, and ρ_1 and ρ_2 are two reflectance levels in $[0, 1]$. By varying γ and θ over the visible wavelength range and ρ_1 and ρ_2 over the $[0, 1]$ interval, we can generate a subset of all possible reflectance spectra, some of which are shown in Figure 6.[7] By varying γ in 10 equal

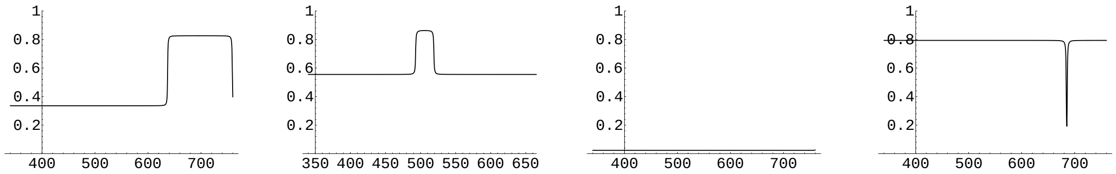

Figure 6: Some example spectra generated by varying the parameters γ, θ, ρ_1 and ρ_2 in equation 22. From left to right, parameter values are $\langle 122, 637, 0.34, 0.83\rangle$, $\langle 26, 493, 0.56, 0.87\rangle$, $\langle 192, 760, 0.02, 0.03\rangle$, and $\langle 2, 684, 0.80, 0.10\rangle$ (see text). On the X axis: wavelength in nm, on the Y axis: relative spectral reflectance.

steps from 0 to 400, θ in 10 equal steps from 380 to 780, and ρ_1 and ρ_2 each in 7 equal steps from 0 to 1, we obtain a set of 4900 spectra that cover the OCS surface and the space enclosed by it fairly well, as shown in Figure 7. In this figure the XYZ coordinates corresponding to the spectra are computed as in equation 7 (p. 3) ff, and these coordinates are the input values for the training set. The output values for the training set are computed analogously using the 6 linear NPP functions. Comparing Figures 2 (p. 4) and 7 we can see that there are no points outside the OCS surface, which is expected of course. Such points do not correspond to physically possible reflectance spectra, and hence are of no interest to us. For details of the learned (optimized) transform, we refer to [5].

An important measure of the quality of the transform is what it does to the OCS surface. Figure 8 shows the OCS surface as transformed by a 4-layer network from the XYZ to the NPP space. It should be compared to Figure 4 (p. 7), which shows the same surface in the (directly computed) NPP space.

4 Conclusion

We have presented a computational definition of the Optimal Color Stimuli surface and a set of points enclosed by it, which correspond to the limit of physically realizable surface reflectance functions and a representative set of surface reflectance functions, respectively. Both of these can serve as a quantitative basis to compare the properties of various color spaces, as we have

[7]This subset is itself infinite since the parameters are real-valued, but its cardinality is obviously less than that of the set of all possible spectra. By discretizing the domains of each of the parameters we can bring the cardinality down to a finite integer number.

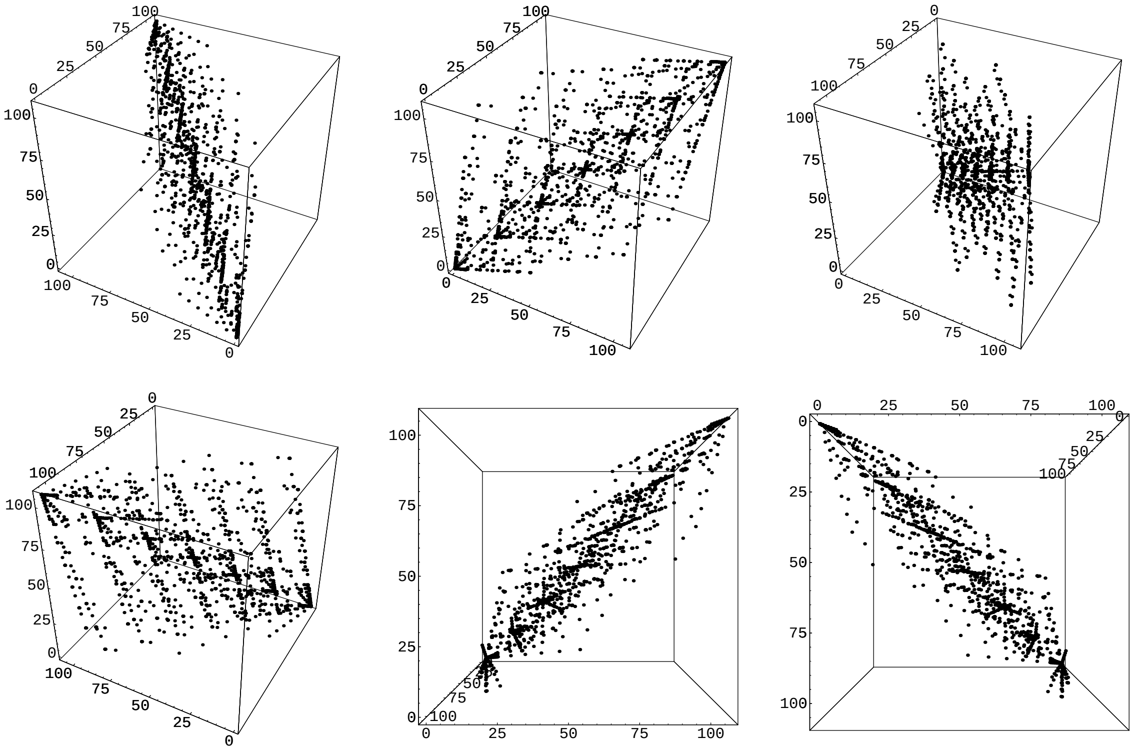

Figure 7: Backpropagation training set input points. Each point corresponds to a spectrum defined by equation 22. Comparing to Figure 2 (p. 4), we can see that the points cover both the OCS surface and the space enclosed by it.

illustrated with the CIE XYZ and L*a*b* spaces, as well as our own neurophysiologically-based NPP space.

A Getting the software

The Mathematica code implementing the described algorithms is available by anonymous ftp from

`ftp.cs.buffalo.edu:/ftp/pub/colornaming`

All software is provided as is, without any express or implied warranty. If you have problems with it, feel free to contact the author.

References

[1] R. M. Boynton. Human color perception. In K. N. Leibovic, editor, *Science of Vision*, chapter 8, pages 211–253. Springer, New York, 1990.

[2] F. S. Hill, Jr. *Computer Graphics*. Macmillan, New York, 1990.

[3] B. K. Horn. *Robot Vision*. MIT Press, Cambridge MA, 1986.

[4] J. M. Lammens. A somewhat fuzzy color categorization model. Submitted to ICCV-95.

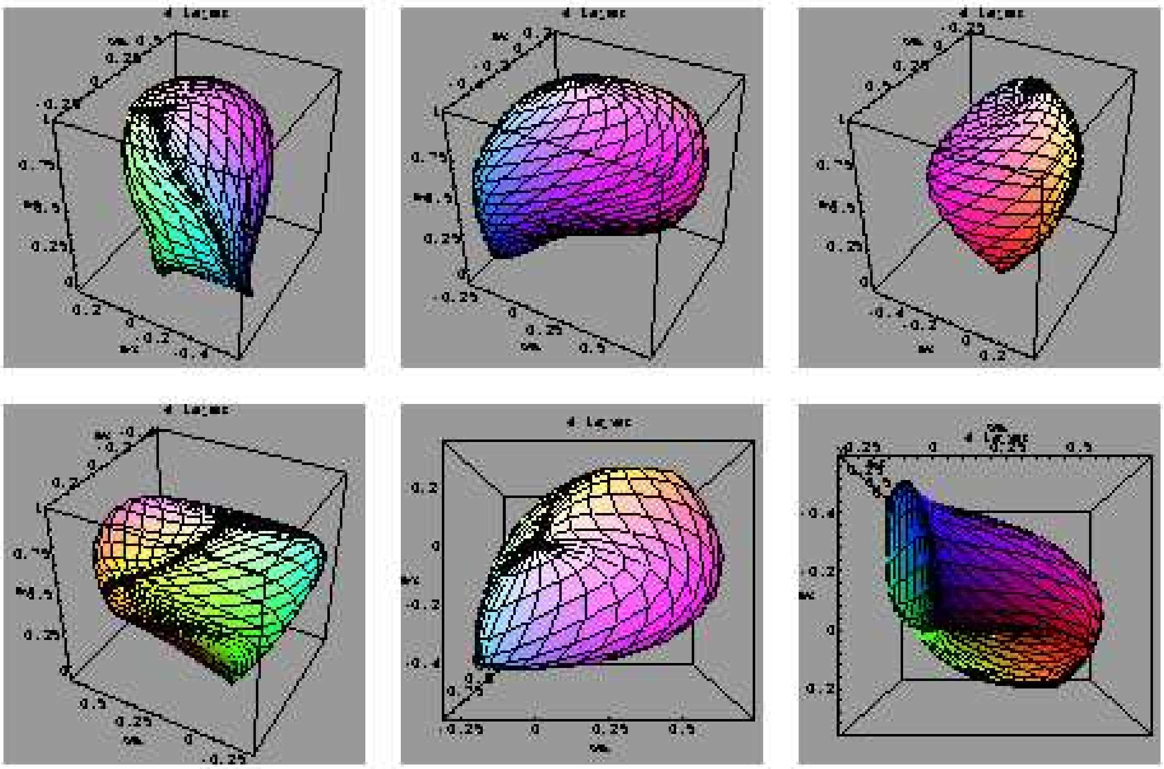

Figure 8: The OCS surface as computed via the 4-layer network XYZ to NPP transform (see text). Compare to Figure 4 (p. 7), which shows the directly computed version.

[5] J. M. Lammens. *A Computational Model of Color Perception and Color Naming.* PhD thesis, State University of New York at Buffalo, 1994. Also available as TR-94-26, CS dept., SUNY/Buffalo.

[6] K. N. Leibovic. Vertebrate photoreceptors. In K. N. Leibovic, editor, *Science of Vision*, chapter 2, pages 16–52. Springer, New York, 1990.

[7] L. T. Maloney. Color constancy and color perception: The linear-models framework. In D. E. Meyer and S. Kornblum, editors, *Attention and Performance XIV*, chapter 3, pages 60–78. MIT Press, 1993.

[8] L. T. Maloney and B. A. Wandell. Color constancy: a method for recording surface spectral reflectance. *Journal of the Optical Sociey of America*, pages 29–33, 1986.

[9] D. F. Rogers. *Procedural elements for computer graphics.* McGraw-Hill, New York, 1985.

[10] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error propagation. In D. Rumelhart, J. McClelland, and the PDP Research Group, editors, *Parallel Distributed Processing*, chapter 1, pages 319–362. MIT Press, Cambridge MA, 1986.

[11] G. Wyszecki and W. S. Stiles. *Color Science.* Wiley, New York, 1982.

[12] S. Zeki. *A Vision of the Brain.* Blackwell Scientific, Oxford, 1993.