Four-Transition 0-1 Functions for Reflectance Spectra

Paul Centore

Abstract

It is known that a reflectance spectrum for an optimal colour takes on the value 0 or 1 at every wavelength, with at most two transitions between those values. This paper shows that any non-optimal colour can be produced by a reflectance spectrum that takes on the value 0 or 1 at every wavelength, with at most four transitions. While the two-transition optimal spectrum is unique, the four-transition non-optimal spectrum is not unique.

1 Introduction

For any illuminant I, there exists an object-colour solid[1] Ω_I, which is a subset of the three-dimensional XYZ colour space introduced by the Commission Internationale de l'Éclairage (CIE).[2] The object-colour solid for I consists of all the XYZ coordinates that can be produced when a human views a surface that is illuminated by I. A surface has a *reflectance spectrum*, which is a function that specifies, for each visible wavelength, i.e. those between 400 and 700 nm, the percentage of incoming light of that wavelength that the surface reflects. A reflectance spectrum takes on values between 0 and 100%, or, more simply, 0 and 1. Typically, many different spectra, making up a *metamer set*, can produce the same XYZ when illuminated by I.

The colours on the boundary of Ω_I are called *optimal colours*. A well-known result,[3,4] referred to here as the Optimal Colour Theorem, states that an optimal colour is produced only by a reflectance spectrum in *Schrödinger form*: the spectrum takes on only the values 0 and 1, and there are at most two transitions from 0 to 1 or vice versa. For convenience, a function that takes on only the values 0 or 1 will be said to have *0-1 form*, and the standard term *support* will denote the set of wavelengths where a spectrum has non-zero reflectance. To provide a further simplification, a three-dimensional geometric interpretation suggests joining the ends of the spectrum to produce a circular domain of wavelengths; on the circular domain, the support of an optimal spectrum is an interval, and, excepting ideal black and ideal white, each optimal spectrum has exactly two transitions. The theorem includes a converse: any spectrum in Schrödinger form must produce an optimal colour. Further work[5-7] has shown that the reflectance spectrum for an optimal colour is unique.

A natural question is whether any similar simple forms occur for non-optimal colours, those that are inside the object-colour solid. This situation is more complicated because any non-optimal colour can be produced by a large metamer set of varied reflectance spectra.

This paper will show, however, that any non-optimal colour can be produced by a 0-1 reflectance spectrum with *four* transitions, as opposed to the two that characterize optimal colours.

Four-transition spectra can arise naturally by adding optimal spectra. Two optimal spectra whose supports do not intersect will be called *disjoint*. Adding two disjoint optimal spectra produces a 0-1 function whose support consists of two intervals. This spectrum has four transitions, one at the endpoint of each interval. Given a non-optimal colour, in the interior of an object-colour solid, we will find two disjoint optimal colours, on the boundary of the solid, whose vectors in XYZ space sum to the vector of the non-optimal colour in that space. It then follows that the sum of the reflectance spectra, which is a four-transition 0-1 function, must be metameric to the non-optimal colour, proving the result.

The selection of the pair of disjoint optimal colours, which is not unique, uses the zonoidal structure of object-colour solids.[8,9] Since zonoids are convex and centrally symmetric, the section through the center that contains the origin, the white point, and the non-optimal colour, is also convex and centrally symmetric. The colours on the boundary of the section must be optimal colours on the original object-colour solid, and the non-optimal colour is interior to the section. The symmetry will allow pairs of disjoint optimal colours to be characterized geometrically, and a construction involving translations will show that any vector in the interior is the sum of two vectors, given by disjoint colours, on the boundary.

The paper is organized as follows. After the introduction, the perception of object or surface colours is formulated geometrically. Object-colour solids are of particular interest; their structure as zonoids, and the inheritance of that structure by sections of the solid, are described. Reflectance spectra are expressed in a circular form, which simplifies the counting of transitions, and implies that any optimal spectrum (excepting black and white) is a 0-1 function with exactly two transitions. An expression of a four-transition 0-1 function as the sum of two disjoint two-transition 0-1 functions is given. The geometric constructions involving a section of an object-colour solid are then used to write an arbitrary non-optimal colour as the vectorial sum of two disjoint optimal colours in XYZ space, from which the main result of the paper follows: the metamer set of an arbitrary non-optimal colour contains a four-transition 0-1 reflectance spectrum. Possible applications are discussed, and the paper closes with a summary.

2 Geometric Formulation

2.1 Reflectance Spectra

An object's colour is perceived when light reflects off that object and enters a viewer's eye. The light must be in the visible spectrum, with wavelengths varying from about 400 to 700 nm, and can be described by a spectral power distribution (SPD), which is a function of the visible spectrum that gives the power of the light at each wavelength. An *illuminant* is an SPD that gives the relative power of a light source at various wavelengths. Illuminants are used rather than absolute SPDs, because the intensity of a light typically varies over a scene, while its SPD's shape remains the same.

When light strikes an object, some of that light is reflected, and the rest is absorbed or transmitted, the proportions of each outcome varying with wavelength. An object's *reflectance spectrum* gives the percentage of light that that object reflects, as a function of the visible wavelengths. A spectrum is bounded between 0 and 100%, or equivalently, between 0 and 1. An object modifies the SPD of incoming light, by reflection in accordance with its spectrum, to produce a new SPD, which arrives at an observer.

The CIE 1931 Standard Observer[2] systematizes human perception of an incoming SPD. This standard gives three colour-matching functions ($\bar{x}(\lambda)$, $\bar{y}(\lambda)$, and $\bar{y}(\lambda)$), each of which is linear on the set of SPDs. The function's outputs are three non-negative CIE coordinates, denoted X, Y, and Z, which can be seen as a vector in $\mathbb{R}^3$. A viewer will perceive that two SPDs have the same colour if and only if their XYZ coordinates are equal. Two reflectance spectra that produce the same colour perception are called *metamers*, and the set of all reflectance spectra that produce the same XYZ when viewed under a certain illuminant I are called the *metamer set* of XYZ for I.

Helpful constructions result by considering maximal monochromatic reflectance spectra (MMRSs), which reflect 100% of the light at one wavelength (or in a narrow band around that wavelength), and 0% of the light anywhere else. The XYZ vector of an MMRS is called a *spectrum locus vector*, and the set of all such vectors, or sometimes of the continuous curve joining their tips, is called the *spectrum locus*. Since the XYZ vector is zero at either end of the visible spectrum, the spectrum locus is topologically a circle; equivalently, the two ends of the visible spectrum can be joined to produce a circle, which serves as the domain for a reflectance spectrum; this interpretation will be used later in this paper. The direction of the locus vector for a particular wavelength is fixed by the colour-matching functions, but its magnitude varies with the illuminant. A particular reflectance spectrum ρ can be seen as a linear combination, with coefficients between 0 and 1, of MMRSs, and will produce a set of CIE coordinates XYZ_ρ. Since the colour-matching functions are linear, XYZ_ρ is likewise a linear combination, with the same coefficients, of the spectrum locus vectors corresponding to those MMRSs.

2.2 Object-Colour Solids

Not all XYZs are possible, because there might not exist an SPD which produces a particular XYZ. Furthermore, not all XYZs involve an object: an observer can also view a light source directly, without any intermediate reflection. The set of all possible XYZs that do result when an illuminant I strikes an arbitrary object is called the *object-colour solid* Ω_I for I, and is a subset of $\mathbb{R}^3$. Even though different illuminants lead to different solids, all object-colour solids share some geometric features. For instance, they all start at the origin, which corresponds to the reflectance spectrum that is identically 0; an object with such a spectrum would reflect no light at all, so it is called an *ideal black*. Each solid terminates at a white point, corresponding to the reflectance spectrum that is identically 1. The origin and white point are extreme vertices of the solid, in between which it is convex, and centrally symmetric about the XYZ that results from the reflectance spectrum that is 0.5 at every wavelength. By linearity, this XYZ lies on the line segment between the origin and the white point, and is halfway between them.

The geometry of object-colour solids admits considerably more development: they are actually zonohedra (or zonoids, in the continuous limit) generated from the spectrum locus vectors.[8,9] To visualize this construction, start with the standard set of 31 wavelengths (from 400 to 700 nm, in increments of 10) and an illuminant I with SPD $I(\lambda)$. When illuminated by I, the MMRS at wavelength λ, denoted $M(\lambda)$, produces the spectrum locus vector

$$s(\lambda) = I(\lambda)\,[\bar{x}(\lambda), \bar{y}(\lambda), \bar{z}(\lambda)] \tag{1}$$

in three-dimensional CIE XYZ space. The set

$$\mathcal{Z} = \left\{ \sum_{\lambda} \rho(\lambda) s(\lambda) \middle| \rho(\lambda) \in [0,1] \forall \lambda \right\} \tag{2}$$

is called the *zonohedron* generated by the spectrum locus under illuminant I. In fact, that zonohedron *is* the object-colour solid:

$$\Omega_I = \mathcal{Z}, \tag{3}$$

to a 31-wavelength approximation. The coefficients in Equation (2) can be seen as the values of a discrete reflectance spectrum ρ over the 31 wavelengths. In fact, the element of $\mathcal{Z}$ corresponding to the sequence of coefficients ρ is just the XYZ vector produced by the reflectance spectrum ρ when viewed under illuminant I. The black point occurs at the origin when the coefficients of ρ are all 0, and the white point occurs when they are all 1.

Finer approximations than 31 wavelengths are possible. With 61 wavelengths, for instance, spaced at 5 nm instead of 10, there would be twice as many spectrum locus vectors, but, because the power $I(\lambda)$ would apply to an interval of 5 nm instead of 10, each vector would be about half the magnitude. The white point would sum up 61 vectors of about half their previous magnitude, so would remain in about the same place. Both object-colour solids would have approximately the same shape, although the solid from 61 wavelengths would have many more vertices, and they would be more closely spaced. In the limit, as the number of wavelengths increases without bound, and the corresponding reflectance spectra become continuous, the solid's vertices (apart from the cone-like vertices at the origin and white point) are smoothed away, and the solid is called a *zonoid*.

2.3 Optimal Colours

As mentioned earlier, a particular XYZ can be produced by a metamer set of reflectance spectra. Since the set of reflectance spectra has many degrees of freedom, while CIE XYZ space only has three, metamer sets are often very large: a surprisingly wide variety of very different spectra can produce the same colour. An important exception to this case occurs for the colours on the boundary of an object-colour solid, which are known as *optimal colours*. The Optimal Colour Theorem, which Schrödinger[3] stated and partially proved in 1920, and MacAdam[4] completely proved in 1935, states that the reflectance spectrum for an optimal colour must have *Schrödinger form*: it takes on only the values 0 and 1, with at most two transitions between those values.

The zonohedral/zonoidal interpretation provides a proof of this theorem.[8] A point of the zonohedron resulting from a set of generating vectors is a vertex if and only if it is the sum of all the generating vectors lying on one side of a plane through the origin. The vertices of an object-colour solid, seen as the zonohedron generated from the spectrum locus vectors, delineate the solid's boundary, and thus the optimal colours. Since the spectrum locus vectors are convex and well-ordered in wavelength,[10] the vectors on one side of a plane result from an interval of wavelengths (when the visible spectrum wraps around to form a circle). The corresponding reflectance spectrum is 1 for the vectors on one side of the plane, and 0 elsewhere, so the spectrum has Schrödinger form. When 31 wavelengths are used, and the solid is a zonohedron, the reflectance spectrum can only transition from 0 to 1 or vice versa at the wavelengths 400, 410, 420, etc. Furthermore, an optimal colour corresponds to a point on the boundary, but that point might not be a vertex; when it is not a vertex, the reflectance spectrum might take on a value strictly between 0 and 1 during a transition. These difficulties disappear in the zonoidal limit, when reflectance spectra are continuous rather than discrete. Then optimal reflectance spectra can have their transitions at any wavelength, and all values are either 0 or 1.

The circular interpretation of the spectrum locus suggests a simplified method of counting transitions, which this paper will use. Schrödinger[3,11] classified the optimal colour spectra into the four forms shown in Figure 1, depending on the number and locations of transition wavelengths. The first form, in the upper left, has one transition, and takes on the value 1 between 400 nm and that transition. The second form, in the upper right, also has one transition, but takes on the value 1 to the right of the transition. The third form, in the lower left, has two transitions, between which it is 1, and outside which it is 0. The fourth form also has two transitions, but is 0 between them, and 1 outside.

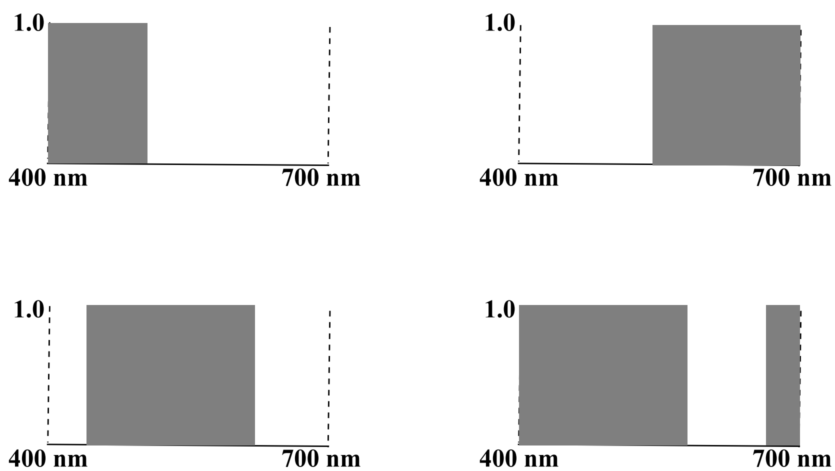

Figure 1: Four Forms for Optimal Colour Functions

These four forms can be conveniently subsumed into one form, by joining the ends of the visible spectrum, at 400 and 700 nm, to form a circle. On the circle, the reflectance function for any optimal colour fills out a sector of an annulus, with one transition wavelength where the sector starts, and another where the sector ends. In the first form, the sector starts at 400 nm (which is also 700 nm). In the second form, the sector ends at 400 (or 700) nm. In the third form, the sector does not cross the 400 nm mark. In the fourth form, the sector starts before 700 nm, continues clockwise across 700 nm, which is also 400 nm, and then fills out some of the clockwise side of the 400 nm mark. An optimal reflectance spectrum takes on value 1 at wavelengths whose spectrum locus vectors are summed up to produce that optimal colour, and the Optimal Colour Theorem says that those wavelengths must lie along an interval of the circular spectrum locus.

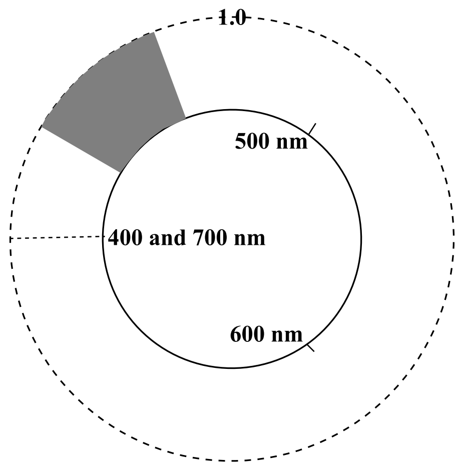

Figure 2: Circular Form for Optimal Colour Functions

In the circular form, any of the four cases in Figure 1 results in exactly two transitions. Equivalently, we could count an extra transition in Figure 1 whenever the spectrum is 0 at the left end and 1 at the right end, or vice versa. With this understanding, it is impossible for a spectrum to have exactly one transition, or indeed any odd number of transitions. Only the spectra for ideal white and ideal black have zero transitions. If we leave aside these two special cases, then the Optimal Colour Theorem says that a reflectance spectrum for an optimal colour must be a 0-1 function with exactly two transitions.

The geometry shows that the converse of the Optimal Colour Theorem also holds: any reflectance spectrum in Schrödinger form produces an optimal XYZ, on the boundary of Ω_I. Further work has strengthened this theorem to include uniqueness: exactly *one* reflectance spectrum produces a particular optimal colour. (The history of the proof of uniqueness is unclear. In a footnote to a 2009 paper,[6] A. Logvinenko refers to a proof in an unpublished manuscript[7] that was still under preparation. An explicit proof of uniqueness,[5] which relies on the empirical finding that no three locus vectors are linearly dependent, was given in 2012.) Uniqueness implies that the metamer set of an optimal colour consists of a single reflectance spectrum.

A natural question is whether similar characterizations exist for non-optimal colours, those that are interior to Ω_I. This paper will provide such a characterization: every metamer set for a non-optimal colour contains at least one 0-1 reflectance spectrum with exactly four transitions. Unlike in the optimal case, however, the four-transition 0-1 spectrum is not unique, and, of course, there are many more metamers that take on values besides 0 and 1.

2.4 Sections of an Object-Colour Solid

The proof of the four-transition 0-1 result will consider not just object-colour solids, but also sections of those solids, in particular the sections given by intersecting the solid with a plane that contains both the origin and the white point. Since such a plane contains the central point, the section must be centrally symmetric. Since the solid and the cutting plane are both convex, the section is also convex. Figure 3 shows an example, with the origin labeled 0 and the white point labeled w. The center is also indicated. For convenience, the boundary of the section can be divided into upper and lower parts.

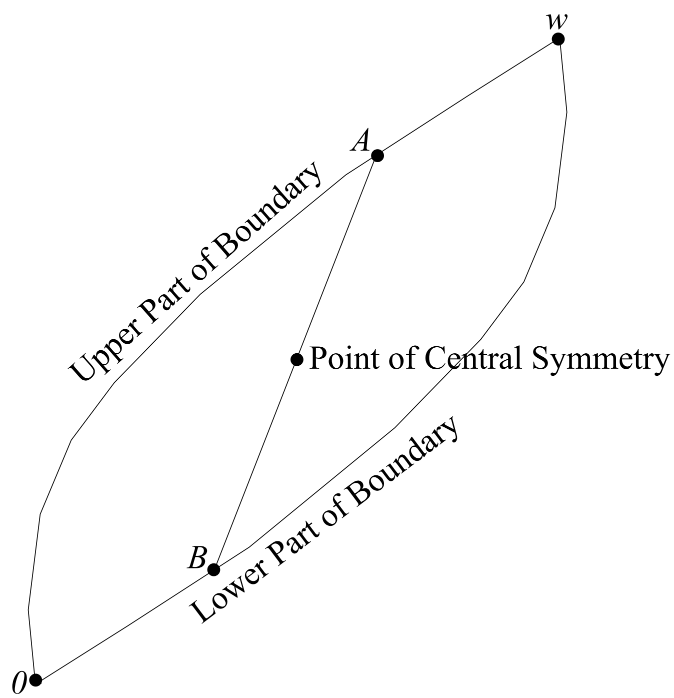

Figure 3: A Section of an Object-Colour Solid, Containing the Origin and White Point

The points on the boundary of the section are also on the boundary of the object-colour solid, so the Optimal Colour Theorem says that they have unique reflectance spectra in Schrödinger form. At the origin, of course, the spectrum is identically 0 (i.e. its support is empty), while at the white point it is identically 1 (i.e. its support consists of all visible wavelengths). As one moves along either part of the boundary from 0 to w, the support must always be an interval (allowing for wraparound), and it expands smoothly, sometimes to the left and sometimes to the right and sometimes to both left and right simultaneously, until it fills the spectrum from 400 to 700 nm. The support at any one point is always a proper subset of the support at a later point on that part of the boundary.

Though the starting and ending intervals for the upper and lower boundaries are the same, the intermediate intervals can be very different. In fact, the symmetry implies a complementary relationship. Let A and B be two diametrically opposite boundary points on the upper and lower parts, as shown in the figure. Since they are symmetric about the center, the average of their two reflectance spectra, ρ_A and ρ_B, must be identically 0.5, and their sum must be identically 1, which is the spectrum for ideal white:

$$\rho_w = \rho_A + \rho_B. \tag{4}$$

The reflectance spectra for A and B are complements: ρ_A is 1 wherever ρ_B is 0, and 0 wherever ρ_B is 1. A corollary is that ρ_B is *disjoint* from any point between 0 and A, where two spectra are said to be disjoint if their supports have empty intersection.

3 Proof of Main Result

This section will prove the main result of this paper: the metamer set of any non-optimal colour contains a four-transition 0-1 reflectance spectrum. First, a four-transition 0-1 function will be constructed as the sum of a pair of disjoint Schrödinger functions. Then, such a pair will be found geometrically, in a section of the object-colour solid, that sum to an arbitrary non-optimal colour. The result follows.

3.1 Constructing Four-Transition 0-1 Functions

Suppose we have two disjoint optimal colours, γ_1 and γ_2. Then the supports of their reflectance spectra ρ_1 and ρ_2 are two non-intersecting intervals, and their sum is a third spectrum:

$$\rho = \rho_1 + \rho_2. \tag{5}$$

Note that the sum of two reflectance spectra is not generally another reflectance spectrum, because the sum could result in values that are greater than 1; the non-intersecting intervals avoid that problem in this case. The non-intersection also implies that ρ must have exactly four transitions, and of course ρ only takes on the values 0 and 1.

Geometrically, γ_1 and γ_2 are two XYZ vectors in $\mathbb{R}^3$, on the boundary of Ω_I. They can be added to produce a third vector

$$\gamma = \gamma_1 + \gamma_2. \tag{6}$$

By linearity, the reflectance spectrum ρ must be in the metamer set of γ. Since ρ has more than two transitions, the Optimal Colour Theorem implies that γ cannot be on the boundary of Ω_I. Since ρ is a physically possible object colour, however, γ must belong to Ω_I, and is therefore a non-optimal, interior colour, for which Equation (6) provides a four-transition 0-1 function.

3.2 Finding Disjoint Optimal Colours

Rather than starting with optimal colours, this section starts with a non-optimal colour γ in the interior of Ω_I, and proves the existence of a four-transition 0-1 metamer for γ by finding disjoint optimal colours γ_1 and γ_2 that sum to γ. Begin by constructing the plane $\mathcal{P}$ which contains γ, the origin, and the white point. (If γ is on the neutral line joining the origin and the white point, then there are many such planes; choose any one.) Construct the section $\Omega_I \cap \mathcal{P}$, which looks like Figure 3; the point γ is in the interior of the section.

The desired vectors γ_1 and γ_2 must satisfy Equation (6). They are optimal, so they are on the boundary of Ω_I, and will also (at least in the zonoidal limit) appear on the boundary of the section. They will be chosen to be disjoint; one will be on the lower boundary of the section, and the other will be on the upper boundary. A geometric test will tell us whether a given γ_1 on the lower boundary is disjoint from a γ_2 on the upper boundary. Start with γ_1 as shown in Figure 4, and construct the complementary colour $\bar{\gamma}_1$. Considering the colours as vectors, we must have

$$w = \gamma_1 + \bar{\gamma}_1, \tag{7}$$

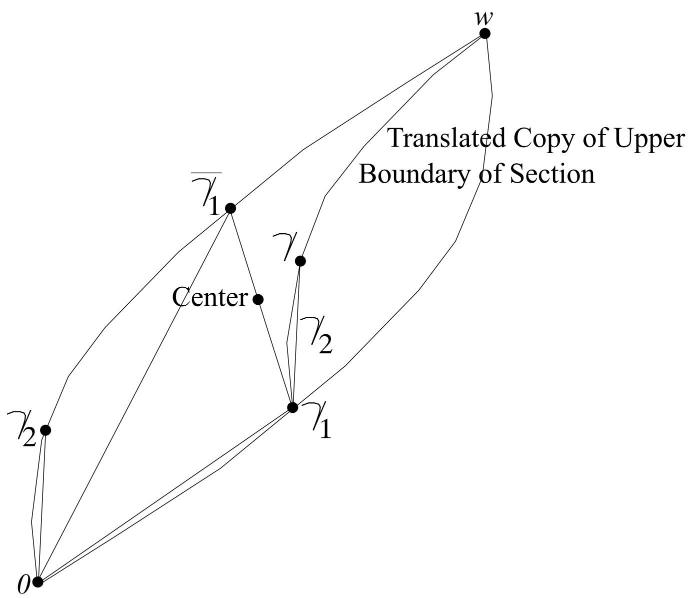

Figure 4: Disjoint Optimal Spectra that Sum to γ

while considering the colours as reflectance spectra gives

$$\rho_w = \rho_1 + \bar{\rho}_1, \tag{8}$$

where ρ_w is the ideal white spectrum, which is identically 1.

Now consider any point γ_2 that is on the upper boundary between 0 and $\bar{\gamma}_1$. Then γ_2 is also optimal, and, as shown earlier, its support must be a proper subset of the support of $\bar{\gamma}_1$. The sum $\gamma_1 + \gamma_2$ is therefore a four-transition, 0-1 function which is contained in the section. As γ_2 varies from 0 to $\bar{\gamma}_1$, the set of sums $\gamma_1 + \gamma_2$ will trace out a translated copy of the part of the upper boundary between 0 and $\bar{\gamma}_1$; the copy starts at γ_1 and ends at w, given by $\gamma_1 + \bar{\gamma}_1$. Any point on this copy can be written as a four-transition, 0-1 function.

The final step in the construction is now straightforward: if one pictures the upper boundary sliding along the lower boundary, it is clear that the entire section will be swept out, and in particular any desired γ inside the section will sooner or later be reached. When a desired γ is reached, then choose as γ_1 the point where the translated copy originates on the lower boundary, and let γ_2 be the difference of γ and γ_1. These optimal colours are disjoint and sum to γ as was to be shown.

While proving the existence of a four-transition, 0-1 metamer, the construction should also make it clear that these metamers are not unique. A simple example of non-uniqueness would involve the center point itself, or indeed any colour on the line segment between 0 and w. There are many choices of a section for these cases, and any section will give a pair of optimal colours that do not occur in the other sections. The sums of these pairs will therefore also be distinct as reflectance spectra, although they all produce the same XYZ. For a point not on the central segment, one could choose a section that contained 0 and that point, but not w. Since many of the sums of pairs of optimal colours on the boundary of that section would still be within the colour solid, they would provide alternate four-transition 0-1 metamers for many non-central colours. In general, then, four-transition 0-1 reflectance spectra would not be unique for non-optimal colours.

4 Discussion

This section explores the limits of the current paper, as well as possible applications. The major limitation is that the current results are based implicitly on human vision, and might not apply to artificial sensors. Optimal colours have two (or zero) transitions because the human spectrum locus is convex and well-ordered in wavelength.[10] A sensor's spectrum locus, that can be defined analogously to the human spectrum locus, might or might not be convex, and its wavelengths might or might not be well-ordered in the locus. As a result, sensors must be analyzed on a case-by-case basis to determine applicable forms for them.

Various applications use different models for reflectance spectra, including 0-1 functions with some number of transitions. In this context, it should be recognized that the four transitions needed for non-optimal colours are a minimum; 0-1 functions with more transitions can also serve as metamers. In fact, the set of 0-1 functions with six or more transitions is dense in the set of four-transition 0-1 functions, in much the same way that the rational numbers are dense in the real numbers: given a four-transition function ρ_4, there always exists a six-transition function ρ_6 that is as close as desired to ρ_4. To construct ρ_6 from ρ_4, take a very narrow interval where ρ_4 is 1, and set an arbitrarily small middle segment of it to 0 (or vice versa from 0 to 1). The middle segment will introduce two new transitions at its endpoints, increasing the total from four to six. By choosing a sufficiently small segment, the resulting 0-1 function can be made as close to ρ_4 as desired. This result extends easily to an arbitrary number of transitions. In general, the set of M-transition 0-1 spectra is dense in the set of N-transition 0-1 spectra, whenever $M > N$.

This density relationship implies that, for numerical computations, there is no harm in allowing more transitions than is strictly necessary, because the higher-transition functions can approximate the lower-transition functions as closely as desired. On the other hand, this paper shows that every higher-transition function is metameric to a four-transition function, so, in at least some cases, numerical computations can be reduced by using only four transitions, with no loss of accuracy.

This observation applies to the problem which indirectly motivated this paper: constructing metamer mismatch bodies (MMBs), which are convex subsets of an object-colour solid.[1] While the boundary of an MMB consists of XYZ points, it is more natural to find extreme reflectance spectra that correspond to boundary XYZs, rather than the XYZs themselves. A standard approach[12-14] to MMB construction is to model a reflectance spectrum as a discrete function over 31 wavelengths, and formulate a linear programming (LP) problem. LP theory then implies that any extreme reflectance spectrum is 0 or 1 everywhere, except on at most three of its 31 wavelengths.

One could naturally object that the 31-wavelength functions are only a subset of the set of all reflectance functions, and that one should consider arbitrary reflectance spectra. This objection motivated Logvinenko, Funt, and Godau[15] to produce a non-LP algorithm for MMB construction that allowed for a wider set of spectra. In particular, they used the set of five-transition (which would have six transitions by this paper's method of counting) 0-1 functions, rather than the discrete wavelength model. The choice of 0-1 functions was motivated by the fact that extreme reflectance spectra are on the boundary of the six-dimensional concatenation of the object-colour solids for two illuminants. Arguments analogous to those used to prove the Optimal Colour Theorem then imply that extreme functions are 0-1. (An alternate route to the same conclusion is to increase the number of wavelengths in the discrete approximation without bound, from 31 to 61 to 301, and so on. Regardless of the number, LP implies that at most three wavelengths have reflectances that are not 0 or 1. In the limit, the proportion of non-0-1 wavelengths becomes vanishingly small, so an extreme spectrum must be 0-1.)

The number of transitions required for these 0-1 functions, however, is somewhat uncertain. A theorem from an unpublished reference[7] associates boundary points of the six-dimensional set with hyperplanes in six-dimensional space: barring special cases, a hyperplane corresponding to a boundary point will intersect the spectrum locus at at least five wavelengths. The associated reflectance spectrum will assign maximal values of 1 to all spectrum locus vectors on one side of the hyperplane, and minimal values of 0 to vectors on the other side; transitions between 0 and 1 occur at those wavelengths where the hyperplane intersects the locus. (Although the authors do not use those terms, the six-dimensional set is in fact a zonoid, and the hyperplane is a higher-dimensional version of our plane through the origin in $\mathbb{R}^3$. The reflectance spectrum that is 1 for all the wavelengths whose locus vectors are on one side of the plane, and 0 elsewhere, is then an optimal, or boundary, spectrum.) The set of hyperplanes is five-dimensional, suggesting that five parameters, in some form or other, are needed to describe the set of boundary points. The transition wavelengths for the 0-1 extreme spectra are a plausible candidate for these parameters, and in fact produced good results.

The previous discussion suggests that four-transition 0-1 functions, rather than five-transition 0-1 functions, might have been sufficient for constructing MMBs. The result is not certain, of course, because many MMBs involve two illuminants instead of one. Nevertheless, this example indicates a possible application for the results of the current paper.

5 Summary

In analogy to the two-transition 0-1 Schrödinger form for optimal colours, the current paper has shown that every non-optimal colour is metameric to the colour produced by a four-transition 0-1 reflectance spectrum. Although the Schrödinger form is unique for optimal colours, the four-transition form is not unique for non-optimal colours. The demonstration of these results relied on the zonoidal structure of object-colour solids, and used the geometrical properties of sections of those solids to express an arbitrary non-optimal colour as the sum, in CIE XYZ space, of two optimal colours with non-intersecting support. The corresponding algebraic sum of the optimal reflectance spectra is then a four-transition 0-1 spectrum as desired. Potential use of this form was discussed.

References

1. G. Wyszecki and W. S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulae, 2nd ed.*, John Wiley & Sons, 1982.
2. Deane B. Judd, "The 1931 I. C. I. Standard Observer and Coordinate System for Colorimetry," JOSA, Vol. 23, pp. 359-374, October 1933.

3. Erwin Schrödinger. "Theorie der Pigmente von grösster Leuchtkraft," Annalen der Physik 4, 62, 1920, pp. 603-622. Translated in 2010 by Rolf G. Kuehni under the name "Theorie [sic] of pigments of greatest lightness."
4. David L. MacAdam. "The Theory of the Maximum Visual Efficiency of Colored Materials," JOSA, Vol. 25, 1935, pp. 249-252.
5. Paul Centore, "Non-Metamerism of Boundary Colors in Multi-Primary Displays," Journal of the Society for Information Display, Vol. 20, Issue 4, pp. 214-220, Apr 2012.
6. A. D. Logvinenko, "An object-color space" (see footnote 1), Journal of Vision (11):5, pp. 1-23, 2009.
7. Alexander D. Logvinenko & V. I. Levin, *Foundations of Colour Science*, unpublished manuscript.
8. Paul Centore, "A Zonohedral Approach to Optimal Colours," Color Research & Application, Vol. 38, No. 2, pp. 110-119, April 2013.
9. Paul Centore, "Geometric Invariants Under Illuminant Transformations," Color Research & Application, Vol. 39, No. 2, pp. 179-187, April 2014.
10. Gerhard West & Michael H. Brill. "Conditions under which Schrödinger object colors are optimal," JOSA, Vol. 73, No. 9, September 1983, pp. 1223-1225.
11. David L. MacAdam. *Color Measurement: Theme and Variations* (see Sect. 7.5), Springer Series in Optical Science, Vol. 27, Springer-Verlag, 1981.
12. Eugene Allen, "Some New Advances in the Study of Metamerism: Theoretical Limits of Metamerism; An Index of Metamerism for Observer Differences," Color Engineering, Vol. 7, No. 1, pp. 35-40, January-February 1969.
13. N. Ohta & G. Wyszecki, "Theoretical chromaticity-mismatch limits of metamers viewed under different illuminants," Journal of the Optical Society of America, Vol. 65, No. 3, pp. 327-333, March 1975.
14. F. J. M. Schmitt, "A method for the treatment of metamerism in colorimetry," Journal of the Optical Society of America, Vol. 66, No. 6, pp. 601-608, June 1976.
15. Alexander D. Logvinenko, Brian Funt, & Christoph Godau, "Metamer Mismatching," IEEE Transactions on Image Processing, Vol. 23, No. 1, pp. 34-43, January 2014.

