Optimal Spectra for Double Object-Colour Solids

Paul Centore

Abstract

Optimal colours for human vision occur on the boundary of a three-dimensional object-colour solid, and result from optimal reflectance spectra, that take on only the values 0 and 1, with at most two transitions between those values. Different illuminants lead to different solids. If there are two illuminants and a single sensing device, then we can construct a six-dimensional double object-colour solid by concatenating colour signals from both illuminants. Colours on the boundary of a double-object solid, and the spectra that generate them, can also be called optimal. This paper shows that, while optimal spectra for double solids take on only the values 0 and 1, there is no maximum number of transitions between those values: given a device, we can always construct two illuminants such that the resulting double object-colour solid has an optimal reflection spectrum with as many transitions as desired.

1 Introduction

Object-colour solids Ω_I, consisting of all the colour outputs that physical objects can produce by reflecting a given illuminant I, occur when studying human colour vision. Formally, a physical object has some *reflectance spectrum* ρ, which is a function with values between 0 and 100%, or equivalently 0 and 1, over the visible spectrum; $\rho(\lambda)$ gives the percentage of incoming light of the wavelength λ that the object reflects. A *monochromatic reflectance spectrum* ρ_λ is 1 at λ (or more commonly, is 1 on the waveband containing λ in some discretization of the visible spectrum), and 0 elsewhere. For a given illuminant I, the *colour signal map* Φ, from the set χ of all reflectance spectra, to the colour output space, which is a subset of $\mathbb{R}^3$, sends ρ to its colour output, when illuminated by I. Each coordinate of the colour output is given by a wavelength-dependent *response function*. The object-colour solid Ω_I for I is then the subset $\Phi(\chi)$ of $\mathbb{R}^3$. *Optimal colours* are those that occur on the boundary of an object-colour solid. In 1920, Erwin Schrödinger[1]showed that any optimal colour for human vision results from an *optimal reflectance spectrum* which takes on only the values 0 and 1 over the visible spectrum, with at most two transitions between those values.

If the human vision system (HVS) is considered as a special case of an imaging device, then the analogue of a human object-colour solid is an *illuminant gamut*, consisting of all the *RGB* outputs that the device produces when imaging physical objects under a given illuminant. An interesting situation occurs in computational colour constancy when a single

device images the same scene twice, but under two different illuminants. Then each illuminant has its own object-colour solid or illuminant gamut. To analyze this case, Eugene Allen (p. 38 of Ref. 2) suggested concatenating colour outputs from both illuminants into a six-dimensional hypervolume, which we will call a *double object-colour solid*. Formally, let Φ_1 and Φ_2 be the colour signal maps for I_1 and I_2. Then the range of the map $\Upsilon : \chi \longrightarrow \mathbb{R}^6$ given by

$$\Upsilon(\rho) = (\Phi_1(\rho); \Phi_2(\rho)) \tag{1}$$

is a sensor's double object-colour solid for I_1 and I_2.

Just like single object-colour solids, the reflectance spectra that produce boundary points on the double solid will be called *optimal spectra*. A recent paper by Logvinenko et al.[3] parametrized optimal spectra for a double object-colour solid with Schrödinger-like functions that took on only the values 0 and 1, but that allowed up to five transitions instead of two. While a reasonable modeling decision (since a six-dimensional solid has a five-dimensional boundary), a natural question is whether there is in fact some maximum number of transitions.

The current paper provides a negative answer. While it can be shown that only the values of 0 and 1 occur, it will also be proven that, for a given device, with a fine enough discretization of the visible spectrum, two illuminants can always be found for which the number of transitions required for some optimal reflectance spectrum is as large as desired.

The proof in this paper relies on the fact that the object-colour solid Ω_I is the Minkowski sum of the sensor's *spectrum locus vectors* under I,[4] where the spectrum locus vectors are $\Phi(\rho_\lambda)$ for the monochromatic spectra ρ_λ. This zonal structure leads to a simple, two-dimensional method of determining optimal reflectance spectra for an arbitrary sensor from its response functions. Extend all the locus vectors indefinitely, and cut them with a plane in $\mathbb{R}^3$, creating a *chromaticity diagram*. For convenience, the plane should cross all three axes at positive values. The resulting diagram is two-dimensional, and one can choose some coordinate system; the *chromaticity* of a point in the diagram is then identical with its coordinate pair.

For any reflectance spectrum ρ, $\Phi(\rho)$ is a vector in $\mathbb{R}^3$ that can be lengthened or shortened until it lies exactly on some point of the chromaticity diagram; the chromaticity of ρ is then defined to be the chromaticity of that point. If ρ_λ is monochromatic, then this paper will refer to the corresponding point in the chromaticity diagram as the *monochromatic point* for λ. Further geometric understanding results from considering an arbitrary reflectance spectrum as the sum of monochromatic reflectance spectra; then the linearity of Φ implies that the set of physically possible chromaticities is the convex hull of the monochromatic points. Unlike an object-colour solid, a device's chromaticity diagram depends only its response functions, and not on the illuminant.

The optimal spectra for an object-colour solid can be found directly[4] from the chromaticity diagram by the following procedure:

1. Indicate all the monochromatic points in the chromaticity diagram.
2. Draw an arbitrary straight line anywhere in the chromaticity diagram.
3. Identify all the monochromatic points on one side of that line, and their wavelengths (or, in a discretization, the wavebands containing those wavelengths).

4. The reflectance spectrum which takes on the value 1 on those wavebands, and the value 0 elsewhere, is optimal.

This construction is not just sufficient but also necessary: every optimal spectrum results from at least one such line.

In the horseshoe-shaped chromaticity diagram for human vision, the monochromatic points appear, ordered by wavelength, around the boundary. A line through the horseshoe divides those points into two sets. The reflectance spectrum that is 1 on the wavelengths on one side, and 0 on the wavelengths on the other side, is optimal. The complementary spectrum, that is 0 on the first side and 1 on the second side, is also optimal. Both these spectra have Schrödinger form, as they must.

This construction can be reinterpreted in terms of linear algebra. Since a chromaticity diagram is a section of $\mathbb{R}^3$, a line in the diagram can be extended to a unique plane through the origin in $\mathbb{R}^3$. The plane through the origin can be taken to be the kernel of a linear functional F, unique up to a multiplicative constant, on the three-dimensional colour space. The planar kernel divides the space into a positive side, where F takes on positive values, and a negative side, where it takes on negative values. The two sides of the line in the chromaticity diagram correspond to this positive/negative division of $\mathbb{R}^3$. The boundary of a zonohedron can be approximated by its vertices, and each vertex is the sum of all the locus vectors that fall on one side of a plane through the origin in the colour space.[5] The monochromatic points for those vectors also appear on one side of the line that results when the planar kernel intersects the chromaticity diagram, so the two- and three-dimensional constructions are equivalent.

Generalizing the constructions from a single object-colour solid to a double object-colour solid leads to a proof of the paper's main result. A double object-colour solid results from concatenating two ordinary solids for the same sensor, but for different illuminants, I_1 and I_2. The locus vectors for this map are still the images of monochromatic spectra, but now they are six-dimensional instead of three-dimensional. The double object-colour solid is the Minkowski sum of the six-dimensional locus vectors. A boundary point can be constructed by summing all the locus vectors on which a linear functional $F : \mathbb{R}^6 \longrightarrow \mathbb{R}$ takes positive values. The concatenation allows us to decompose F into two other functionals

$$F = F_1 + F_2, \tag{2}$$

where F_1 is 0 on the last three coordinates of $\mathbb{R}^6$, and F_2 is 0 on the first three coordinates. Then F_1 can be seen as a linear functional on the three-dimensional set of locus vectors for I_1, while F_2 is a linear functional on the three-dimensional set of locus vectors for I_2. The kernel of F_1 produces a line in the chromaticity diagram for I_1, and all the monochromatic points on one side of that line take on positive values; the points on the other side take on negative values. A similar statement holds for F_2 in the chromaticity diagram for I_2. Since the chromaticity diagrams for I_1 and I_2 are identical, however, we can draw both lines in one chromaticity diagram, and ask about the signs of the monochromatic points when F_1 and F_2 are summed to produce F. If a particular monochromatic point is positive for both F_1 and F_2, then it is also positive for their sum F. Similarly, a point that is negative for F_1 and F_2 is negative for F. These statements hold regardless of how strong or weak the illuminant is at the wavelength of interest.

A monochromatic point of wavelength λ, however, might be positive on F_1 and negative on F_2, or vice versa, in which case its value on F does depend on $I_1(\lambda)$ and $I_2(\lambda)$. The paper will show that choosing $I_1(\lambda)$ and $I_2(\lambda)$ appropriately allows us to make $F(\Upsilon(\rho_\lambda))$ positive or negative as desired. Since F is arbitrary, we can choose F_1 and F_2 as desired, and thus, assuming a fine enough discretization of the visible spectrum, produce a large number of monochromatic points that are positive on F_1 and negative on F_2 (or vice versa). By assigning appropriate values to the illuminants at the wavelengths for those points, we can produce an optimal reflectance spectrum that is 1 at some wavelength or waveband λ, 0 at the next waveband, then 1 at the waveband after that, and so on in alternating fashion. Such a spectrum can transition between 0 and 1 as many times as desired. This construction proves the paper's negative result: there is no maximum number of transitions for a double object-colour solid.

The paper is organized as follows. Minkowski sums and zonal structures are described first, to lay the mathematical groundwork. Next, standard geometric colour constructions such as spectrum locus vectors and colour signal maps are defined, and their relevant properties are stated. Double object-colour solids, the focus of this paper, are then defined, and their structure is related to previously established concepts. With these mathematical concepts in place, the paper's main result, that an optimal reflectance spectrum for a double object-colour solid has no maximum number of 0-1 transitions, is proven with an easily visualized two-dimensional chromaticity diagram. The paper concludes with a discussion of possible applications and a summary.

2 Zonotopes and The Minkowski Sum

2.1 The Minkowski Sum

The *Minkowski sum*, also called the *vector sum* or *linear sum*, of two sets $\mathbf{A}$ and $\mathbf{B}$ in $\mathbb{R}^n$, is defined as

$$\mathbf{A} \oplus \mathbf{B} = \{a + b | a \in \mathbf{A}, b \in \mathbf{B}\}. \tag{3}$$

This sum can be extended in the obvious way to any number of sets, and is readily seen to be commutative and associative, so that the order of the summands is immaterial. When all the summands in a Minkowski sum are convex, then the sum itself is also convex (Sect. 2.2.1 of Ref. 4).

Low-dimensional spaces offer appealing visual interpretations of Minkowski sums. In $\mathbb{R}^2$, for instance, $\mathbf{A}$ and $\mathbf{B}$ can be pictured as two flat shapes. The shape of their Minkowski sum, $\mathbf{A} \oplus \mathbf{B}$, is the shape consisting of all the area swept out by $\mathbf{A}$ when a fixed point a_1 of $\mathbf{A}$ slides over $\mathbf{B}$, in the sense that, for each b in $\mathbf{B}$, a translated copy of $\mathbf{A}$ is placed such that a_1 lies directly on top of b; the Minkowski sum is the union of all such copies. (While this method constructs the *shape* of $\mathbf{A} \oplus \mathbf{B}$, it will likely not give the correct *location* of $\mathbf{A} \oplus \mathbf{B}$, which could be very far from both $\mathbf{A}$ and $\mathbf{B}$.) A similar sliding construction works for polyhedra in $\mathbb{R}^3$, and indeed for analogous sets in any dimension. See Chap. 2 of Ref. 4 for some detailed examples.

2.2 Zonotopes

Zonotopes, known as *zonohedra* in three-dimensional space, are a special case of Minkowski summation, in which all the summands are line segments which start at the origin. Such a line segment can be naturally represented by the vector overlaying it, so a zonotope is a Minkowski sum of *generating vectors*, or just *generators*. To form the Minkowski sum of a set of the generators $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n$, choose a point a_1 in the first segment, another point a_2 in the second segment, and so on, and add them as vectors; the Minkowski sum is the set of all such sums of points. Since a_i belongs to the i^{th} line segment, which extends from the origin to the vector $\mathbf{v}_i$, we can write

$$a_i = \alpha_i \mathbf{v}_i, \tag{4}$$

for some α_i between 0 and 1. As a whole, the segments generate the zonotope $\mathcal{Z}$ given by

$$\mathcal{Z} = \left\{ \sum_{i=1}^{n} \alpha_i \mathbf{v}_i \middle| 0 \leq \alpha_i \leq 1 \ \forall i \right\}. \tag{5}$$

Even though line segments are one-dimensional sets, the zonotope they generate can, and likely will, be multi-dimensional. The Minkowski sum of two linearly independent vectors, for example, is a parallelogram, and the sum of three linearly independent vectors is a parallelepiped. The construction of zonotopes gives them some characteristic properties: a zonotope is a convex, centrally symmetric polytope. If all the generators have non-negative components, as will happen in our colour constructions, then the origin is a vertex; a terminal vertex, created by summing all the generators, is symmetric to the origin about the zonotope's center.

A vertex of a zonotope also has a special form:[5] it can be written as the sum of all generating vectors that lie on one side of a hyperplane $\mathcal{H}$ through the origin:

$$\text{vertex} = \sum_{i=1}^{N} \epsilon_i \mathbf{v}_i \middle| \epsilon_i = 1(\text{if } \mathbf{v}_i \text{ is on the given side of } \mathcal{H}) \text{ or } 0 \text{ (otherwise)}, \tag{6}$$

Conversely, suppose we have a hyperplane through the origin. Then the sum of all the generators on one side of the hyperplane is a vertex; the sum of the generators on the other side is also a vertex, that is diametrically opposite the first vertex. Rather than use a hyperplane, one can use a linear functional F, unique up to a multiplicative constant, whose kernel is that hyperplane: $F(\mathcal{H}) = 0$. Then all the generators on which F is positive lie on one side of the plane, and all the generators on which F is negative lie on the other side. Given any functional F, then, the sum of all the positive generators is a vertex, and the sum of all the negative generators is the diametrically opposite vertex. Formally,

$$\text{first vertex} = \sum \mathbf{v}_i \middle| F(\mathbf{v}_i) > 0, \tag{7}$$

$$\text{second vertex} = \sum \mathbf{v}_i \middle| F(\mathbf{v}_i) < 0, \tag{8}$$

Since such a functional can be found for any vertex, we have a convenient way of characterizing a zonotope's vertices, which will be used later in the paper.

3 Geometric Objects of Colour Science

In the setting of interest for this paper, a scene consisting of various physical objects is lit by some consistent illumination and imaged by a sensor with a three-dimensional output. The term "object" should be understood broadly, to include any physical surface; even a lawn, for instance, would be considered an object for our purposes. Likewise the term "sensor" or "device" should also be construed broadly, to include even the human visual system. This section will model the basic elements of the setting of interest as functions over the visible spectrum, and use the previous section to make some relevant zonohedral constructions.

3.1 Discretization

Colour applies to the visible electromagnetic spectrum, whose wavelengths vary from about 400 and 700 nm. Practical applications typically discretize the spectrum into wavebands of equal width. A common choice is 31 10-nm bands. A function of wavelength, such as a reflectance spectra or a light's spectral power distribution (SPD), is then defined by a table over the finite set of wavebands. A waveband is often indicated by a representative wavelength; for instance, the waveband between 435 and 445 nm could be represented by the wavelength 440 nm.

Though they calculate in discrete terms, workers often think in continuous terms. The continuous case is the limit of the discrete case as the number of wavebands increases without bound and wavebands become increasingly narrow. As we will see later, wavebands often lead to generating vectors, which produce zonotopes. As the discretization becomes finer, the number of generators increases, but the magnitude of each generator decreases, and the zonotope as a whole does not change shape much. In the limit, the zonotopes converge to a convex, centrally symmetric *zonoid*.[6] While a zonotope is a polytope with corners and edges, a zonoid can be smooth, with no corners or edges. This paper will consistently assume an evenly spaced discretization over the visible spectrum, although with arbitrary narrow wavebands, so only the discrete case is needed.

3.2 Illuminants

A physical light can be described mathematically by a *spectral power distribution* (SPD), which is a function that specifies the power contained in each waveband. For instance, an SPD might specify 5 Watts in the waveband between 575 and 585 nm, 7 Watts between 585 and 595, and so on. An SPD can be broken down into a shape and a magnitude. If one SPD is a scalar multiple of another, then their magnitudes differ, but their shapes, which define the relative amount of power in each waveband, are the same. An *illuminant* $I(\lambda)$ is a relative, rather than an absolute, SPD, and conveys shape rather than magnitude. Often a single light source, such as the sun, emits light with an SPD of a certain shape. The sun might light some objects frontally, while lighting other objects at more raking angles, so different objects are lit with different intensities. These differences, however, change the magnitude of the SPD, but not its shape, so an illuminant function is a natural descriptor. To avoid degenerate cases, this paper will make the further assumption, satisfied by natural lights, that all illuminants are positive over the visible spectrum,

3.3 Reflectance Spectra and Object Colours

Unlike a light source, a typical physical object reflects light produced elsewhere, rather than emitting light of its own. This paper deals with the common situation where light from some source falls on an object, reflects off that object, and then reaches a sensing device. The SPD of the light at the source is consistent with some illuminant $I(\lambda)$. Reflecting off the object modifies that SPD, and the modified SPD arrives at the sensor. A function ρ called the *reflectance spectrum* describes the modification. For each waveband, with a representative λ, the value $\rho(\lambda)$ is between 0 and 100%, or equivalently between 0 and 1, and gives the percentage of the incoming light, in that waveband, that the object reflects diffusely. The set of all reflectance spectra will be denoted χ. Mathematically, the modified SPD is $\rho(\lambda)I(\lambda)$, the pointwise product of an illuminant and a reflectance spectrum.

An arbitrary reflectance spectrum ρ can be written as a linear, indeed a zonal, combination of monochromatic reflectance spectra. For a given discretization, define

$$\rho_i(\lambda) = \begin{cases} 1, & \text{if } \lambda \text{ is in the } i^{\text{th}} \text{ waveband} \\ 0, & \text{otherwise.} \end{cases} \tag{9}$$

Then ρ can be written as

$$\rho = \sum_i \alpha_i \rho_i, \tag{10}$$

where each coefficient α_i is between 0 and 1. The set χ of all spectra is thus a zonotope, whose generating vectors are the monochromatic spectra.

An *object colour* is a colour output that can result when a light source reflects off an object. The set of object colours is a proper subset of the larger set of light-source colours. Two different illuminants I_1 and I_2 give rise to two different sets of object colours, which we will soon see are both zonohedra, called *object-colour solids*, in a sensor's output space.

3.4 Colour Signal Maps

Colour information results when light impinges on a sensing device, such as the HVS or a typical camera, which responds by producing a three-component vector, called a *colour signal*, in $\mathbb{R}^3$. In this paper, sensors be assumed to be linear, in the sense that each coordinate in the output vector is a linear transformation of the incoming SPD.

For human vision, the 1931 Standard Observer[7] defined by the Commission Internationale de l'Éclairage (CIE) specifies a colour signal with three output coordinates denoted X, Y, and Z, which are calculated from the CIE colour-matching functions $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, and $\bar{z}(\lambda)$, over the visible spectrum. When a light consistent with illuminant $I(\lambda)$ reflects off an object with reflectance spectrum $\rho(\lambda)$ and reaches a human eye, the CIE colour signal is given by:

$$X = \sum_\lambda \rho(\lambda)I(\lambda)\bar{x}(\lambda), \tag{11}$$

$$Y = \sum_\lambda \rho(\lambda)I(\lambda)\bar{y}(\lambda), \tag{12}$$

$$Z = \sum_\lambda \rho(\lambda)I(\lambda)\bar{z}(\lambda), \tag{13}$$

where λ takes on all the values in a set of representative wavelengths for some discretization. A camera's colour signal coordinates are usually denoted R, G, and B, for red, green, and blue, and are similarly calculated from three response curves $r_1(\lambda)$, $r_2(\lambda)$, and $r_3(\lambda)$, which vary from device to device:

$$R = \sum_{\lambda} \rho(\lambda) I(\lambda) r_1(\lambda), \tag{14}$$

$$G = \sum_{\lambda} \rho(\lambda) I(\lambda) r_2(\lambda), \tag{15}$$

$$B = \sum_{\lambda} \rho(\lambda) I(\lambda) r_3(\lambda). \tag{16}$$

In the limit, as the discretization becomes increasingly fine, the discrete sums in Equations (11) through (16) approach integrals over the visible spectrum.

Regardless of the sensor, the colour signal map for a particular illuminant can be written as a function $\Phi : \chi \to \mathbb{R}^3$, from the set χ of reflectance spectra to the sensor's colour output space, which is a subset of $\mathbb{R}^3$. For a given artificial sensing device and illuminant, the codomain $\Phi(\chi)$ of the set of all object colours is called an *illuminant gamut*. An illuminant gamut is a subset of $\mathbb{R}^3$, and the same sensor will have different illuminant gamuts for different illuminants. When the sensor is the human visual system, the term *object-colour solid* is used in place of *illuminant gamut*; for simplicity, we will use the former term for both cases. Like an artificial sensor, human vision has different solids for different illuminants. Object-colour solids will soon be shown to be zonohedra, whose zonal structure will be used in proving the paper's main result.

3.5 The Spectrum Locus

A *spectrum locus vector* is the image of a monochromatic spectrum ρ_i under a colour signal map Φ. Formally, the i^{th} spectrum locus vector is given by

$$\Phi(\rho_i) = (I(\lambda_i) r_1(\lambda_i), I(\lambda_i) r_2(\lambda_i), I(\lambda_i) r_3(\lambda_i)) \tag{17}$$

$$= I(\lambda_i) \cdot (r_1(\lambda_i), r_2(\lambda_i), r_3(\lambda_i)) \tag{18}$$

$$= I(\lambda_i) \cdot \sigma(\lambda_i), \tag{19}$$

where λ_i is a representative wavelength for the i^{th} waveband, and we have introduced the notation $\sigma(\lambda_i) = (r_1(\lambda_i), r_2(\lambda_i), r_3(\lambda_i))$. The set of all such vectors is called the *spectrum locus*. Figure 1 shows an example. A few of the locus vectors have been labeled with representative wavelengths. The axes would be X, Y, and Z, for human vision, but some other R, G, and B for an artificial sensor.

Equation (18) shows that the magnitude of a spectrum locus vector depends only on the illuminant, while its direction depends only on the sensor's response curves. For a given sensor and discretization, then, a spectrum locus vector might lengthen or contract if the illuminant changes, but will not change the direction in which it is pointing.

Together, Equations (10) and (18) demonstrate the zonohedral structure of an object-colour solid, which consists of all vectors of the form $\Phi(\rho)$, for an arbitrary reflectance

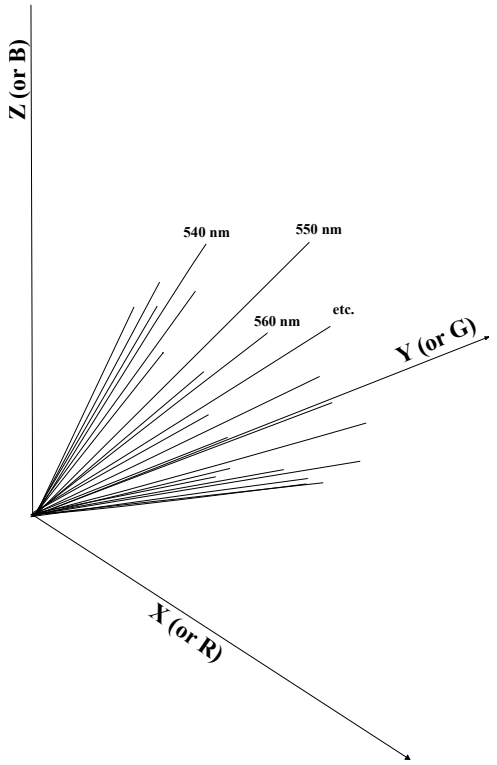

Figure 1: A Spectrum Locus

spectrum ρ. Formally, write

$$\Phi(\rho) = \Phi\left(\sum_i \alpha_i \rho_i\right) \tag{20}$$

$$= \sum_i \alpha_i \Phi(\rho_i) \tag{21}$$

$$= \sum_i \alpha_i \cdot I(\lambda_i) \cdot \sigma(\lambda_i). \tag{22}$$

Every coefficient α_i is between 0 and 1, so Equation (22) expresses $\Phi(\rho)$ as a zonal combination of the spectrum locus vectors given by Equation (18). The object-colour solid depends on the illuminant because the magnitude of a spectrum locus vector varies with the illuminant, even though the response vectors $\sigma(\lambda_i)$ do not.

As mentioned earlier, a zonohedron, which is always a polyhedron, can approach a zonoid, which can be a smooth solid, as the number of generators increases, as long as the generators maintain their approximate directions and shrink appropriately in magnitude. Suppose that the discretization used when constructing an object-colour solid becomes finer, perhaps with wavebands shrinking from 10 nm to 5 nm. Then there will be twice as many wavebands, but any light that reaches a sensor will have about half the power per waveband that it had previously. As a result, there will be twice as many spectrum locus vectors in about the same directions, at about half their previous magnitudes. The zonohedral object-colour solid will have approximately the same size, shape, and terminal point, but many more vertices,

and edges that are about half their previous lengths. As the wavebands become infinitely fine, most or all of the vertices and edges will disappear, leaving a smooth zonoid. Although this paper will restrict itself to the discrete formulation, which is sufficient to prove the main result, the continuous formulation will be occasionally considered as a conceptually important limiting case.

3.6 The Spectrum Cone and Chromaticity Diagram

The magnitudes of the spectrum locus vectors in Figure 1 depend on the illuminant, but other useful constructions result from extending each locus vector to infinity, producing a set of rays that starts at the origin. The convex hull of the rays is called the *spectrum cone*. For human vision, the rays all lie on the boundary of the cone, but, for an arbitrary sensor, some of the rays could lie inside. The spectrum cone is generated solely from the response curves, and is independent of the illuminant, because it needs only the directions of the locus vectors and not their magnitudes. Note that the spectrum cone usually lies in the positive octant of $\mathbb{R}^3$, because output coordinates, both XYZ and RGB, are typically positive.

An instructive picture results from slicing the cone with a plane, producing a profile. Many planes will work, although the chosen plane should intersect every ray in the cone. This paper uses the standard plane $X+Y+Z=1$, or equivalently $R+G+B=1$. Figure 2 shows the locus vectors extended to rays that are truncated when they intersect this plane. Being a convex solid, the spectrum cone intersects the plane in a convex polygonal area that is shaded grey. This area (and sometimes the plane as a whole) is called the *chromaticity diagram*. Since the spectrum cone only depends on the response functions, and not on the illuminant, the chromaticity diagram is also independent of illuminant, and applies solely to the sensor.

Since a spectrum locus vector results from a monochromatic reflectance spectrum, the points where locus rays intersect the chromaticity diagram will be called *monochromatic points*, and each monochromatic point can be assigned the representative wavelength of the locus vector that generates it. The vertices of the polygonal boundary are all monochromatic points, but monochromatic points can appear inside the area as well. For human vision, the chromaticity diagram takes a familiar horseshoe shape, and the monochromatic points all occur on the curved section of the boundary, where they are typically labeled with representative wavelengths.

3.7 Optimal Colours

Colours on the boundary of the object-colour solid are called *optimal colours*. Given a sufficiently fine discretization, every boundary point of the zonohedral object-colour solid is as close as desired to a vertex, so any optimal colour is sufficiently well approximated by a vertex. Every vertex of a zonohedron is a sum of some subset $\mathcal{S}$ of the generators; the coefficient of each generator in $\mathcal{S}$ is set to 1, while the coefficients of all the other generators are set to 0 (see Sect. 2.8.1 of Ref. 4). A subset $\mathcal{S}$ leads to a vertex if and only if $\mathcal{S}$ consists of all the generators that are on one side of a plane $\mathcal{H}$ through the origin. Every plane through the origin can be seen as the kernel of a linear functional F, defined up to a multiplicative constant. F takes on positive values for all the vectors on one side of $\mathcal{H}$, and negative

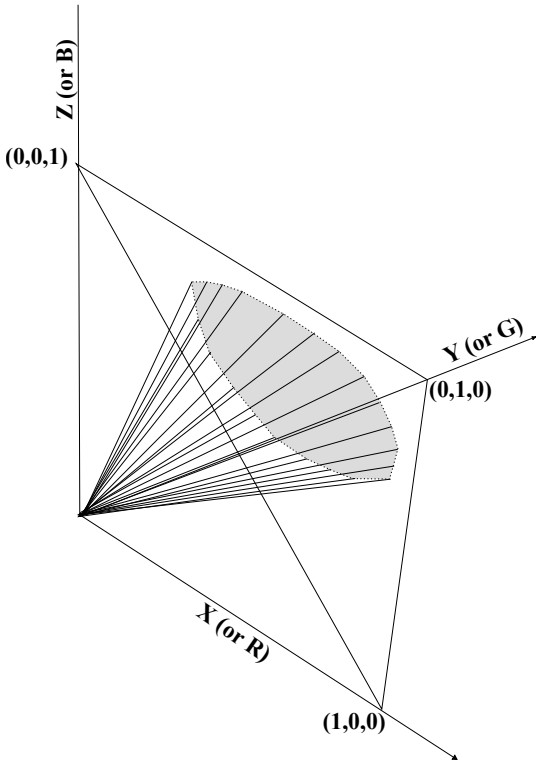

Figure 2: A Spectrum Cone and Chromaticity Diagram

values on the other side. Furthermore, if $F(\mathbf{v})$ is positive for some vector $\mathbf{v}$, then $F(\alpha\mathbf{v})$ is also positive, whenever the coefficient α is positive; an analogous statement holds if $F(\mathbf{v})$ is negative. Rather than requiring spectrum locus vectors to be on one side of $\mathcal{H}$, then, we could just as well require the corresponding rays to be on one side of $\mathcal{H}$.

Figure 3 shows an example, using the spectrum cone and chromaticity diagram from the previous figure. The grey triangle is the intersection of the planar kernel $\mathcal{H}$ of a functional with the positive octant, bounded by the plane used for the chromaticity diagram. The functional is positive on one side of its kernel, and negative on the other side. The monochromatic points on the positive side have been indicated with grey dots, while those on the negative side have been indicated with black dots. Each monochromatic point has a corresponding locus vector, so there is such a vector for each grey dot. Summing up the spectrum locus vectors corresponding to grey dots gives a vertex, which is also an optimal colour, of the object-colour solid. Similarly, the sum of the locus vectors corresponding to black dots gives another optimal colour, that is diametrically opposite, through the solid's center of symmetry, to the first optimal colour. Any optimal colour can be produced by a planar kernel, and likewise all planar kernels produce optimal colours.

The chromaticity diagram allows this construction to be simplified from three dimensions to two. Figure 4 shows the triangular intersection of the positive octant and the plane $X + Y + Z - 1$, from Figure 3. This intersection in fact is the chromaticity diagram, and the kernel of the functional reduces to a straight line. While the kernel must go through the origin in three-dimensional space, any line in the chromaticity diagram can occur. The requirement to be on one side of the planar kernel in $\mathbb{R}^3$ now reduces to being on one side

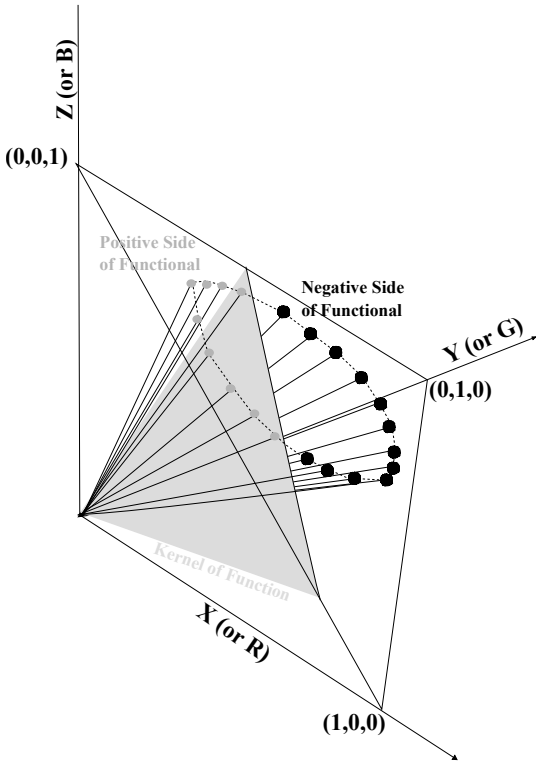

Figure 3: Generating a Zonohedron Vertex from the Spectrum Locus

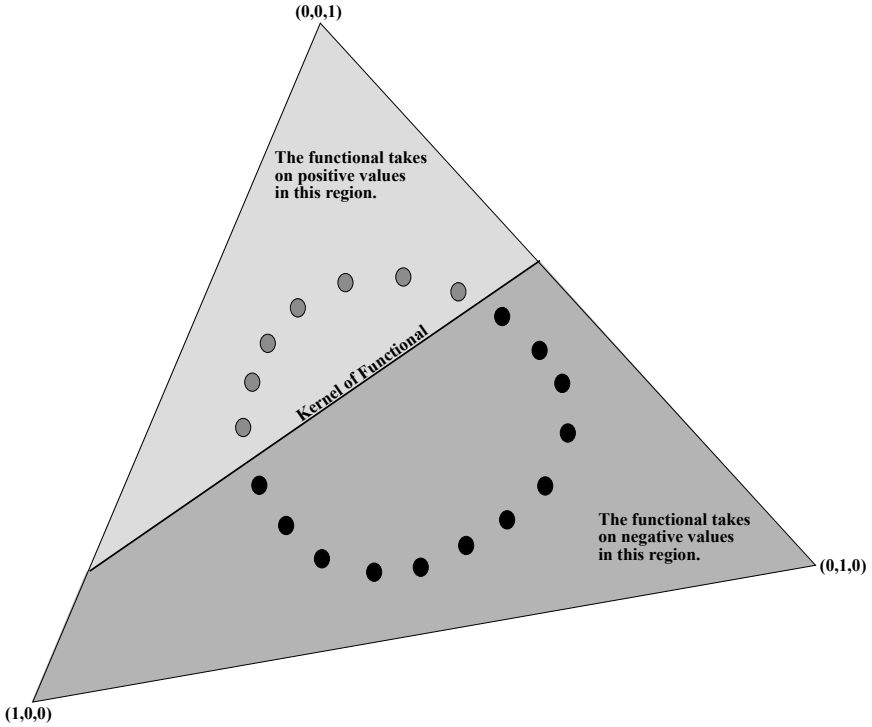

Figure 4: Generating a Zonohedron Vertex from the Chromaticity Diagram

of the line in the chromaticity diagram. The grey dots in the chromaticity diagram still correspond to the grey spectrum locus vectors, which sum to a vertex, but it is easier to work in a two-dimensional space than in $\mathbb{R}^3$.

3.8 Optimal Reflectance Spectra

Since an optimal colour belongs to an object colour solid, any optimal colour must result from some reflectance spectrum, called an *optimal spectrum*. The previous constructions indicate how to characterize optimal spectra. Suppose some straight line cuts a chromaticity diagram, as in Figure 4. The line defines a unique plane $\mathcal{H}$ that contains both the line and the origin. Then select all the monochromatic points on one side of the line, and construct the reflectance spectrum in which all the corresponding wavebands reflect at 100%, and all other wavebands reflect at 0%. In the zonohedron, the colour corresponding to that reflectance spectrum is the sum of all the locus vectors on one side of $\mathcal{H}$, and so is a vertex, and thus an optimal colour.

The positions of monochromatic points in the chromaticity diagram determine the form of optimal spectra, and that form varies with the sensor. For human vision, the monochromatic points are all on the boundary of the horseshoe-shaped chromaticity diagram; furthermore, the points are well-ordered by wavelength, falling in an ordered sequence from 400 to 700 nm. This special form implies[8] Schrödinger's optimal colour theorem:[1,9,4] An optical reflectance spectrum for human vision takes on only the values 0 and 1, with at most two transitions between those values. Conversely, every reflectance spectrum of this form produces an optimal colour. For an artificial camera or sensor, however, whose response functions differ from human vision, some monochromatic points might occur inside the convex hull of the other points, and the points might not appear in order. Then the optimal spectra for that sensor does not have to follow Schrödinger form, or any other convenient form for that matter. In fact, the main result of this paper is that optimal spectra for a more general object called a double object-colour solid do *not* have any nice form, and in fact show an indefinitely large number of transitions between 0 and 1.

4 Double Object-Colour Solids

An object-colour solid for a sensor or the HVS consists of all the colour signals that result from physical objects under a particular illuminant. An important practical case occurs when a sensor makes an image under one illuminant I_1, which produces an unwanted colour cast. For instance, the image colours might tend to be unrealistically reddish, because I_1 itself is reddish. To see how the image would have appeared under a more neutral illuminant I_2, some colour correction algorithm is needed. To develop such algorithms, Eugene Allen suggested (p. 38 of Ref. 2) a six-dimensional hypervolume that concatenates the colour signal maps for both illuminants into a single vector with six components. We shall call such hypervolumes *double object-colour solids*. This section will define double object-colour solids mathematically, and then prove the paper's main result that an optimal reflectance spectrum for a double object-colour solid can always be found, by choosing an appropriate I_1 and I_2, that contains indefinitely many transitions between 0 and 1.

4.1 Mathematical Definitions

Suppose that we have a sensing device with three response curves (r_1, r_2, and r_3) and two illuminants (I_1 and I_2). The two illuminants give two associated colour signal maps, Υ_1 and Υ_2 :

$$\Upsilon_1(\rho_i) = (I_1(\lambda_i)r_1(\lambda_i), I_1(\lambda_i)r_2(\lambda_i), I_1(\lambda_i)r_3(\lambda_i)) = I_1(\lambda_i)\sigma(\lambda_i), \tag{23}$$
$$\Upsilon_2(\rho_i) = (I_2(\lambda_i)r_1(\lambda_i), I_2(\lambda_i)r_2(\lambda_i), I_2(\lambda_i)r_3(\lambda_i)) = I_2(\lambda_i)\sigma(\lambda_i), \tag{24}$$

where λ_i is a representative wavelength for the i^{th} waveband in some discretization of the visible spectrum, and ρ_i is a monochromatic spectrum defined as in Equation (9). Concatenate Υ_1 and Υ_2 into a joint colour signal map $\Upsilon : \chi \to \mathbb{R}^6$:

$$\Upsilon(\rho_i) = (\Upsilon_1(\rho_i); \Upsilon_2(\rho_i)) \tag{25}$$
$$= (I_1(\lambda_i)r_1(\lambda_i), I_1(\lambda_i)r_2(\lambda_i), I_1(\lambda_i)r_3(\lambda_i); I_2(\lambda_i)r_1(\lambda_i), I_2(\lambda_i)r_2(\lambda_i), I_2(\lambda_i)r_3(\lambda_i)) \tag{26}$$
$$= (I_1(\lambda_i)\sigma(\lambda_i); I_2(\lambda_i)\sigma(\lambda_i)). \tag{27}$$

The semi-colon in Equations (26) and (27) could equally as well be replaced with a comma, but a semi-colon will emphasize that the destination space $\mathbb{R}^6$ is divided naturally into two three-dimensional subspaces: the first three coordinates refer to the first illuminant, and the second three refer to the second illuminant. While Equations (23) through (27) apply to monochromatic spectra, we can use linearity to find an expression for the arbitrary spectrum ρ given by Equation (10):

$$\Upsilon(\rho) = \sum \alpha_i \Upsilon(\rho_i). \tag{28}$$

The range $\mathcal{Z} = \Upsilon(\chi)$ of Υ is the double object-colour solid for the sensor and the two illuminants. Since all the coefficients α_i are between 0 and 1, the set $\mathcal{Z}$ is in fact the zonotope generated by the six-dimensional vectors in Expression (27). Under our assumption that all sensor outputs are non-negative, $\mathcal{Z}$ must be a subset of the non-negative octant of $\mathbb{R}^6$, and exhibit the previously discussed properties of convexity, central symmetry, a vertex at the origin, and so on. Single and double object-colour solids share a common zonal structure, but the former exist in three dimensions, while the latter require six dimensions. Optimal colours for a single illuminant occur on the boundary of that illuminant's three-dimensional object-colour solid. Optimal colours for a pair of illuminants are similarly defined to be those colours that occur on the boundary of the double object-colour solid for the two illuminants. Analogously to the single case, optimal reflectance spectra for a pair of illuminants are mapped by Υ to boundary points of the double object-colour solid.

A natural question is whether the optimal spectra of double object-colour solids exhibit a simple, Schrödinger-like form. Apart from theoretical interest, this question arose in a recent paper by Logvinenko et al.,[3] where the boundary of a double object-colour solid needed to be parametrized. Since a six-dimensional convex body has a five-dimensional boundary, the set of reflectance spectra of interest were taken to be 0-1 functions with up to five transitions. While this set worked well enough for their application, it was unclear whether five transitions were always sufficient, or if more would sometimes be needed. The main result of this paper, to be proven next, is negative: *no number* of transitions is sufficient for all situations. More

precisely, for any given sensor and response curves, and any positive integer n, a pair of illuminants can always be found such that at least one optimal reflectance spectrum exists that requires more than n transitions between 0 and 1. (One assumes implicitly for this result, of course, that a sufficiently fine discretization has been chosen.)

4.2 Proof of Main Result

Suppose that we are given a sensing device and its three response curves, along with a discretization of the visible spectrum. Then a sensor chromaticity diagram can be constructed, on which the monochromatic points can be marked. The chromaticity diagram is independent of the illuminant, so the monochromatic points will not move, no matter how the illuminant changes.

Now suppose that there is a double object-colour solid for the sensor, with respect to illuminants $I_1(\lambda)$ and $I_2(\lambda)$. Then the solid is a zonotope in $\mathbb{R}^6$, and each vertex is an optimal "colour," with a corresponding optimal reflectance spectrum. To a vertex $\mathbf{v}$ we can associate a subset $\mathcal{S}$ of the set of six-dimensional spectrum locus vectors, such that $\mathbf{v}$ is exactly the sum of all the vectors in $\mathcal{S}$. Furthermore, to each vertex we can associate a (non-unique) linear functional F, such that F is positive on every vector in $\mathcal{S}$, and negative on every vector not in $\mathcal{S}$. The kernel of F is a five-dimensional hyperplane $\mathcal{H}$ that divides $\mathbb{R}^6$ into two halves; all the vectors in $\mathcal{S}$, and no vectors not in $\mathcal{S}$, are in one half.

The fact that $\mathbb{R}^6$ is a concatenation of two copies of the three-dimensional colour output spaces allows us to write

$$F = F_1 + F_2, \tag{29}$$

where F_1 is a linear functional that is 0 on the last three coordinates (those that correspond to I_2) of $\mathbb{R}^6$, and F_2 is 0 on the first three coordinates (those that correspond to I_1). F_1 can then be seen as a linear functional on the first copy of the three-dimensional colour output space, where the corresponding two-dimensional hyperplane $\ker(F_1)$ divides the space into two halves. There is also an equivalent corresponding line L_1 that divides the chromaticity diagram into two halves. The monochromatic points on the "positive" half correspond to the three-dimensional spectrum locus vectors on which F_1 is positive. Similar constructions exist for a line L_2, derived from F_2, but on the second copy of the sensor's output space.

Rather than starting with a linear functional on $\mathbb{R}^6$, one could reverse the above constructions: choose lines L_1 and L_2, extend them to three-dimensional functionals F_1 and F_2, and then add F_1 and F_2 to get a functional F on $\mathbb{R}^6$, which would lead to a vertex, and thus an optimal colour, for the double object-colour solid. The reversibility of the constructions shows that starting with two lines in the chromaticity diagram is equivalent to starting with a functional.

Regardless of the illuminants (though they must satisfy the mild assumption of being positive on the visible spectrum), the chromaticity diagrams for the two copies of colour output space are identical, with the monochromatic points in exactly the same positions. The lines L_1 and L_2, on the other hand, are likely different. Figure 5 shows an example. Each line divides the diagram into a positive and a negative side, as indicated, and also divides the monochromatic points into positive and negative sets. Each monochromatic point can be classified into one of three categories:

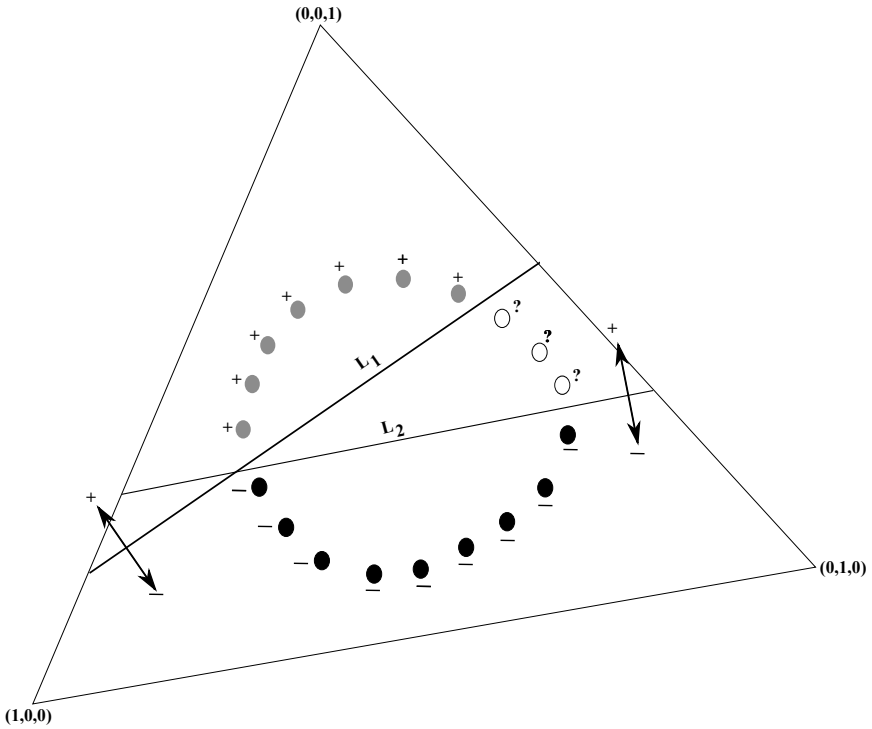

Figure 5: Chromaticity Diagram with Two Lines Arising from a Functional F on $\mathbb{R}^6$

1. Points that are on the positive side of both lines,
2. Points that are on the negative side of both lines, and
3. Points that are on the positive side of one line, and the negative side of the other line.

Figure 5 indicates the points in the first category with a + sign, the points in the second category with a − sign, and the points in the third category with a question mark. If a point is in the first category, then F_1 and F_2 are both positive on the spectrum locus vector corresponding to that point, so F, which is the sum of F_1 and F_2, is also positive there. Similarly, F is negative on the locus vectors corresponding to the second category. The sign of F on a point in the third category is uncertain, however, because F_1 and F_2 have different signs for such a point.

An optimal reflectance spectrum for the optimal colour produced by F is 1 on all the wavebands where F is positive, and 0 elsewhere. Optimal reflectance spectra thus take on only the values 0 and 1. Each monochromatic point corresponds to a waveband, so the optimal reflectance spectrum must be 1 for wavebands whose monochromatic points are labeled with a + sign, and 0 for wavebands whose points are labeled with a − sign. A transition between 0 and 1 only occurs when adjacent wavebands in the spectrum lead to a positive and a negative monochromatic point. In the most conservative case (the case with the fewest transitions), the monochromatic points are well-ordered in wavelength, so that the positive points form a block of adjacent wavebands with value 1, and the negative points form a block of adjacent wavebands with value 0. A transition in Figure 5, then, would occur at the far right, where the positive points meet the negative points.

For the points in the third category, however, those indicated with question marks, we

will show how to create as many transitions as desired, simply by choosing appropriate values for I_1 and I_2 on those wavebands. Let λ be a representative wavelength for a waveband with a question mark. Then

$$\begin{aligned} F(\lambda) &= F_1(\lambda) + F_2(\lambda) && (30)\\ &= F_1(I_1(\lambda)\sigma(\lambda)) + F_2(I_2(\lambda)\sigma(\lambda)) && (31)\\ &= I_1(\lambda)F_1(\sigma(\lambda)) + I_2(\lambda)F_2(\sigma(\lambda)) \quad \text{(by linearity)}. && (32) \end{aligned}$$

The sign of $F(\lambda)$ can then be determined by asking when $I_1(\lambda)F_1(\sigma(\lambda)) + I_2(\lambda)F_2(\sigma(\lambda))$ is less than or greater than 0. Assume without loss of generality that $F_1(\sigma(\lambda)) > 0$ and $F_2(\sigma(\lambda)) < 0$. Then simple algebraic manipulations give

$$\begin{cases} F(\lambda) > 0, & \text{when } \frac{I_1(\lambda)}{I_2(\lambda)} > -\frac{F_2(\sigma(\lambda))}{F_1(\sigma(\lambda))} \\ F(\lambda) < 0, & \text{when } \frac{I_1(\lambda)}{I_2(\lambda)} < -\frac{F_2(\sigma(\lambda))}{F_1(\sigma(\lambda))} \end{cases} \tag{33}$$

Since illuminants take on only positive values, it is easy to see that one could arrange to have $F(\lambda)$ greater or less than 0 as desired, simply by adjusting $I_1(\lambda)$ or $I_2(\lambda)$.

The figure shows three adjacent third-category wavebands. Even if they don't appear by happenstance, however, such a set can always be found by choosing appropriate lines L_1 and L_2, from which a functional F on $\mathbb{R}^6$, and thus an optimal colour, can be identified. If one also allows a sufficiently fine discretization, then the number of third-category wavebands can be made as large as desired. Give the names λ_1, λ_2, λ_3, etc to their representative wavelengths. The previous paragraph shows how these wavebands can take on value 1 or 0 in the support of the optimal reflectance spectrum, simply by adjusting $I_1(\lambda_1)$, $I_1(\lambda_2)$, $I_1(\lambda_3)$, etc. (or $I_2(\lambda_1)$, $I_2(\lambda_2)$, $I_2(\lambda_3)$, etc.) so that F takes on positive or negative values as desired, in accordance with Equation (33). If the monochromatic points are well-ordered, then arranging for $F(\lambda_1)$ to be positive, $F(\lambda_2)$ to be negative, $F(\lambda_3)$ to be positive, and so on, forces transitions between adjacent wavelengths. Even if the monochromatic points are not well-ordered, however, one can always arrange for a third-category waveband to take on a value opposite to an adjacent waveband, and thus cause a transition. In this way, then, an optimal reflectance spectrum can be found with as many transitions as desired. Thus, unlike the Schrödinger form for single object-colour solids for human vision, the optimal spectra for double object-colour solids have no maximum number of transitions, as was to be demonstrated.

4.3 Discussion

Since this paper's main result is negative rather than positive, it is likely of more theoretical interest than practical. Two issues mitigate the difficulties suggested by the result: first, most applications involve "well-behaved" illuminants and sensors, so the worst cases envisioned by the constructions might never occur, and second, a finite number of transitions likely provides a good-enough approximation in many cases.

In contrast to the optimal colour theorem for human object-colour solids, which says that an optimal spectrum never has more than two transitions, regardless of the illuminant, this paper shows that no number of transitions for optimal spectra for double object-colour solids

would be sufficient for an arbitrary pair of illuminants. In practice, however, the illuminants are often already chosen; for instance, a photograph might be taken under Illuminant F1 and then converted to Illuminant D65. Other restrictions might also exist, e.g. that the two illuminants are well-behaved enough to have a limited number of crossings, or that one is a scalar multiple of the other. These special cases might rule out some of the constructions in the proof. While no general rule could be found, because optimal spectra depend on monochromatic points, whose positions vary unpredictably from device to device, it is plausible that cases exist in which the number of transitions in fact is bounded. Each case would have to be investigated individually, however, to find the maximum.

The second point involves approximating the set of all reflectance spectra rather than considering every spectrum. With a typical discretization of about 30 wavebands, the set of all spectra is 30-dimensional, too large to search exhaustively, and too large for many optimization algorithms. Instead, a smaller-dimensional set with a simple form could be chosen, even though it does not contain every spectrum—if it contains a metamer for every reflectance spectrum, it might cover colour space well enough for many practical purposes. Some special forms are Logvinenko's object colour spectra[10] and Centore's four-transition 0-1 spectra.[11] These more manageable forms require three and four dimensions, respectively.

The computational advantages and geometric rationale for 0-1 functions make them natural candidates for double object-colour solids. By our zonal constructions, in fact, the reflectance spectra corresponding to the boundary of a double solid must be 0-1 functions. In a recent paper on metamer mismatch bodies (MMBs), Logvinenko, Funt, and Godau[3] used 0-1 functions for just this situation. Since the boundary of a double solid is five-dimensional, the authors took the natural course of allowing up to five transitions. While the current paper shows that no number of transitions is sufficient for all sensing devices and illuminants, it is still likely that the choice of five transitions is adequate for practical cases, in which some approximation is expected—indeed, the very act of discretization introduces an approximation. A simple approach to insuring that enough transitions were used might be to keep increasing the number of transitions until the differences were negligible for practical considerations. For instance, calculate MMBs using four transitions, then five, then six, and so on, until the variation in the results is too small to be of concern.

While providing some understanding, then, and of course potentially contributing to further development, this paper's negative result, that no number of transitions is sufficient to produce all optimal reflectance spectra for double object-colour solids, probably makes no difference for practical applications, where a small number of transitions would usually provide a good enough approximation.

5 Summary

In the context of imaging or vision, a colour signal map for a given illuminant determines which element of colour space an object's reflectance spectrum is mapped to, when viewed under that illuminant. For a single device or vision system, the range of a colour signal map is a subset of three-dimensional colour space called an illuminant gamut, or, for human vision, an object-colour solid. Each illuminant leads to a different solid. A double object-color solid is defined for a three-dimensional imaging device (one case of which is the human vision

system), and a pair of illuminants. The double object-colour solid is the six-dimensional hypervolume obtained by concatenating, for each possible reflectance spectrum, the three-dimensional single solids. In colour space, the colours on the boundary of an object-colour solid are called optimal colours, and the reflectance spectra which produce optimal colours are called optimal reflectance spectra. For a single solid for human vision, optimal spectra possess Schrödinger form: they take on only the values 0 and 1, with at most two transitions between those values.

This paper investigated optimal spectra for double object-colour solids, that is, those spectra whose two colour signals (one for each illuminant), when concatenated, appear on the boundary of the six-dimensional double solid. A series of geometric constructions showed that, like Schrödinger form, optimal spectra for double solids take on only the values 0 and 1. The main question of interest was whether an optimal spectrum has a maximal number of transitions between those two values, and if so, what that maximum is. Further geometric constructions reduced the problem to two arbitrary lines in the two-dimensional device chromaticity diagram. The number of transitions in a spectrum depends on how those lines divide up the monochromatic points, and on the relative strength of the two illuminants on some of those points. The existence (or not) of a transition was found to depend on a simple algebraic relationship. By increasing the fineness of the discretization as needed, and choosing illuminant values to satisfy the algebraic relationship for specially chosen monochromatic points, two illuminants could be created for which some optimal reflectance spectrum would have an arbitrarily large number of transitions. Thus Schrödinger form does not generalize to double object-colour solids: there is no maximum number of transitions on optimal spectra for double solids. While this result is of theoretical interest, it likely has little practical import, because spectra with a moderate number of transitions, such as five or six, should provide adequate approximation in most cases, and in any event, the approximation can be refined steadily until the result varies as little as desired.

References

1. Erwin Schrödinger. "Theory of pigments of greatest lightness," translated by Rolf G. Kuehni with a technical introduction by Michael H. Brill, accessed Feb 2018 from http://www.iscc-archive.org/resources/translations.php (original article: "Theorie der Pigmente von grösster Leuchtkraft," Annalen der Physik 4, 62, 1920, pp. 603-622).
2. 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.
3. Alexander D. Logvinenko, Brian Funt, & Christoph Godau, "Metamer Mismatching," IEEE Transactions on Image Processing, Vol. 23, No. 1, pp. 34-43, January 2014.
4. Paul Centore, *The Geometry of Colour*, 221 pages, 2017.
5. H. S. M. Coxeter. *Regular Polytopes, 3rd ed.,* Dover Publications, 1963.
6. Ethan D. Bolker, "A Class of Convex Bodies," Transactions of the American Mathematical Society, Vol. 145, pp. 323-345, November 1969.
7. Deane B. Judd, "The 1931 I. C. I. Standard Observer and Coordinate System for Colorimetry," JOSA, Vol. 23, pp. 359-374, October 1933.

8. Gerhard West & Michael H. Brill, "Conditions under which Schrödinger object colors are optimal," Journal of the Optical Society of America, Vol. 73, No. 9, pp. 1223-1225, September 1983.
9. David L. MacAdam. "The Theory of the Maximum Visual Efficiency of Colored Materials," JOSA, Vol. 25, 1935, pp. 249-252.
10. A. D. Logvinenko, "An object-color space," Journal of Vision (11):5, pp. 1-23, 2009.
11. Paul Centore, "Four-transition 0-1 functions for reflectance spectra," Color Research & Application, Vol. 43, No. 4, pp. 612-618, August 2018.

