Geometric Invariants Under Illuminant Transformations

Paul Centore

Abstract

An object colour's CIE XYZ coordinates can change when it is viewed under different illuminants. The set of XYZ coordinates for all object colours, which is called the object-colour solid, likewise varies under different illuminants. This paper shows that, despite these changes, some properties are invariant under illuminant transformations. In particular, as long as the illuminant is nowhere zero in the visible spectrum, optimal colours take the same Schrödinger form, and no two optimal colours are metameric. Furthermore, all object-colour solids have the same shape at the origin: they all fit perfectly into the convex cone (which we will call the spectrum cone) generated by the spectrum locus. The spectrum cone, itself, does not vary when the illuminant changes. The object-colour solid for one illuminant can be transformed into the solid for another illuminant, by an easily visualized sequence of expansions and contractions of irregular rings, called zones.

Keywords: colour solid, illuminant, spectrum locus, spectrum cone, optimal colour, zonohedron

1 Introduction

An illuminant can have a significant effect on colour vision, and on computations for colour processing. When viewing an object, the stimulus that reaches the eye is a combination of the illumination and that object's reflectance function. From the visual stimulus alone, there is no *a priori* way of disentangling the illuminant's contribution from the object colour's contribution.

Despite these difficulties, this paper will present some colour properties that are independent of illuminant. The invariance will be shown using geometric constructions based on the CIE colour-matching functions.[1] The constructions are all located in the positive octant of $\mathbf{R}^3$, and involve the three CIE coordinates denoted X, Y, and Z. Only vector space properties of $\mathbf{R}^3$ are needed; no Euclidean notions such as distance or angle are required. The subset of $\mathbf{R}^3$ that contains visual stimuli is generated by positive linear combinations of XYZ coordinates for monochromatic stimuli. The XYZ coordinates for a monochromatic stimulus are called a spectrum locus vector, and the set of all such vectors is called the spectrum locus. The convex cone generated by the spectrum locus will be called the spectrum cone, and every visual stimulus is located somewhere in this solid cone. This spectrum cone is an invariant: it does not change even when the illuminant changes.

Another important construction in $\mathbf{R}^3$ is an object-colour solid. Every illuminant has a corresponding object-colour solid, which is the set of all XYZ vectors that can result from an object colour, when lit by that illuminant. Geometrically, an object-colour solid is the zonohedron generated from the spectrum locus vectors for an illuminant. From this zonohedral structure, one can conclude that optimal colours, which are the colours on the boundary of the solid, all have reflectance functions in Schrödinger form: the reflectance percentage for any wavelength is either 0 or 100%, with at most two transitions between these two values. This form is also an invariant: no matter what the illuminant, optimal colours have Schrödinger reflectance functions. Furthermore, no two optimal colours are metameric, again regardless of illuminant.

Finally, the zonohedral structure provides a geometrically intuitive way of transforming the object-colour solid for one illuminant into the solid for another illuminant. The transformation involves zones, which take the form of irregular rings around the solid, and consist of translations of the spectrum locus vectors. To transform a solid, each zone is stretched or compressed in sequence. The magnitude of the stretch or compression depends on the relative magnitudes of the spectrum locus vectors for the two illuminants. It is shown geometrically that this transformation sequence does not affect the shape of an object-colour solid at the origin, nor at the white point. In addition, it is shown that every object-colour solid fits perfectly into the spectrum cone at the origin, regardless of the illuminant, so that fact constitutes another invariant.

These statements of invariance should be qualified by a mild assumption on the illuminant. It will be assumed throughout this paper that the illuminant is nowhere zero. In other words, it has some power in every part of the visible spectrum. This assumption is true of the broadband illuminants which are typically encountered in natural situations. If an illuminant were zero somewhere in the visible spectrum, then a viewer would be effectively blind to reflectance functions that were non-zero only where the illuminant is zero. In order to consider the entire visible spectrum, then, we assume that illuminants are nowhere zero.

2 Geometric Colour Constructions

2.1 XYZ Tristimulus Space

A visual stimulus consists of a distribution, $\delta(\lambda)$, of electromagnetic power, between about 400 and 700 nm (the visible spectrum), that impinges on an eye. If the distribution is constant over a reasonably sized area of the visual field, and does not vary over time, then a human will attribute a colour to that area. It is possible that two different areas, despite having different spectral distributions, will have the same colour. In 1931, the Commission Internationale de l'Éclairage[1] (CIE) experimentally determined three colour-matching functions, called $\bar{x}(\lambda), \bar{y}(\lambda)$, and $\bar{z}(\lambda)$, from which three tristimulus values can be computed for any $\delta(\lambda)$:

$$X = \int_{400}^{700} \bar{x}(\lambda)\delta(\lambda)d\lambda, \tag{1}$$

$$Y = \int_{400}^{700} \bar{y}(\lambda)\delta(\lambda)d\lambda, \tag{2}$$

$$Z = \int_{400}^{700} \bar{z}(\lambda)\delta(\lambda)d\lambda. \tag{3}$$

Two colour stimuli, $\delta_1(\lambda)$ and $\delta_2(\lambda)$, are perceived as identical if and only if their three tristimulus values are equal.

Leaving aside special cases such as fluorescence, the term *colour* can be applied to either light sources or objects. A light source can impinge directly on an eye, without reflecting off any intermediate objects; in that case, the spectral distribution of the stimulus is just the spectral distribution of the light source itself. An object, on the other hand, can be seen only when illuminated by a light source. An object reflects different percentages of different wavelengths that impinge upon it. An object's colour properties are summed up by its reflectance function, $r(\lambda)$, which is bounded between 0 and 100%. An object is visible when rays from a light source, say of distribution $s(\lambda)$, reflect off the object, of reflectance function $r(\lambda)$, before reaching the eye. Since the object reflects different percentages of the light, depending on wavelength, the stimulus that reaches the eye has a new distribution, given by $\delta(\lambda) = r(\lambda)s(\lambda)$.

Because of these differences, the CIE tristimulus values defined in Equations (1) to (3) are applied only to light sources, while modified tristimulus values are applied to object colours. The modified values account for two facts. First, reflectance functions have a natural maximum, that occurs when $r(\lambda)$ is identically 100%. Perceptually, an object that reflects 100% of light of any wavelength, is an ideal white. Second, human perception involves lightness adaptation, in which the lightness of any object colour is judged relative to white. The second colour-matching function, $\bar{y}(\lambda)$, was specially chosen to incorporate lightness information. To account for lightness adaptation, the tristimulus values for an object colour are normalized by dividing by the maximum possible lightness:

$$X = \frac{\int_{400}^{700} \bar{x}(\lambda)r(\lambda)s(\lambda)d\lambda}{\int_{400}^{700} \bar{y}(\lambda)s(\lambda)d\lambda}, \tag{4}$$

$$Y = \frac{\int_{400}^{700} \bar{y}(\lambda)r(\lambda)s(\lambda)d\lambda}{\int_{400}^{700} \bar{y}(\lambda)s(\lambda)d\lambda}, \tag{5}$$

$$Z = \frac{\int_{400}^{700} \bar{z}(\lambda)r(\lambda)s(\lambda)d\lambda}{\int_{400}^{700} \bar{y}(\lambda)s(\lambda)d\lambda}. \tag{6}$$

With these definitions, the Y-value for ideal white is always 1, and the Y-value for any other colour is always less than 1.

To simplify expressions throughout this paper, we will often assume, without any loss of generality, that $s(\lambda)$ has been multiplied by an appropriate scalar to produce a normalized light source, $s_n(\lambda)$, such that

$$\int_{400}^{700} \bar{y}(\lambda)s_n(\lambda)d\lambda = 1. \tag{7}$$

This simplification eliminates the denominators from Equations (4) through (6).

It is natural to plot XYZ coordinates in $\mathbf{R}^3$. Since reflectance functions, spectral distributions, and colour-matching functions are all non-negative, it follows that any tristimulus value, for either a light source or an object colour, is non-negative. Therefore, we may restrict XYZ coordinates to the positive octant of $\mathbf{R}^3$. In addition to being a coordinate system, $\mathbf{R}^3$ can also be viewed as a vector space, with the standard linear addition and scalar multiplication of vectors. The vector space structure will give insight into colour relationships.

The set of spectral distributions can be viewed as a subset (though not a subspace) of the vector space of functions over the interval from 400 to 700 nm. This vector space is infinite-dimensional, but in practice it can be approximated adequately by a finite-dimensional space. A natural approach, which we shall follow here, is to divide the interval $[400, 700]$ into N channels of equal width. For practical purposes, 30 channels, each of which is 10 nm wide, is usually sufficient. A useful "basis" then consists of the 30 functions which are 1 on one of the 30 channels, and 0 elsewhere. Formally, the basis function for the ith channel can be written as

$$\Delta_i(\lambda) = \begin{cases} 1(\text{in some units}), & \text{if } \lambda \text{ is in the } i\text{th channel,} \\ 0, & \text{otherwise.} \end{cases} \tag{8}$$

Of course finer divisions, say into 50 or 100 channels, can also be used. This "basis" is not a vector space basis, but it will be seen that it has similar properties.

Similarly, denote the reflectance basis functions for the ith channel by

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

and the source basis functions by

$$S_i(\lambda) = \begin{cases} 1(\text{in some units}), & \text{if } \lambda \text{ is in the } i\text{th channel,} \\ 0, & \text{otherwise,} \end{cases} \tag{10}$$

Since $R_i(\lambda)$ is a ratio rather than a quantity, no units are needed. Although the functions in Equations (8) through (10) are formally identical, they are different in both conceptual and practical terms.

With these vector space structures, the assignment of tristimulus values for light sources, given in Equations (1) to (3), is a linear transformation. The "sum" of two sources, $s_1(\lambda)$ and $s_2(\lambda)$, would be $s_1(\lambda) + s_2(\lambda)$. The tristimulus values, (X_1, Y_1, Z_1) and (X_2, Y_2, Z_2), of the two sources, would be added as vectors in $\mathbf{R}^3$, to get $(X_1 + X_2, Y_1 + Y_2, Z_1 + Z_2)$, which are the trisitmulus values of $s_1(\lambda) + s_2(\lambda)$. If one fixes a light source, $s(\lambda)$, then it can similarly be seen that the assignment of tristimulus values for object colours, given in Equations (4) to (6), is also a linear transformation. Two reflectance functions, $r_1(\lambda)$ and $r_2(\lambda)$, can only be "added" in this case if their total value at any wavelength λ is 1 or less. This will always be the case for basis functions, no two of which are positive on the same channel.

The addition of two light sources is easily interpreted physically. For example, one could simultaneously turn on an incandescent and a fluorescent bulb: the resulting light would be the sum of the lights from the two bulbs individually. Object colours cannot be added in the same sense. For example, suppose there were two paints, with reflectance functions $r_1(\lambda)$ and $r_2(\lambda)$. Then a third paint would be the sum of the first two, if it had reflectance function $r_1(\lambda) + r_2(\lambda)$. In general, of course, there is no way to combine the first two paints to get such a third paint. The two paints can be added formally, however, which is sufficient for our purposes.

It is tempting to impose the standard Euclidean inner product on the vector space of XYZ coordinate vectors. For example, one typically draws the X, Y, and Z axes at right angles. Apart from convenience, this practice has no physical or perceptual justification, because the choice of the basis X, Y, and Z is arbitrary. Accordingly, no notion of angle or distance should be attached to XYZ space *a priori*. The vector space does, however, have a natural structure as a differentiable manifold, obtained by pulling back the differential structure on $\mathbf{R^3}$, along the coordinate assignment map. Thus, it is possible to apply concepts from calculus, such as the tangency of curves through the same point. The vector space structure, given by vector addition and scalar multiplication, along with the differentiable structure, is sufficient to derive invariants of illuminant transformations.

2.2 The Spectrum Locus and the Spectrum Cone

The set of XYZ vectors that can result from an arbitrary visual stimulus, $\delta(\lambda)$, is called the *spectrum cone*. The spectrum cone is the image of all visual stimuli under the linear transformation given by the colour-matching functions. An XYZ stimulus vector corresponding to a basis function $\Delta_i(\lambda)$ is called a *spectrum locus vector*. The set of all spectrum locus vectors is called the *spectrum locus*. It will be seen that the spectrum cone is the convex cone generated by the spectrum locus.

Spectrum locus vectors are interpreted differently for arbitrary visual stimuli, and for stimuli that result from object colours. For an arbitrary visual stimulus, the ith spectrum locus vector, σ_i, is given by $\sigma_i = (X_i, Y_i, Z_i)$, where

$$X_i = \int_{400}^{700} \bar{x}(\lambda)\Delta_i(\lambda)d\lambda, \tag{11}$$

$$Y_i = \int_{400}^{700} \bar{y}(\lambda)\Delta_i(\lambda)d\lambda, \tag{12}$$

$$Z_i = \int_{400}^{700} \bar{z}(\lambda)\Delta_i(\lambda)d\lambda. \tag{13}$$

Since $\Delta_i(\lambda)$ can be defined with arbitrary units, the magnitude of σ_i is correspondingly arbitrary, although its direction is fixed.

Visual stimuli can be considered as a subset, consisting of positive functions, of the vector space of all functions on $[400, 700]$. The set of positive functions is not a vector subspace, so it does not have a basis in the standard sense. It is possible, however, to express any visual stimulus as a *non-negative* sum of the basis functions $\Delta_i(\lambda)$, defined in Equation (8). Formally, any $\delta(\lambda)$ can be approximated by

$$\delta(\lambda) \approx \sum_{i=1}^{N} \alpha_i \Delta_i(\lambda), \alpha_i \geq 0 \ \forall i. \tag{14}$$

Applying Equations (11) through (13) to Equation (14) gives an expression for $\delta(\lambda)$ in XYZ coordinates:

$$(X_\delta, Y_\delta, Z_\delta) \approx \sum_{i=1}^{N} \alpha_i \sigma_i. \tag{15}$$

In any vector space V, the *convex cone* C, generated by a set of vectors $\{\mathbf{v_i} \in V, i = 1..n\}$, is

$$C = \left\{ \sum_{i=1}^{n} \alpha_i \mathbf{v_i} \middle| \alpha_i \geq 0 \ \forall i \right\}. \tag{16}$$

Geometrically, each vector $\mathbf{v_i}$ lies on a unique positive ray that starts at the origin and continues to infinity. The coefficient α_i, which is non-negative, specifies a location on this ray. The summation in Equation (16) constructs the convex hull of these rays. Though it is called a "cone," a convex cone is probably not circular, but is more likely to be some irregular shape. Its vertex is always at the origin of V, and it is a filled solid.

A comparison of Equations (15) and (16) shows that the convex cone of the spectrum locus vectors is identical with the set of all visual stimuli. In XYZ space, then, the set of all visual stimuli is an irregular cone whose vertex is at the origin. Figure 1 shows the spectrum locus, consisting of the XYZ vectors corresponding to the individual functions $\Delta_i(\lambda)$, and Figure 2 shows the resulting convex cone. For a more descriptive three-dimensional interpretation, Figure 2 plots the intersection of the spectrum cone with the unit cube.

While the direction of the vectors in Figure 1 is important, no importance should be attached to their magnitudes. The reason is that any function $\Delta_i(\lambda)$ could be replaced with $\alpha_i \Delta_i(\lambda)$, where $\alpha_i > 0$. This replacement would not change the shape of the spectrum cone in Figure 2. Though it is drawn as a finite set, the spectrum cone in Figure 2 in fact is infinite, extending arbitrarily far from the origin.

2.3 Object-Colour Solids

The previous section considered the spectrum locus for arbitrary visual stimuli. It is also possible to consider the spectrum locus, when restricted to visual stimuli that

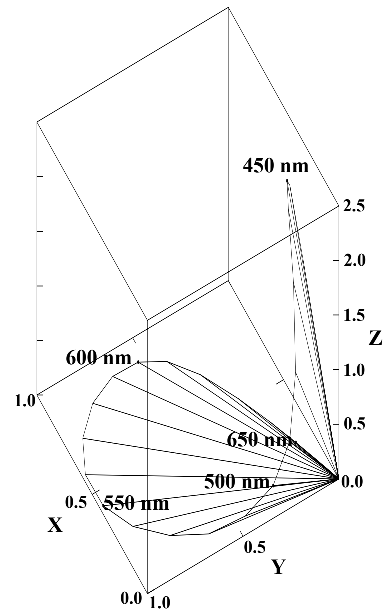

Figure 1: The Spectrum Locus

result from object colours, in the presence of a fixed illuminant. The assumption of a fixed illuminant is often physically realistic, as it simply says that the ambient light is not varying.

A visual stimulus $\delta(\lambda)$ that results from an object colour, can be expressed as

$$\delta(\lambda) = s_n(\lambda)r(\lambda). \tag{17}$$

The set of all object colours corresponds to the set of all reflectance functions. Any reflectance function $r(\lambda)$ can be approximated by

$$r(\lambda) \approx \sum_{j=1}^{N} \alpha_j R_j(\lambda), 0 \leq \alpha_j \leq 1 \ \forall j, \tag{18}$$

and the normalized illuminant can be approximated by

$$s_n(\lambda) \approx \sum_{j=1}^{N} \beta_j S_j(\lambda), \beta_j > 0 \ \forall j. \tag{19}$$

Substituting Equations (18) and (19) into Equation (17) gives

$$\delta(\lambda) \approx \sum_{j=1}^{N} \alpha_j \beta_j S_j(\lambda) R_j(\lambda), 0 \leq \alpha_j \leq 1 \ \forall j. \tag{20}$$

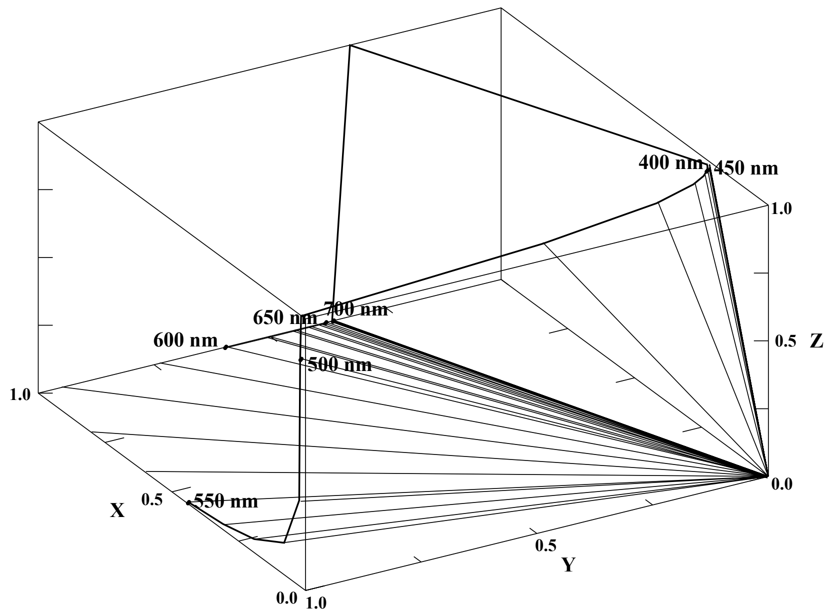

Figure 2: The Spectrum Cone

The ith spectrum locus vector, τ_i, when considering object colours, is the XYZ expression of the reflectance function $R_i(\lambda)$, which is 100% on the ith channel, and 0 elsewhere. In Equation (20), this function is achieved when $\alpha_i = 1$ and all other α's are 0, simplifying to

$$\delta(\lambda) \approx \beta_i S_i(\lambda). \tag{21}$$

By a direct calculation, τ_i is given by (X_i, Y_i, Z_i), where

$$X_i = \beta_i \int_{400}^{700} \bar{x}(\lambda) S_i(\lambda) d\lambda, \tag{22}$$

$$Y_i = \beta_i \int_{400}^{700} \bar{y}(\lambda) S_i(\lambda) d\lambda, \tag{23}$$

$$Z_i = \beta_i \int_{400}^{700} \bar{z}(\lambda) S_i(\lambda) d\lambda. \tag{24}$$

No denominators appear in this expression, because the illuminant has been normalized.

By comparing Equations (22) through (24) to Equations (11) through (13), and using the fact that $S_i(\lambda)$ and $\Delta_i(\lambda)$ are formally equivalent, we see that, for every i, σ_i

and τ_i are in the same direction, but might differ in magnitude. While the magnitude of a spectrum locus vector is arbitrary for a general visual stimulus, the spectrum locus vector for an object colour is as long as possible, when using a fixed illuminant. The vectors in Figure 1, for example, are the longest that could be obtained when Illuminant C impinges on an object colour.

An object-colour solid, for a fixed illuminant, is the set of all XYZ vectors that can result when that illuminant impinges on a surface. Equation (18) expresses an arbitrary reflectance function in terms of coefficients between 0 and 1. For a light source, by contrast, the coefficients could be any non-negative numbers. It follows from linearity, and the derivation of the spectrum locus vectors, that the set of all XYZ vectors resulting from object colours is given by

$$\text{Object-Colour Solid } = \left\{ \sum_{i=1}^{N} \alpha_i \tau_i \middle| 0 \leq \alpha_i \leq 1 \ \forall i \right\}. \tag{25}$$

Rather than being a convex cone, an object-colour solid is the zonohedron generated by the spectrum locus vectors for a fixed illuminant. Each illuminant therefore has its own object-colour solid. Reference 2 presents more details of zonohedral constructions. As an example, Figure 3 shows the zonohedral object-colour solid for Illuminant C, when $N = 30$.

Zonohedra have a considerable amount of structure, that is relevant to illuminant transformations. In particular, each edge of a zonohedron is a translated copy of one of the generating vectors, so each edge of an object-colour solid is a translated spectrum locus vector. If no three spectrum locus vectors are linearly independent, then each face of an object-colour solid is a parallelogram.[3]

These observations allow a *zone* to be defined as follows: If an edge of a zonohedron is chosen, then the zone corresponding to that edge is the set of all the zonohedron's faces, at least one of whose edges is a translated copy of the original edge. Figure 4 shows an example. The figure shows a simple zonohedron, constructed from four generating vectors, $\mathbf{v_1}$ through $\mathbf{v_4}$. Generating vector $\mathbf{v_2}$ can be seen emanating from the origin. All the parallelogram faces that have a copy of $\mathbf{v_2}$ on their boundary have been highlighted in grey. The zone makes an irregular ring around the zonohedron.

Any zone can be constructed in steps, starting from one edge. Choose a face, F_1, which contains that edge. Since F_1 is centrally symmetric, there is a corresponding copy of the starting edge, on the opposite side of F_1. In addition to F_1, another face, F_2, will contain the new edge. On the opposite side of F_2 is another copy of the edge, which joins another face, F_3. Continue in this fashion until the original edge is returned to; since there are only finitely many edges, a return is certain. The

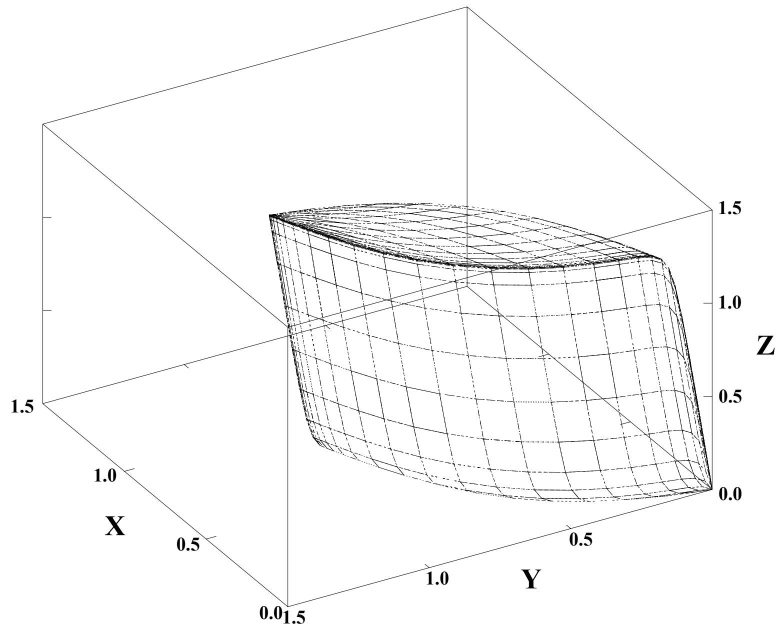

Figure 3: Zonohedral Object-Colour Solid for Illuminant C

set of all faces traversed in this journey is a zone. Furthermore, a zone contains every copy of the original edge that occurs on the boundary of the zonohedron. Each parallelogram face belongs to two zones, and any two zones contain a pair of faces (on opposite sides of the zonohedron) in common. In many zonohedra, and in particular in zonohedral object-colour solids, every zone contains a face, one of whose vertices is the origin. We will use zones in the next section, to describe geometrically how to transform from the object-colour solid for one illuminant, to the object-colour solid for another illuminant.

Every vertex of a zonohedron is a sum of generating vectors, but not all sums of generating vectors appear on the boundary. For example, $\mathbf{v_1} + \mathbf{v_3}$ and $\mathbf{v_2} + \mathbf{v_4}$ in Figure 4 are both in the interior. A sum of generating vectors is on the vertex if and only if the set of vectors in the sum consists of all the vectors on one side a plane through the origin.

3 Illuminant Transformations

Suppose that the CIE coordinates of an object colour, when viewed in a certain illuminant, are given by (X_1, Y_1, Z_1). When the illuminant is changed, the same

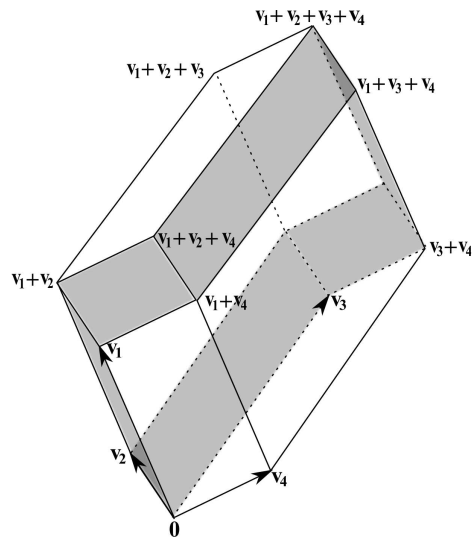

Figure 4: Example of a Zone

object colour will likely have different CIE coordinates, (X_2, Y_2, Z_2). Likewise, the object-colour solid of the first illuminant will differ from the object-colour solid of the second illuminant. Despite these changes, some colour properties will remain invariant even when the illuminant changes. In particular, optimal colours will retain their Schrödinger form, and no two optimal colours will be metameric. In addition, all object-colour solids will fit flushly into the spectrum cone at the origin, although away from the origin they will have different shapes. The spectrum cone, itself, will not change when the illuminant changes. This section will prove these statements, and in addition will demonstrate a geometrically natural way to transform one object-colour solid into another. A technical condition is necessary for most of these results: we will assume throughout this section that the illuminants we use are nowhere zero on the visible spectrum.

3.1 Invariance of the Spectrum Cone

Though the spectrum locus changes when the illuminant changes, the spectrum cone does not. The reason is that the illuminant affects the magnitudes of spectrum vectors, but not their directions. The spectrum cone is the convex cone defined by Equation (16). If any $\mathbf{v_i}$ were replaced by a positive multiple of $\mathbf{v_i}$, the set C would be identical. Replacement by a positive multiple changes magnitude, but not direction, and only direction is needed to define a convex cone.

We will now show that the ith spectrum locus vectors, for any two nowhere-zero illuminants, differ only in magnitude. Equations (22) through (24) give the XYZ coordinates for the spectrum locus vector for the ith channel, for a particular, normalized illuminant $s_n(\lambda)$. Now let us consider a second illuminant, $t_n(\lambda)$, also normalized. In analogy to Equation (19), we can write

$$t_n(\lambda) \approx \sum_{j=1}^{N} \gamma_j S_j(\lambda), \gamma_j > 0 \ \forall j. \tag{26}$$

By following through the derivation of Equations (22) through (24), using Equation (26), we get a new set of spectrum locus vectors:

$$X_{it} = \gamma_i \int_{400}^{700} \bar{x}(\lambda) S_i(\lambda) d\lambda, \tag{27}$$

$$Y_{it} = \gamma_i \int_{400}^{700} \bar{y}(\lambda) S_i(\lambda) d\lambda, \tag{28}$$

$$Z_{it} = \gamma_i \int_{400}^{700} \bar{z}(\lambda) S_i(\lambda) d\lambda. \tag{29}$$

Taking the ratio of the components of the two spectrum locus vectors,

$$\frac{X_{it}}{X_i} = \frac{Y_{it}}{Y_i} = \frac{Z_{it}}{Z_i} = \frac{\gamma_i}{\beta_i}. \tag{30}$$

Since these ratios are constant, therefore the vectors are in identical directions, but differ only in magnitude. The directions of spectrum locus vectors are therefore invariant under an illuminant transformation. It follows that the spectrum cone, which is the convex cone generated by the spectrum locus vectors, is also invariant under an illuminant transform, as was to be shown.

The invariance of the spectrum cone requires a nowhere-zero illuminant. If one illuminant were zero in some channel, then its spectrum locus vector for that channel (for both light sources and object colours) would also be zero. Other illuminants, however, would not be zero on that channel, so they would have non-zero spectrum locus vectors for that channel. The convex cones they generate would contain those spectrum locus vectors, but the convex cone for the original illuminant would not contain them. Thus the spectrum cones would differ, and invariance would not hold.

3.2 Invariance of Optimal Colours

Originally, an optimal object colour was defined as a colour of maximal luminosity, given its chromaticity.[4] An important result is the Optimal Colour Theorem,[4,5] which states that an object colour is optimal if and only if its reflectance function has Schrödinger form, that is, it takes on only the values 0 and 100%, with no more than two transitions between those values. A modern definition, equivalent to the original one, is that an optimal colour is any colour whose XYZ coordinates are on the boundary of the object-colour solid.[2] Since the object-colour solid will change when the illuminant changes, it would seem that different illuminants would lead to different optimal colours. In fact, this is not the case. We will show that the reflectance functions of optimal colours always have Schrödinger form, regardless of the illuminant (under the assumption that the illuminant is nowhere zero). In addition, we will show that no two Schrödinger reflectance functions are metameric, again regardless of the illuminant.

A recent proof[2] of the Optimal Colour Theorem involves the following characterization: every vertex of a zonohedron is the sum of all the generators on one side of a plane through the origin; conversely, the sum of all the generators on one side of a plane through the origin is necessarily a vertex. In the current context, the zonohedron of interest is an object-colour solid, and the generators are the spectrum locus vectors. Any particular spectrum locus vector corresponds to a basic reflectance function $R_i(\lambda)$, for some i. One can make the approximation (which becomes exact as $N \to \infty$) that optimal colours, when expressed in XYZ coordinates, are identical with the vertices of the zonohedral object-colour solid. Then, by the characterization of vertices, optimal colours, in XYZ space, are sums of spectrum locus vectors. Since spectrum locus vectors correspond to reflectance functions of the form $R_i(\lambda)$, the reflectance function of an optimal colour must be a sum of such R_i's. One sees empirically that the spectrum locus vectors are cyclical and well-ordered in XYZ space. The set of such vectors that are on side of a plane through the origin therefore correspond to contiguous R_i's, allowing for wraparound from 700 nm to 400 nm. The sum of a set of contiguous R_i's is a reflectance function of Schrödinger form.

Though Schrödinger[4] and MacAdam[5] do not say so, the Optimal Colour Theorem requires that the illuminant be nowhere zero. If an illuminant were zero in some channel, then the reflectance function could take on any value in that channel, without affecting the colour perception. The Schrödinger form would then have to be modified to account for this extra freedom.

The fact that an illuminant transformation changes only the magnitudes of spectrum locus vectors, and not their directions, implies that optimal colours all have the same Schrödinger form. Any plane through the origin divides the spectrum locus into two groups of vectors, and each group sums to an optimal colour of Schrödinger form. Even if the illuminant is transformed, the directions of the spectrum locus vectors remains constant, so that plane subdivides the new vectors into the same two groups. The sum of each group of vectors, which is an optimal reflectance function, remains unchanged, then, even when the illuminant changes.

Non-metamerism of optimal colours follows similarly. It has been proven[3] that optimal colours are never metameric, as long as no three of the generating spectrum locus vectors are linearly dependent. Determinations of linear dependence and independence rely only on the directions of the vectors involved, and not on their magnitudes. Calculations show that no three spectrum locus vectors are linearly dependent for the CIE standard observer, under illuminant E, whose density is equal-energy over the visible spectrum. The independence relations are unchanged when the illuminant is transformed, so two optimal colours will not be metameric.

Again, this result requires that the illuminant be nowhere zero. If some channel were zero, then the corresponding spectrum locus vector would be the zero vector, and any three vectors, one of which is the zero vector, is trivially linearly dependent. As in the Optimal Colour Theorem, the reflectance function could take on any value in that channel, without affecting colour perception—every different reflectance value leads to another metamer for that colour.

3.3 Transformation of Object-Colour Solids

When an illuminant changes, the object-colour solid also changes. The previous section showed that an illuminant change does not affect the reflectance functions corresponding to the colours on the boundary. In this section, we show that the shape of the solid at the origin is also unaffected: the object-colour solid for a nowhere-zero illuminant always fits perfectly inside the spectrum cone. We also present a sequence of steps for transforming the colour solid for one illuminant into the colour solid for another illuminant. Each step involves stretching or contracting one zone of the colour solid. Unlike the other results, the transformation can occur even if the illuminant is zero on some channels.

An example will demonstrate these assertions. Figure 5 shows two illuminants, s_1 and s_2, at the far left and far right of the upper row. For simplicity, only four channels are used, centered at 437.5 nm, 512.5 nm, 587.5 nm, and 662.5 nm. Calculations will use the values of the illuminant and the colour-matching functions just at these four wavelengths. The middle three plots show steps in transforming from s_1 to s_2. In the first step, shown in the second plot from the left, we change the first channel from its value for s_1, which is 70, to the value of the first channel for s_2, which is 100. This change is shown with a thick line. The other three channels of s_1 are left unchanged. In the second step, we change the value of the second channel from its s_1 value, which is 75, to its s_2 value, which is 110, while leaving the other channels unchanged. The middle plot shows this change with a thick line. After performing this procedure on all four channels, we arrive at s_2, on the far right.

The middle row of Figure 5 shows the spectrum loci for s_1 and s_2, and for the three intermediate illuminants. The directions of the four spectrum locus vectors are identical for all the illuminants, but they differ in magnitude. The magnitudes are partially determined by the normalizing requirement that the Y-value of ideal white is 1. The bottom row of Figure 5 shows the object-colour solids for the five illuminants in the top row. On the far left of the bottom row is the object-colour solid for s_1. It is a zonohedron, so every edge is a translate of a spectrum locus vector. The zone corresponding to the first channel has been highlighted. It consists of all faces that contain a copy of the generator $\mathbf{v_1}$, which has also been highlighted.

We will study the effect on the colour solid as the illuminants in the top row transform from s_1 to s_2. The second plot in the middle row of Figure 5 shows the spectrum locus vectors after the first step, and the second plot in the bottom row shows the corresponding object-colour solid. The new solid differs from the original solid in two ways. First, a different normalizing factor has been used. This factor changes the overall size of the solid, but not its shape. The second difference is that the spectrum locus vector for the first channel is significantly longer for s_2 than for s_1, taking on values of 100 and 70, respectively. The shape of the colour solid changes as a result: the first zone is stretched in the direction of the first spectrum locus vector.

Similar changes occur throughout the remaining three steps. The spectrum locus vector for the channel under consideration will become shorter or longer, depending on the values of the illuminants in that channel. If it becomes longer, then the zone for that vector stretches along the direction of the vector. The two pieces of the zonohedron on either side of the zone are unaltered individually, but move farther apart, in the direction of the vector. The zonohedron as a whole stretches (irregularly) along that zone. If the spectrum locus vector becomes shorter, then the solid contracts (similarly irregularly) along that zone. Figure 5 shows the sequence of stretches and contractions, along the highlighted zones, that transform the s_1 solid into the s_2 solid.

This algorithm, which transforms the object-colour solid when the illuminant is transformed, can be extended to any number of channels, rather than just four. In fact, the illuminants can even be zero on some channels. Transforming from a non-zero channel to a zero channel would mean eliminating a zone entirely, while transforming from a zero channel to a non-zero channel would mean adding a new zone.

The second assertion to be shown is that any object-colour solid, for a nowhere-

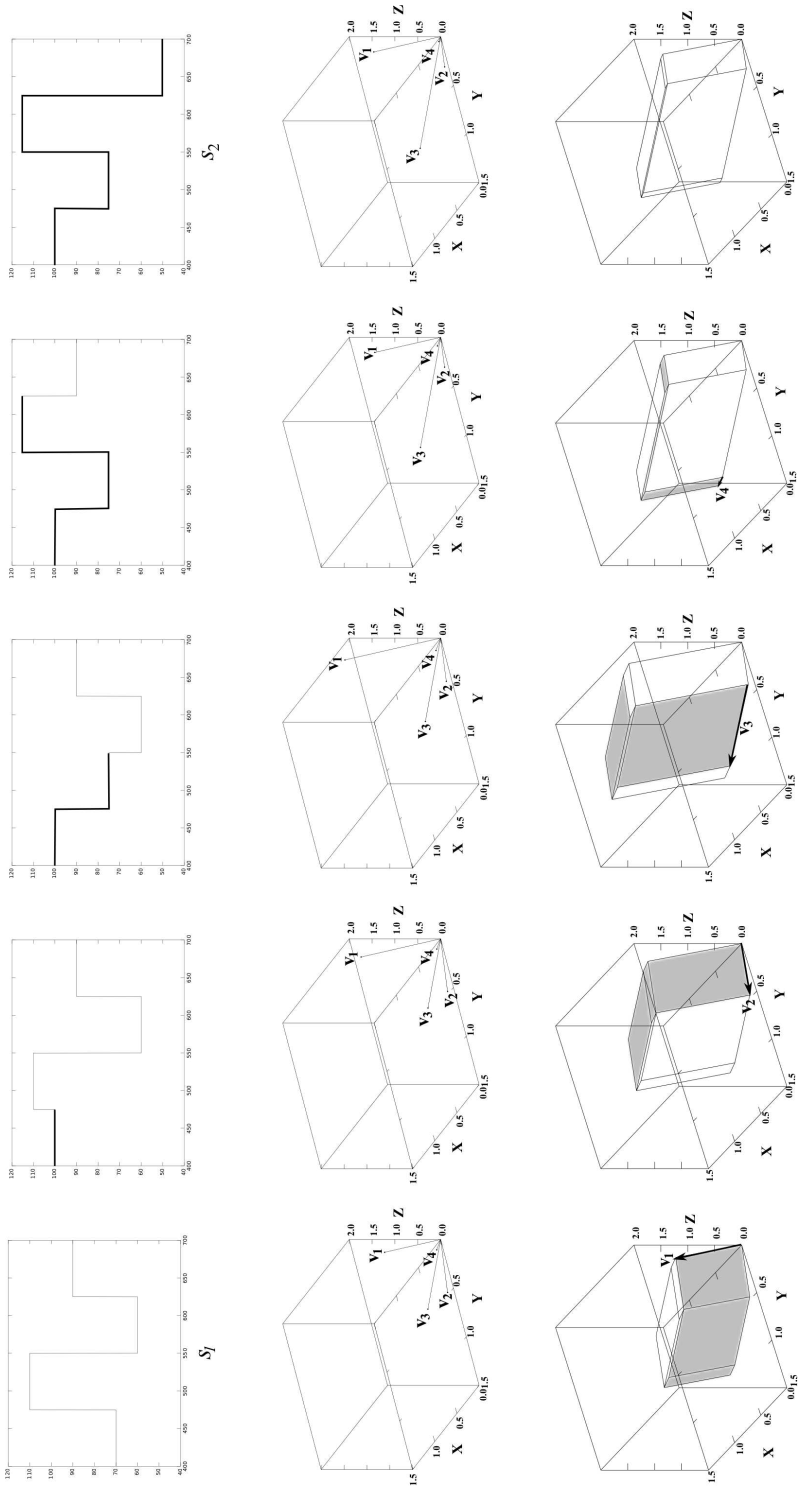

Figure 5: Transformation of an Object-Colour Solid

zero illuminant, fits perfectly into the spectrum cone at the origin. A perfect fit means that any ray contained in the boundary of the spectrum cone is tangent to the colour solid at the origin. For example, if the Illuminant C colour solid shown in Figure 3 were drawn on the same axes as the spectrum cone in Figure 2, then each spectrum locus vector would have tangential contact with the solid at the origin. Since the cone is irregular, it would be impossible to move the colour solid and still maintain this contact everywhere.

The colour solids in Figure 5 provide an easy way to see this result. All the colour solids in that figure are zonohedra, whose generators are spectrum locus vectors. As noted before, these vectors differ in magnitude, but not in direction. It can be seen that the edges that meet at the origin are just the spectrum locus vectors. (Reference 2 derived this result more formally; it depends on the fact that the spectrum locus is cyclic.) The spectrum cone in this simple example would be an irregular square pyramid whose apex is at the origin, and whose base is infinitely far away. The colour solid's faces, that meet at the origin, are Minkowski sums of two locus vectors, so they are contained in the planar region spanned by positive combinations of those vectors. The same planar regions are faces of the convex cone generated by the locus vectors. Since this convex cone is just the spectrum cone, the solid's faces at the origin have full contact with the spectrum cone's faces. Thus the object-colour solid makes a flush fit inside the spectrum cone. As an example, the object-colour solid for Illuminant C, shown in Figure 3, can be seen to fit inside the spectrum cone, shown in Figure 2.

Again, the number of channels is arbitrary; the same arguments hold with a more practical number, such as 30. It is necessary, however, that no channel be zero. If an illuminant has a zero channel, then it will not fit snugly inside the spectrum cone generated from an illuminant where that channel is non-zero, because the second illuminant's spectrum cone will contain a spectrum locus vector that the first illuminant's spectrum cone lacks.

As an illuminant changes, then, the section near the origin remains fixed in location and orientation, although it will stretch or contract in the direction of different spectrum locus vectors. The ideal white point, which is the farthest point from the origin, is constrained to lie on the plane $Y = 1$, but its X- and Z-coordinates will vary for different illuminants. If the illuminant changed continuously, then the resulting series of object-colour solids would also vary continuously, but always remain fixed snugly within the spectrum cone at the origin. Meanwhile, the white point, far away from the origin, could be tracing a very complicated path on the plane $Y = 1$, while the solid as a whole lengthened and shortened along different rings that encircle it.

Another interesting corollary comes from the fact that a zonohedron has a 180^o

rotational symmetry. It follows from this fact that the white and black points have the same three-dimensional shape, just rotated in space. Since the shape around the black point is invariant, the shape around the white point is also invariant, though its location in XYZ space changes as the illuminant changes.

4 Summary

This paper has identified some properties in XYZ space that do not change, even when an illuminant changes, under the mild assumption that the illuminant is nowhere zero on the visible spectrum. The spectrum cone, which is the set of all XYZ triples that could result from a visual stimulus, is independent of illuminant. Amongst object colours, the reflectance functions of optimal colours take the same Schrödinger form, regardless of the light source. Furthermore, no two optimal colours are metameric. An object-colour solid, which is defined in terms of an illuminant, will occupy a different subset of XYZ space when the illuminant changes. Still, every object-colour solid will fit perfectly into the spectrum cone at the origin of XYZ space, and this property holds for every nowhere-zero illuminant. A colour solid's change in shape, under an illuminant change, can be described by a sequence of stretches and contractions, each of which occurs only on a zone of the solid. The magnitude of the stretch or contraction depends on the different values of the original and new illuminant, at the spectrum locus vector that corresponds to that zone.

1. Deane B. Judd. "The 1931 I. C. I. Standard Observer and Coordinate System for Colorimetry," JOSA, Vol. 23, October 1933, pp. 359-374.
2. Paul Centore, "A Zonohedral Approach to Optimal Colours," Color Research and Application, to appear, 2012.
3. Paul Centore, "Non-metamerism of boundary colors in multi-primary displays," Journal of the Society for Information Display, Vol. 20/4, April 2012, pp. 214-220.
4. Erwin Schrödinger. "Theorie der Pigmente von grösster Leuchtkraft," Annalen der Physik 4, 62, 1920, pp. 603-622. In 2010, Rolf G. Kuehni prepared an English translation, "Theorie [sic] of pigments of greatest lightness," which includes a technical introduction by Michael H. Brill.
5. David L. MacAdam. "The Theory of the Maximum Visual Efficiency of Colored Materials," JOSA, Vol. 25, 1935, pp. 249-252.

