

# Zonohedral Gamuts For Colour Constancy

Paul Centore

© February 12, 2016

## Abstract

*Forsyth's 1990 gamut-based illuminant estimation (GBIE) is an important colour constancy algorithm. When a sensing device (assumed to contain three individual sensors) makes an image of physical objects under a single illuminant, GBIE uses the RGB outputs of the image to estimate the illuminant. The set of all RGB outputs in the image is called the image gamut. The set of all RGB outputs that could result when a particular illuminant shines on an object of arbitrary local colour is called the illuminant gamut for that illuminant. Forsyth's algorithm uses the fact that a particular illuminant can be a possible light source for an image only if that illuminant's gamut contains that image's gamut. Implementations to date have used training sets of reflectance spectra to approximate illuminant gamuts. A main result of the current paper is a method for calculating illuminant gamuts exactly, rather than approximating them: geometric constructions prove that illuminant gamuts are zonohedra, and that they are generated by the device's spectrum locus vectors, which are the RGB outputs for monochromatic reflectance spectra. On a more negative note, the same geometric constructions also show that some common GBIE practices are theoretically unsound. In particular, many implementations assume that a linear image of an illuminant gamut is again an illuminant gamut; also, containment tests are sometimes performed in chromaticity space rather than RGB space. The underlying geometry contraindicates both these practices. Finally, a new GBIE algorithm is suggested, that calculates zonohedral gamuts explicitly, avoids contraindicated practices, and takes advantage of the geometry of colour constancy.*

**Keywords:** Computational Colour Constancy, Illuminant Estimation, Zonohedron, Gamut, Sensor Response, Image, GBIE

## 1 Introduction

Colour constancy<sup>1</sup> is the ability to identify an object's local colour, regardless of the lighting conditions in which that object is viewed. Though human colour constancy is robust and effective in everyday situations, its mechanisms are far from understood. A related problem is *computational* colour constancy: can algorithms give imaging devices the same colour constancy that humans exhibit? Colour constancy is a difficult problem mathematically because the visual stimulus from a viewed object is the product of the light source and that object's reflectance spectrum. Mechanical sensors or human eyes, however, only receive the

visual stimulus, and cannot disentangle the factors in the product, at least for any particular stimulus.

Though disentanglement is impossible for any individual stimulus, it might still be possible when given the set of stimuli corresponding to a scene as a whole. In 1990, Forsyth<sup>2</sup> introduced *gamut-based illuminant estimation* (GBIE) to implement this approach. *Illuminant estimation* (IE) makes inferences about the light source under which a scene is viewed. Since the light source is one multiplicative factor contributing to visual stimuli, information about the light source can lead to information about the other factor, object colours, which could produce colour constancy.

GBIE uses the fact that universe of object colours is limited. Even if a scene contained every possible object colour, the visual stimuli making up an image of the scene, when viewed under a single light source, would only fill a limited *illuminant gamut*. Furthermore, this illuminant gamut would be different for different light sources. Forsyth’s main insight is that an *image gamut*, which is the set of stimuli that appear in an image of a scene, must be contained in the gamut of the illuminant under which that image was made. Image gamuts can be extracted directly from image data, and illuminant gamuts can be estimated, so containment tests provide some information about what light sources are possible. Since 1990, Forsyth’s gamut-based approach to IE has spawned, and continues to spawn, many implementations and variations.<sup>3</sup>

This paper focuses on Forsyth’s GBIE approach in the context of a mechanical sensing device, such as an RGB camera, that makes images of scenes of physical objects, under a single light source. The light source itself is assumed not to appear in the image, even as a highlight. The device is assumed to contain three independent sensors, each of which responds linearly. The problem is to make inferences about the light source under which an image was made, using only the image gamut (and, of course, the sensors’ properties).

The current paper’s main contribution is an explicit form for illuminant gamuts: they are shown to be zonohedra in the device’s *RGB* output space. Figure 1 shows an example of a zonohedral gamut. As the paper will explain in detail later, a zonohedron is the Minkowski sum of a set  $\mathcal{S}$  of  $n$  vectors  $\mathbf{v}_i$  in  $\mathbb{R}^3$ . Formally,

$$Z(\mathcal{S}) = \left\{ \sum_{i=1}^n \alpha_i \mathbf{v}_i \mid 0 \leq \alpha_i \leq 1 \ \forall i \right\}. \quad (1)$$

A zonohedral gamut starts at the origin of *RGB* space, which corresponds to an ideal black. The gamut’s terminal point (the farthest point from the origin) is the sensor’s *RGB* response to a perfectly diffusing white, under the illuminant for that gamut. As can be seen in the figure, a zonohedron is centrally symmetric. A zonohedron is a polyhedron. Each edge is a translated copy of some  $\mathbf{v}_i$ , and, under some mild conditions, each face is a parallelogram. The  $\mathbf{v}_i$ ’s in this paper are the spectrum locus vectors, which are the *RGB* responses to objects with monochromatic reflectance spectra. At the origin, a zonohedron is often shaped like a brilliant-cut diamond, with many edges converging on the vertex there. In many practical cases, the converging edges consist exactly of the spectrum locus vectors.

The zonohedral form allows the gamut for a given illuminant to be calculated, rather than estimated. The vertices, edges, and faces can be determined from the spectrum locus vectors, which can themselves be found from the sensors’ response functions. In the cyclic case (to

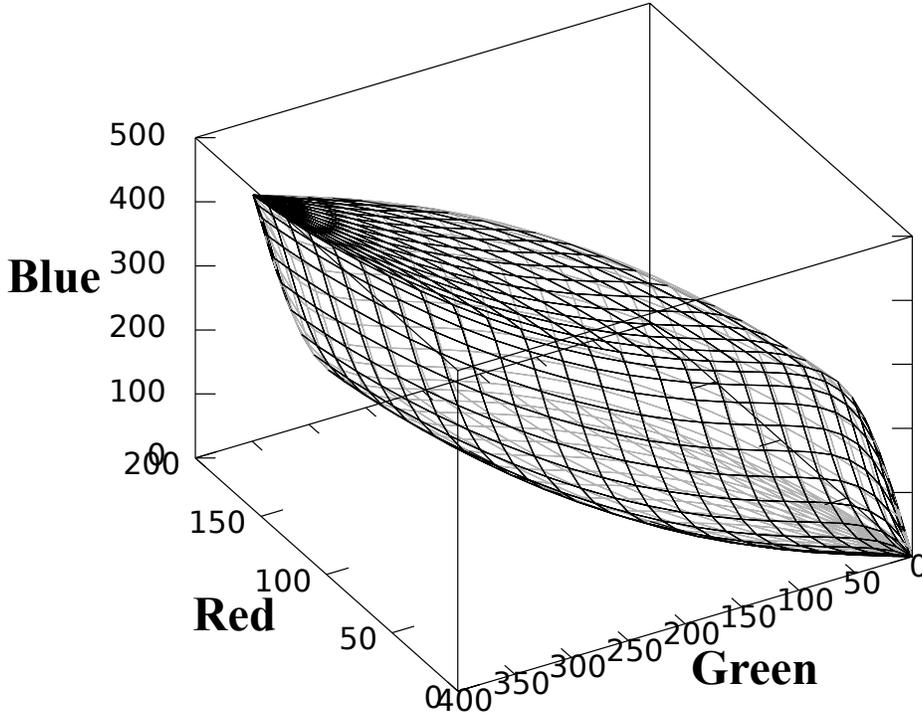


Figure 1: A Zonohedral Illuminant Gamut in  $RGB$  Space

be defined later), the vertex-edge-face structure can be read off from an easily constructed table. Calculating illuminant gamuts can only improve GBIE algorithms: efforts to date have estimated gamuts from training sets of reflectance spectra. Since training sets are limited, the estimated gamuts are incomplete at best, and any geometric structure goes unnoticed. The zonohedral approach eliminates both these shortcomings.

A further important insight is the relationships between different illuminant gamuts. Figure 2 shows two gamuts, for Illuminants A and E, for the same imaging device. A counterintuitive result is that any two gamuts (provided their illuminants have positive power at every wavelength) fit together perfectly in some neighborhood of the origin—they are both shaped like brilliant-cut diamonds there, whose converging edges are spectrum locus vectors. The terminal points of the zonohedral gamuts, however, can be far apart, as the figure shows.

The identical shapes at the origin imply some negative results. Forsyth’s original paper expressed a transformation between two illuminant gamuts as the sum of a linear operator and a residual error term; he and further researchers have largely ignored the error term, treating a linear transformation of an illuminant gamut as another illuminant gamut. Commonly, one starts with a known *canonical gamut*, and looks for a linear transformation such that the transformed canonical gamut contains a given image gamut; then one can calculate the  $RGB$ ’s that the object would produce under the canonical illuminant. In 1996, Finlayson<sup>4</sup> relied on this linear framework to restrict containment tests to two-dimensional chromaticity sets, instead of three-dimensional  $RGB$  sets, to simplify computations.

Unfortunately, the current paper shows that the linear framework is theoretically untenable. The gamuts in Figure 2, and in fact all illuminant gamuts, have identical edges (which

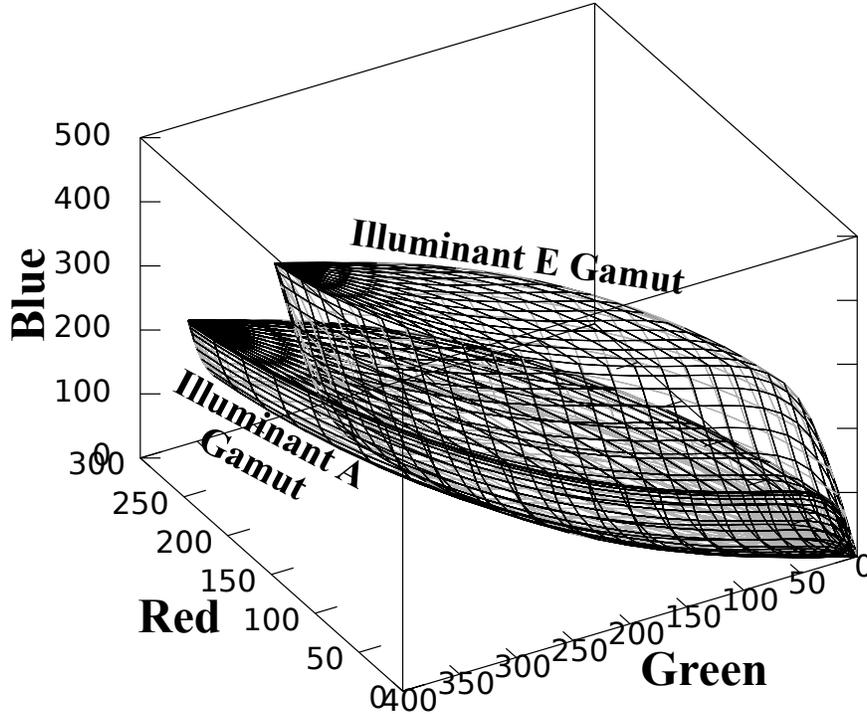


Figure 2: Zonohedral Gamuts in  $RGB$  Space for Illuminants A and E

are all spectrum locus vectors) at the origin. To send one gamut to another gamut, a linear transformation must therefore send the set of spectrum locus vectors to itself. Given that those vectors typically lie on an asymmetric, irregular cone, likely only trivial transformations satisfy this condition. Selecting a transformation without regard to the locus vectors would almost certainly not produce another illuminant gamut. An unwelcome conclusion is that many GBIE implementations have been based on unsound premises: even if a linear image of a canonical gamut contains the image gamut, we have not really estimated an illuminant. Finlayson’s chromaticity containment criterion, which is derived from the linear framework, is similarly untenable; in fact, the geometric constructions will show that it is vacuous, because it is always satisfied.

The geometric constructions in this paper, on the other hand, should provide a sound foundation for further implementations. Exact illuminant gamuts, rather than approximations, can be calculated directly, for any illuminant. A new GBIE algorithm will be suggested, that avoids linear gamut transformations, but still benefits from Finlayson’s simplification, by handling an illuminant’s chromaticity and power level separately.

This paper is organized as follows. First, the elements of the IE problem (illuminants, objects, and sensors) are formulated mathematically, as is the problem itself. Next, these formulations are used in geometric constructions, such as the sensor spectrum locus, the sensor spectrum cone, and the sensor chromaticity diagram, culminating in the zonohedral illuminant gamut. The geometric constructions are then applied to the IE problem. Zonohedral calculations are recommended for illuminant gamuts, and the linear transformation framework and chromaticity containment tests are shown to be theoretically unsound. Finally, a suggested approach is given for a new GBIE algorithm that takes advantage of the

geometric constructions.

## 2 Problem Formulation

This paper deals with illuminant estimation (IE) for a mechanical sensing device, such as a camera, which produces images of scenes consisting of physical objects. To make an image in the first place, some illumination must fall on the scene. The objects in the scene reflect that illumination, and the reflected light enters the sensing device, which converts it into an image. The illuminant estimation problem is to determine the illuminant solely from such an image. This section presents formal descriptions of the problem’s three elements: illuminants, objects, and sensors. Later sections will use this formalism to produce geometric constructions, and apply them to illuminant estimation.

### 2.1 Illuminants

Light is electromagnetic radiation in the visible spectrum, from about 400 to 700 nm. Light can be described formally by a spectral power distribution (SPD), which is a function  $I(\lambda)$  that specifies the power at each wavelength  $\lambda$  in the visible spectrum. The set of SPDs can be viewed as a subset (though not a subspace) of the vector space of functions over the interval from 400 to 700 nm. The vector space structure allows two SPDs to be added (corresponding to physical superimposition) to produce a new SPD, and also allows an SPD to be multiplied by a positive constant, again producing a new SPD.

While the vector space is infinite-dimensional, in practice it can be approximated adequately by a 31-dimensional space of discrete functions, each of which has positive physical power at only one wavelength, and no power at other wavelengths. Such functions are called *monochromatic*. The 31 wavelengths  $\lambda_i$  are the multiples of 10 nm, between 400 and 700 nm. Formally, the  $i^{\text{th}}$  such “basis” function is given by

$$I_i(\lambda) = \begin{cases} 1(\text{in some units}), & \text{if } \lambda = \lambda_i \\ 0, & \text{otherwise,} \end{cases} \quad (2)$$

and an arbitrary nowhere-zero SPD  $I$  can be written as a positive linear combination of the  $I_i$ ’s:

$$I(\lambda) = \sum_{i=1}^{31} P_i I_i(\lambda) \quad (P_i > 0 \forall i). \quad (3)$$

The qualification *nowhere-zero* avoids degenerate cases in which an SPD has no power at some wavelength. Nowhere-zero SPDs might also be thought of as broadband SPDs, and any natural SPDs are nowhere-zero. Removing the nowhere-zero restriction is possible, but technically complicated, so it will be left in for now.

A *light source* is a physical object, such as a light bulb or the sun, that produces light. In colour science, an *illuminant* is an idealized description of a physical light source. Figure 3 shows Illuminant A, which is intended to model incandescent lighting. An illuminant is specified by a relative SPD, while a light source is specified by an absolute SPD. Multiplying

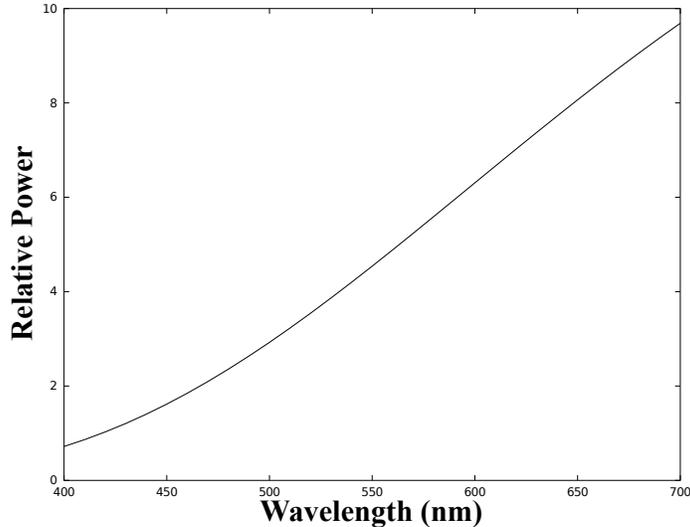


Figure 3: Relative SPD for Illuminant A

a light source’s SPD by a positive factor would change the power level, but not the illuminant. Unlike colour science, computational colour constancy sometimes (but not always) uses the term *illuminant* as a synonym for light source. In this paper, the terms will be distinguished when needed, but otherwise the meaning can be inferred from context. In the problem of interest, the lighting for a scene will be assumed to be consistent with a single illuminant. The light, however, can be totally or partially diffused, or have both directional and diffuse components, etc., so the power level of the light can be different at different points of the scene.

## 2.2 Objects and Scenes

A sensing device makes an image of an illuminated scene, which contains various physical objects. The term *object* should be understood broadly, to include natural items such as trees and lawns, as well as artificial surfaces such as painted canvases. Obviously, an object’s colour affects the image that the device produces.

The objects considered in this paper have a physically invariant *local colour*, that depends only on the object’s *reflectance spectrum*. A reflectance spectrum,  $S(\lambda)$ , gives the percentage of light that the object would reflect if it were illuminated solely with light of wavelength  $\lambda$ .  $S(\lambda)$  takes on values between 0 and 100%. Figure 4 shows an example reflectance spectrum, for chromeoxide green, an artist’s pigment.

As with SPDs, the monochromatic reflectance spectra

$$S_i(\lambda) = \begin{cases} 1, & \text{if } \lambda = \lambda_i, \\ 0, & \text{otherwise,} \end{cases} \quad (4)$$

defined at the same 31 wavelengths, form a “basis” for an arbitrary spectrum:

$$S(\lambda) = \sum_{i=1}^{31} \alpha_i S_i(\lambda) \quad (0 \leq \alpha_i \leq 1 \forall i). \quad (5)$$

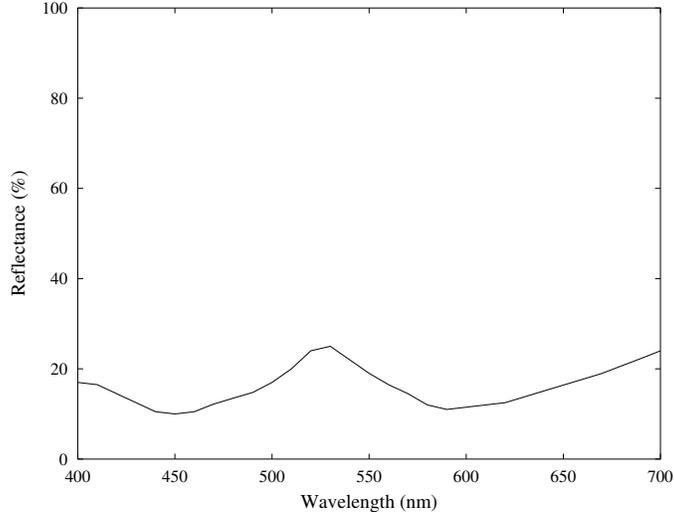


Figure 4: Reflectance Spectrum for Chromeoxide Green

Note that Equation (5) restricts each coefficient  $\alpha_i$  to be between 0 and 1.

The scenes considered in this paper will consist only of Lambertian objects. No highlights or gloss effects will occur, and no light sources will be visible. While somewhat artificial, these restrictions still define an important case that is an active area of research.

### 2.3 Sensors

In the problem of interest, light with SPD  $I(\lambda)$  reflects off an object with reflectance spectrum  $S(\lambda)$ , and the reflected light enters a sensor device. We shall call the reflected light a *sensor stimulus*, or just a stimulus. The stimulus,  $C(\lambda)$ , is calculated by multiplying the illuminating light’s SPD by the reflectance spectrum, wavelength by wavelength:

$$C(\lambda) = I(\lambda)S(\lambda) \quad (6)$$

$$= \sum_{i=1}^{31} \alpha_i P_i I_i(\lambda) S_i(\lambda). \quad (7)$$

Once the stimulus  $C(\lambda)$  enters the sensor device, it impinges, independently, on three individual sensors, or receptors. Typically, one of the sensors is most sensitive to the red part of the spectrum, another is most sensitive to the green part, and the third is most sensitive to blue, so we speak of red, green, and blue (RGB) sensors, making up an RGB device. Each sensor has a wavelength-dependent *response function*, or *curve*, denoted  $\rho_r(\lambda)$ ,  $\rho_g(\lambda)$ , and  $\rho_b(\lambda)$ . Figure 5 shows a hypothetical example of three response curves for one sensor device.

On a practical level, sensor response curves, if they are not already known, can be determined by experiment. Monochromatic light of wavelength  $\lambda_i$  can be shone on a white tile, and an image made. The entries of the *RGB* vector for the tile, after normalizing for the light power and the tile’s reflectance value at  $\lambda_i$ , give the values for  $\rho_r(\lambda_i)$ ,  $\rho_g(\lambda_i)$ , and  $\rho_b(\lambda_i)$ .

The sensor usually converts the light that impinges on it into some form of electricity, which is the sensor’s output. This paper will assume, as is generally the case, that each

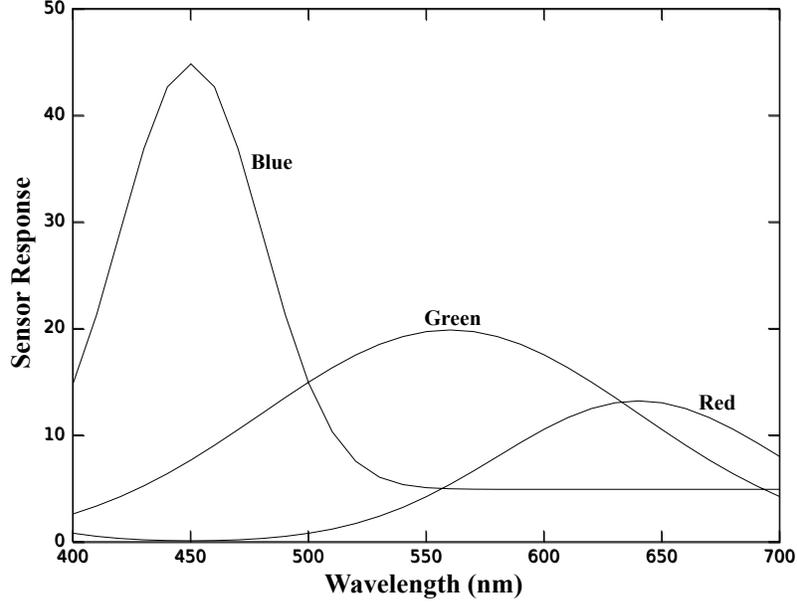


Figure 5: Some Example Response Curves

sensor responds linearly at each wavelength, so that its output is a constant factor times the impinging stimulus. These constants are given by the sensor response functions, whose outputs will be denoted  $R$ ,  $G$ , and  $B$ . The linearity allows us to write

$$R(C) = \sum_{i=1}^{31} \alpha_i P_i I_i(\lambda) S_i(\lambda) \rho_r(\lambda_i), \quad (8)$$

$$G(C) = \sum_{i=1}^{31} \alpha_i P_i I_i(\lambda) S_i(\lambda) \rho_g(\lambda_i), \quad (9)$$

$$B(C) = \sum_{i=1}^{31} \alpha_i P_i I_i(\lambda) S_i(\lambda) \rho_b(\lambda_i). \quad (10)$$

Since  $R$ ,  $G$ , and  $B$  are not functions of wavelength, and since the terms  $P_i(\lambda)$  and  $S_i(\lambda)$  are 1 when  $\lambda = \lambda_i$ , and 0 otherwise, we can drop those terms, getting

$$R(C) = \sum_{i=1}^{31} \alpha_i P_i \rho_r(\lambda_i), \quad (11)$$

$$G(C) = \sum_{i=1}^{31} \alpha_i P_i \rho_g(\lambda_i), \quad (12)$$

$$B(C) = \sum_{i=1}^{31} \alpha_i P_i \rho_b(\lambda_i). \quad (13)$$

The total output for all three sensors is  $(R, G, B)$ , which is a vector in a three-dimensional vector space.  $R$ ,  $G$ , and  $B$  will also denote a set of axes in that vector space. The components

of sensor output vectors are assumed to be non-negative, so only the positive octant,  $\mathcal{O}$ , of the vector space is needed. Since each component in  $(R, G, B)$  is a linear function of a stimulus, the output vector as a whole is a linear transformation of the set of stimuli. In fact, we can combine Equations (11) through (13) to get

$$(R, G, B)(C) = \sum_{i=1}^{31} \alpha_i P_i(\rho_r(\lambda_i), \rho_g(\lambda_i), \rho_b(\lambda_i)). \quad (14)$$

A sensing device generally produces a two-dimensional image, consisting of a rectangular grid of pixels. Each pixel corresponds to a different direction in which the device “looks at” the scene. An image therefore consists of a multitude of sensor output vectors, corresponding to the points of the scene.

## 2.4 The Illuminant Estimation Problem

A light source, a scene consisting of objects, and a sensor device with three individual sensors, are the physical elements of the illumination estimation problem. The device as a whole outputs a multi-pixel image of that scene under that light source. The illumination estimation problem is to infer the light source from such an image. This problem is challenging because the sensor stimulus for each pixel is the product of an SPD and a reflectance spectrum, and there seems to be no way to disentangle the SPD, which contains illumination information, from the reflectance spectrum, which contains object information.

Even though no pixel by itself provides sufficient information for IE, Forsyth<sup>2</sup> argued that the outputs from the set of pixels as a whole *can* provide sufficient information. In the aggregate, the set of pixel outputs, as a subset of  $RGB$  space, is called the *image gamut*. Another kind of gamut, called the *illuminant gamut*, can also be constructed. Consider the set of all possible object colours: imagine a sort of comprehensive Color Checker—a physical surface painted with a very large grid of flat squares, such that every reflectance spectrum (or at least a sufficiently close approximation) occurs in one of the squares. Suppose that the sensor device makes an image of that grid under a particular light source. Then the resulting set of pixel outputs, which again is a subset of  $RGB$  space, is called that light source’s illuminant gamut. When an actual scene is illuminated by that light source, the sensor output for every pixel corresponding to an object in that scene, and therefore the image gamut as a whole, is contained in the illuminant gamut for that source. Forsyth’s approach is to calculate the image gamut for a particular image, and then find an illuminant gamut that both contains the image gamut, and is a good fit for it. The light source that produces such an illuminant gamut should be a good estimate for the actual light source. This approach is called *gamut-based illuminant estimation* (GBIE).

This paper works within the GBIE framework. In the next section, the recently presented formulations of sources, objects, and sensors will be used to construct gamuts mathematically. The section after that will apply those constructions to GBIE implementations to date, suggesting some improvements.

### 3 Geometric Constructions

The previous section formulated the physical phenomena, sensor properties, and mathematical relationships that determine the outputs of a sensing device. The illuminant gamuts used in GBIE were also introduced. The current section will build on the formulations, step by step, to construct an illuminant gamut geometrically. The first step is the *sensor spectrum locus*, which is the set of vectors, in  $RGB$  space, that are sensor outputs, for a particular light source, of maximal monochromatic reflectance spectra. The second step is the *sensor spectrum cone*, which is the total set of  $RGB$ s that the device can produce; this set includes the  $RGB$  produced by any reflectance spectrum, under any light source. The third step is the *sensor chromaticity diagram*, which is the set of all possible chromaticities of sensor  $RGB$  outputs. The chromaticity of an  $RGB$  vector depends on the ratios of its  $R$ ,  $G$ , and  $B$  components, so is constant along rays in  $RGB$  space. A somewhat surprising fact is that the chromaticity diagram and spectrum cone are both independent of illumination: they depend only on the sensor response curves. The spectrum locus, however, depends on both response curves and illumination.

These three steps culminate in the final construction, an illuminant gamut, which also depends on both response curves and illumination. The chief new result of this paper is that each illuminant gamut is a zonohedron. A zonohedron is a convex polyhedron that is the *Minkowski sum* of the vectors in the spectrum locus. This result is important because it allows illuminant gamuts to be calculated explicitly. GBIE approaches to date estimate illuminant gamuts from a limited training set of reflectance spectra, and use no information, except convexity, about the gamut’s form. The zonohedral approach, on the other hand, encompasses *all* theoretically possible reflectance spectra, and uses the information that the gamut is not just convex, but in fact zonohedral. The next section will apply the zonohedral approach to GBIE implementations.

#### 3.1 The Sensor Spectrum Locus

For a given light source  $I$ , a device’s sensor spectrum locus, denoted  $\mathcal{L}_I$ , is the set of  $RGB$  output vectors that result from maximal monochromatic reflectance spectra, when imaged under  $I$ . Such a spectrum reflects 100% of the light at the wavelength  $\lambda_j$ , and no light at any other wavelength; in terms of Equation (5),  $\alpha_j$  is 1 and all other  $\alpha$ ’s are 0. From Equation (3), the power of the light source at  $\lambda_j$  is  $P_j$ . The resulting sensor stimuli  $C_j$  for the 31 wavelengths are therefore given by

$$C_j = P_j I_j(\lambda), \quad j = 1, 2, \dots, 31. \quad (15)$$

Evaluating the output vectors for the  $C_j$ ’s with Equations (11) through (13) gives

$$\mathcal{L}_I = \left\{ P_i(\rho_r(\lambda_i), \rho_g(\lambda_i), \rho_b(\lambda_i)) \mid i = 1, 2, \dots, 31 \right\}. \quad (16)$$

The expression  $P_i(\rho_r(\lambda_i), \rho_g(\lambda_i), \rho_b(\lambda_i))$  is just the  $RGB$  output for the  $i^{\text{th}}$  maximal monochromatic reflectance spectrum, and is called the  $i^{\text{th}}$  *sensor spectrum locus vector*. The set of spectrum locus vectors as a whole is called the *sensor spectrum locus*.

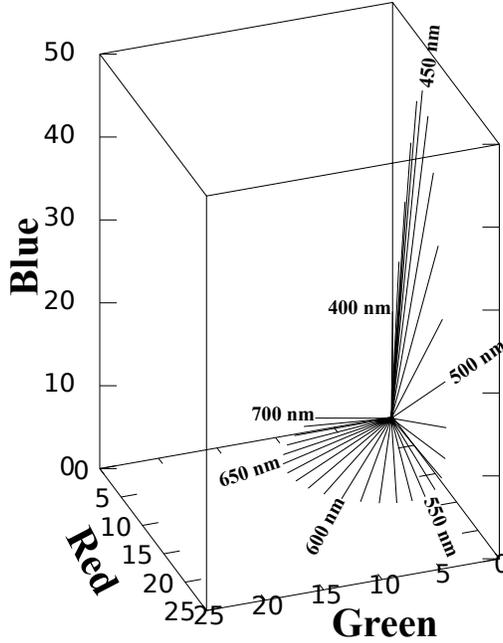


Figure 6: The Sensor Spectrum Locus for the Response Curves in Fig. 5, and Illuminant E

In  $RGB$  space, a sensor spectrum locus can be drawn as a sequence of 31 vectors that start at the origin. Since device outputs are non-negative, the sensor spectrum locus is restricted to the non-negative octant  $\mathcal{O}$  of  $RGB$  space. Figure 6 shows the locus of the sensor response curves in Fig. 5. The light source in this example is assumed to set each  $P_i$  to 1. (Such constant-power light sources are said to be consistent with Illuminant E.)

Suppose another light source was used, whose power level at 500 nm was 2 instead of 1. Then the 500 nm vector in Fig. 6 would be in the same direction, but twice as long. Similarly, the vector would be half the length if the 500 nm power level were  $1/2$ . The vectors for the other wavelengths would also shrink or expand, depending on the power levels for the new source—but in all cases the vectors would maintain their directions. A device’s spectrum locus will therefore differ for different light sources, but the locus vectors’ directions are the same for all sources.

### 3.2 The Sensor Spectrum Cone

A device can produce images of objects under a multitude of light sources. The totality of  $RGB$  outputs that can be theoretically produced, for any light source, occupies the sensor spectrum cone, which is a subset of  $\mathcal{O}$  in  $RGB$  space. We will show that the spectrum cone is the convex cone of the spectrum locus.

The set of all non-negative linear combinations of a set of  $n$  vectors,  $\{\mathbf{v}_i\}$ , in a vector space is called a *convex cone*, and is formally defined as

$$\left\{ \sum_{i=1}^n \beta_i \mathbf{v}_i \mid \beta_i \geq 0 \ \forall i \right\}. \quad (17)$$

Geometrically, a convex cone is constructed by extending each vector  $\mathbf{v}_i$  into a semi-infinite

ray, and then taking the convex hull of all those rays. (The convex hull of a set  $\mathcal{S}$  can be defined to be the smallest convex set that contains  $\mathcal{S}$ .) Equivalently, the extended vectors define a roughly conical surface in  $\mathcal{O}$  that radiates outward from the origin, and the convex cone consists of all vectors “inside” that surface. Every vector inside the cone can be extended to a ray, which is also inside the cone, so a convex cone equals the union of all the rays it contains.

Now suppose that the  $\mathbf{v}_i$ ’s are taken to be the locus vectors  $\mathcal{L}_I$  for the light source  $I$ , with coefficients  $P_i$ . Then the convex cone of  $\mathcal{L}_I$  is given by

$$\mathcal{C}_I = \left\{ \sum_{i=1}^{31} \beta_i P_i (\rho_r(\lambda_i), \rho_g(\lambda_i), \rho_b(\lambda_i)) \mid \beta_i \geq 0 \forall i \right\}. \quad (18)$$

$\beta_i$  and  $P_i$  are both non-negative, so their product is non-negative. Thus any coefficient  $P_i$  could be absorbed into the single coefficient  $\beta_i$ , to produce

$$\mathcal{C}_I = \left\{ \sum_{i=1}^{31} \beta_i (\rho_r(\lambda_i), \rho_g(\lambda_i), \rho_b(\lambda_i)) \mid \beta_i \geq 0 \forall i \right\}. \quad (19)$$

This argument would apply to any light source. Assume for instance that there is a second light source  $I_2$ , with coefficients  $P_{i2}$ . Then the expression for  $\mathcal{C}_{I_2}$  would be identical to Equation (18), except that  $P_i$  would be replaced by  $P_{i2}$ . After absorbing each  $P_{i2}$  into  $\beta_i$ , we see that the expression for  $\mathcal{C}_{I_2}$  is identical to the expression for  $\mathcal{C}_I$  in Equation (19). An important conclusion is that the sensor spectrum cone, which is the convex cone of the locus vectors, is the same, no matter which light source is used for the locus. The convex cone, then, depends only on the sensor response curves. Thus we will drop the subscript in Equation (19) and just denote the sensor spectrum cone by  $\mathcal{C}$ .

We will now show that  $\mathcal{C}$  consists exactly of the *RGB* sensor outputs for any local colour, under any illumination. A local colour is defined by its reflectance spectrum,  $S(\lambda)$ , as given in Equation (5). A light source,  $I$ , is defined by its SPD,  $I(\lambda)$ , as given in Equation (3). Equation (14) gives the entries of the *RGB* output. The terms  $\alpha_i$  and  $P_i$  are both non-negative, so they can be combined into one non-negative  $\beta_i$ . This adjustment makes it clear that the *RGB* output vector is of the form given in Equation (19), so the *RGB* vector is in the sensor spectrum cone. Conversely, given a cone vector of the form in Equation (19), one could let  $\alpha_i = 1$ , and  $P_i = \beta_i$  (as long as  $\beta_i$  is not 0; if  $\beta_i$  is 0, then let  $\alpha_i = 0$ , and  $P_i = 1$ , to avoid zero power at the  $i^{\text{th}}$  wavelength). These new coefficients could be used to define a reflectance spectrum and light source whose *RGB* output is the initial cone vector. Thus every cone vector is contained in the set of *RGB* sensor outputs. The two-way containment proves that  $\mathcal{C}$  consists exactly of all possible *RGB* sensor outputs, as was to be shown.

The spectrum cone can be visualized by cutting it with a plane. Figure 7 shows the spectrum locus vectors from Figure 6, extended to rays. These rays, and any rays inside of them, make up the spectrum cone. In the figure, the rays have been truncated at the plane  $R + G + B = 1$ , to show the cone’s profile, which is the convex hull of the rays’ intersections with  $R + G + B = 1$ . Since the spectrum locus vectors are all non-negative, the sensor spectrum cone really is a geometrical cone, though not a circular one; while its profile is convex, it is also likely irregular.

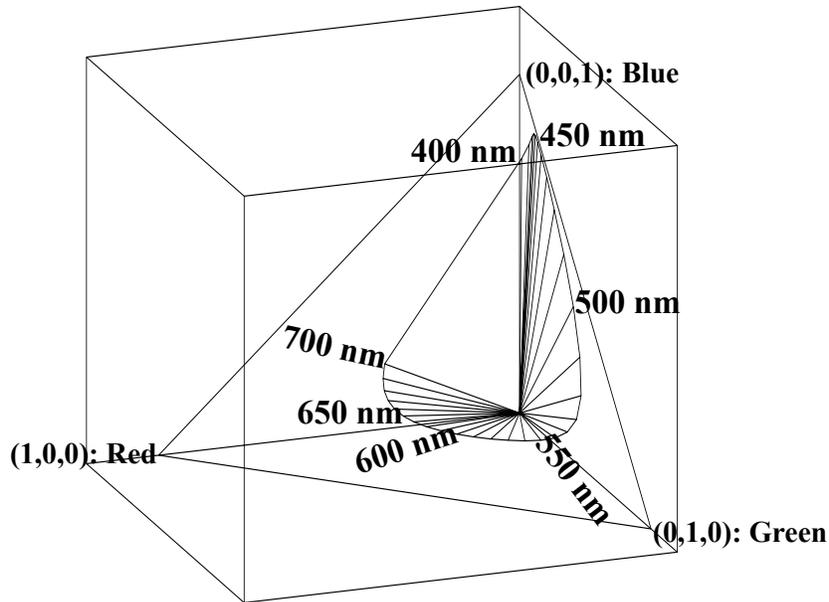


Figure 7: Profile of the Spectrum Cone

### 3.3 The Sensor Chromaticity Diagram

The dichotomy between locus vector direction and extent suggests that we can decompose a three-dimensional locus vector into a magnitude (which requires one dimension) and a direction (which requires two dimensions). The two-dimensional direction will be called *chromaticity*. The set of chromaticities can be displayed conveniently in a sensor chromaticity diagram, whose construction will now be given.

Chromaticities have an important physical interpretation, in terms of shadows. Even though an object in a scene has a single local colour, the sensor outputs for different points on that object will likely differ, because some parts of the object are in light, and other parts are in various degrees of shadow. The set of all sensor outputs for that local colour, in all degrees of shadow and light, is called a *shadow series*. The sensor stimuli for a shadow series share a common shape, or relative SPD, which is multiplied by larger factors when the object is in light, and by progressively smaller factors as the object is more deeply in shadow.<sup>5</sup> Chromaticity depends only on the relative SPD, so a shadow series has constant chromaticity.

The sensor stimuli of a shadow series impinge on the sensors, producing a set of *RGB* output vectors. Since the sensors are linear, the *RGB* vectors all fall along a single ray, called a *chromaticity ray*, through the origin. In fact, it is not hard to see that the sensor spectrum cone consists exactly of the chromaticity rays. Although the device reduces the series from a 31-dimensional SPD to a 3-dimensional *RGB*, the series' common relative SPD causes the *R*, *G*, and *B* entries of each vector on the chromaticity ray to have common ratios. These ratios can be completely described by two chromaticity coordinates, that will be used to construct the sensor chromaticity diagram.

As before, extend each spectrum locus vector in Figure 6 to a ray, and normalize those rays by cutting them with the plane  $R + G + B = 1$ , as shown in Figure 8. The plane

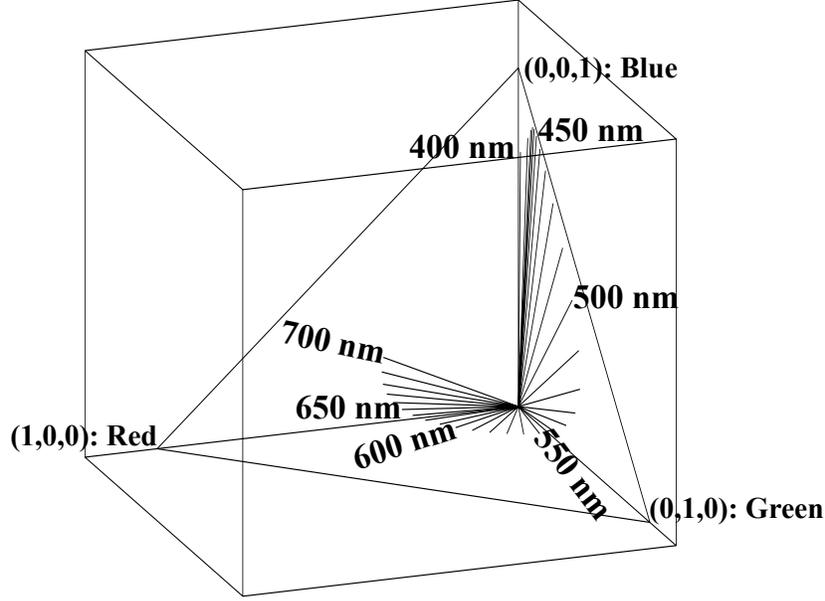


Figure 8: Normalized Spectrum Locus Vectors, that Intersect the Plane  $R + G + B = 1$

$R + G + B = 1$  intersects  $\mathcal{O}$  in the triangle with vertices  $(1, 0, 0)$ ,  $(0, 1, 0)$ , and  $(0, 0, 1)$ . The vertex  $(1, 0, 0)$  would correspond to a sensor stimulus (if any existed) that produced a pure  $R$  output, without any  $G$  or  $B$  output; the other two vertices have similar interpretations.

We can specify the ray for an output vector  $(R, G, B)$  by its intersection point,  $(r, g, b)$ , with the plane  $R + G + B = 1$ :

$$(r, g, b) = \frac{1}{R + G + B}(R, G, B). \quad (20)$$

The point  $(r, g, b)$  is on the ray and  $r + g + b = 1$ . Since  $b$  can always be calculated by  $b = 1 - r - g$ , the first two coefficients,  $r$  and  $g$ , are sufficient as chromaticity coordinates.

Figure 9 shows the triangle in Figure 8, with chromaticity coordinates. Only one point of each ray remains. Each monochromatic stimulus, via its spectrum locus vector, is located in the triangle, as labeled. The vertices correspond to pure red, green, and blue chromaticities; these are mathematical ideals that are likely not producible physically. The *chromaticity diagram*, shown in grey in the figure, is the convex hull of the points corresponding to the locus vectors. Apart from its smallest-convex-set definition, the convex hull of a set of  $n$  vectors  $\{\mathbf{v}_i\}$  can also be defined by

$$\text{co}(\{\mathbf{v}_i\}) = \left\{ \sum_{i=1}^n \alpha_i \mathbf{v}_i \mid 0 \leq \alpha_i \leq 1 \ \forall i, \sum_{i=1}^n \alpha_i = 1 \right\}. \quad (21)$$

Choose a point  $\mathbf{v}$  in the grey chromaticity diagram. Since each black point  $\mathbf{v}_i$  in Figure 9 corresponds to a monochromatic reflectance spectrum  $S_i$ , and since each point in the grey chromaticity diagram can be written as  $\sum \alpha_i \mathbf{v}_i$  for some set of  $\alpha_i$ 's, all less than 1, it follows that  $\mathbf{v}$  will occur as the chromaticity of an object with reflectance spectrum  $\sum \alpha_i S_i$ . Therefore the chromaticity diagram gives the entire set of chromaticities that can be produced.

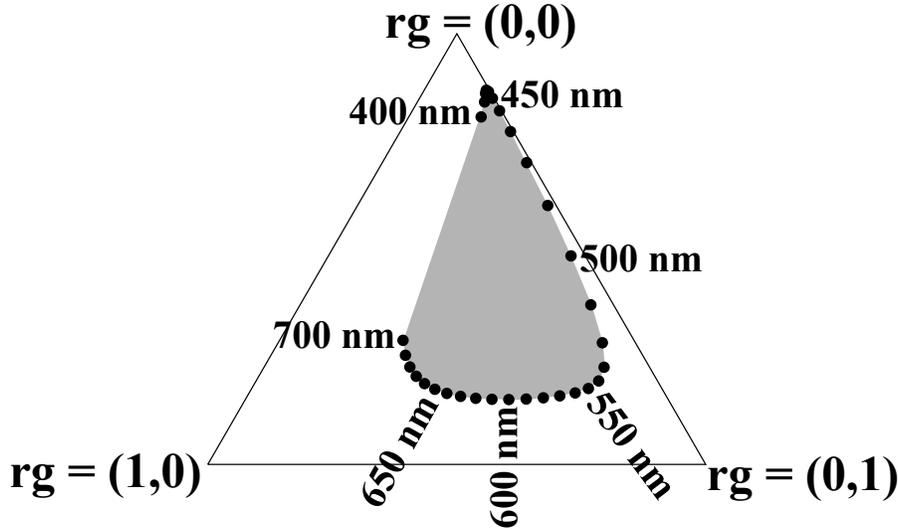


Figure 9: Chromaticity Diagram

Like the spectrum cone, but unlike the spectrum locus, the chromaticity diagram for a sensing device does not depend on the illuminant—instead, it depends only on the sensor response curves. Mathematically, this result follows from the fact that no particular power level is needed to define a ray, so the coefficients  $P_i$  in Equation (16), provided they are all nowhere zero, are not required when defining a chromaticity diagram.

A more surprising consequence is that *any* nowhere-zero illuminant, if it reflects off the right reflectance spectrum, can produce *any* chromaticity in the device’s chromaticity diagram. If some  $P_i$ ’s are very small, then the output for some chromaticities might be very dark, but nevertheless all chromaticities will be present. Furthermore, increasing the light source’s power level, without changing the source’s relative SPD, could make those dark chromaticities arbitrarily bright. Increasing the power level slides  $RGB$  outputs away from the origin, along the chromaticity rays, but does not change the fact that the spectrum cone consists precisely of all chromaticity rays, and the same set of rays occurs for all nowhere-zero illuminants.

In Figure 9, the intersection points are said to be *cyclic*, because they form a sequence around the boundary of their convex hull. In the terms of West and Brill,<sup>6</sup> a cyclic set is “convex and well ordered in wavelength.” Algebraically, no point is a positive linear combination of any subset of the other points. Geometrically, no point is in the convex hull of the other points. Cyclic locus vectors simplify some further constructions and calculations. In practice, the spectrum loci of many sensing devices are cyclic, simply because the device’s sensor responses mimic human photoreceptor responses, which are known to be cyclic.<sup>6</sup>

### 3.4 Minkowski Sums and Zonohedra

After introducing the concepts of Minkowski sums and zonohedra, the previous constructions will be used to construct zonohedral illuminant gamuts in  $RGB$  space. Some previous colour science papers<sup>7,8</sup> can be consulted for further details.

The Minkowski sum (also called the *vector sum*<sup>9</sup>) of two sets,  $\mathbf{A}$  and  $\mathbf{B}$ , in  $\mathbb{R}^n$ , is defined

as

$$\mathbf{A} \oplus \mathbf{B} = \{a + b | a \in \mathbf{A}, b \in \mathbf{B}\}. \quad (22)$$

The addition on the right side of Equation (22) is ordinary vector addition in  $\mathbb{R}^n$ . Minkowski addition is commutative:  $\mathbf{A} \oplus \mathbf{B}$  is the same as  $\mathbf{B} \oplus \mathbf{A}$ . It is also associative:  $(\mathbf{A} \oplus \mathbf{B}) \oplus \mathbf{C}$  is the same as  $\mathbf{A} \oplus (\mathbf{B} \oplus \mathbf{C})$ . Thus we can unambiguously write  $\mathbf{A} \oplus \mathbf{B} \oplus \mathbf{C}$ , allowing the Minkowski sum to be defined for an arbitrary number of sets.

Geometrically, the Minkowski sum of  $\mathbf{A}$  and  $\mathbf{B}$  is the shape covered by all the copies of  $\mathbf{A}$ , whose centers touch  $\mathbf{B}$  at at least one point. In fact, one could visualize  $\mathbf{A} \oplus \mathbf{B}$  as the set of points that is covered whenever a copy of  $\mathbf{A}$  is added to every point in  $\mathbf{B}$ . By commutativity, the same new set would result by adding a copy of  $\mathbf{B}$  to every point in  $\mathbf{A}$ .

Zonohedra result from a special case of Minkowski summation, when the summands are all line segments starting at the origin. Such a line segment can be equally well represented by a vector whose tail is the origin and whose head is the farthest point of the line segment. Suppose that we have a set of  $n$  non-negative vectors,  $\{\mathbf{v}_i\}$ , in  $\mathbb{R}^3$ . Then the *zonohedron* generated by  $\{\mathbf{v}_i\}$  is the Minkowski sum of the line segments corresponding to those vectors:

$$Z(\{\mathbf{v}_i\}) = \mathbf{v}_1 \oplus \mathbf{v}_2 \oplus \dots \oplus \mathbf{v}_n. \quad (23)$$

Equivalently,

$$Z(\{\mathbf{v}_i\}) = \left\{ \sum_{i=1}^n \alpha_i \mathbf{v}_i \mid 0 \leq \alpha_i \leq 1 \forall i \right\}. \quad (24)$$

The vectors  $\mathbf{v}_i$  are referred to as *generating vectors*, or simply *generators*.

Zonohedra offer considerable geometric structure. A zonohedron is always a convex polyhedron. Furthermore, each edge of that polyhedron is a translated copy of some generator, and, if no three generators are coplanar, then every face is a parallelogram. Every vertex of a zonohedron can be represented as a sum of generators, and this representation is unique. Algebraically, vertices occur only (but not always) when every  $\alpha_i$  in Equation (24) is either 0 or 1. A zonohedron has a vertex at the origin, when every  $\alpha_i$  is 0, and its *terminal vertex* occurs when every  $\alpha_i$  in Equation (24) is 1; the terminal vertex is simply the sum of all the generators. A zonohedron is convex, and centrally symmetric about the point

$$\sum_{i=1}^n \frac{1}{2} \mathbf{v}_i. \quad (25)$$

The problem of determining a zonohedron's vertices, edges, and face, when given only its generators, is referred to as *solving* the zonohedron. Ref. 10 suggests some algorithms for the general case. When the generators are cyclic, however, pp. 113-115 of Ref. 7 describes a simple algorithm involving an easily constructed table.

### 3.5 The Zonohedral Illuminant Gamut

The previous constructions and definitions, combined with the assumed linearity of the sensing device, have laid the groundwork for this paper's main result: an illuminant gamut

is a zonohedron in  $RGB$  space. More particularly, this section will show that an illuminant gamut is the Minkowski sum of the sensor spectrum locus vectors.

For a sensing device and light source  $I$ , the illuminant gamut consists of all the  $RGB$  vectors that could occur in an image of any object, when illuminated by  $I$ . An object's colour is defined by a reflectance spectrum with 31 coefficients  $\alpha_i$ , all between 0 and 1 (see Equation (5)). Illumination by  $I$ , with power coefficients  $P_i$ , produces the sensor stimulus  $C$ , given by Equation (7). Equation (14) gives the  $RGB$  output for  $C$ :

$$(R, G, B)(C) = \sum_{i=1}^{31} \alpha_i [P_i(\rho_r(\lambda_i), \rho_g(\lambda_i), \rho_b(\lambda_i))]. \quad (26)$$

As the  $\alpha_i$ 's vary over all allowable combinations, producing all possible reflectance spectra, Equation (26) will produce all  $RGB$ 's in the illuminant gamut.

The form of Equation (26) shows that it can be interpreted as a zonohedron. Write

$$\mathbf{v}_i = P_i(\rho_r(\lambda_i), \rho_g(\lambda_i), \rho_b(\lambda_i)). \quad (27)$$

Then  $\mathbf{v}_i$  is the  $i^{\text{th}}$  spectrum locus vector. Furthermore, all the  $\alpha_i$ 's must be between 0 and 1, but are otherwise arbitrary. The total output vector for  $C$ , on the left side of Equation (27), is therefore in the zonohedron generated by the  $\mathbf{v}_i$ 's. Since each  $\alpha_i$  is arbitrary, the converse follows: each vector in the zonohedron results from an object colour. Formally, let

$$Z_I = \left\{ \sum_{i=1}^m \alpha_i [P_i(\rho_r(\lambda_i), \rho_g(\lambda_i), \rho_b(\lambda_i))] \mid 0 \leq \alpha_i \leq 1 \forall i \right\}. \quad (28)$$

Then  $Z_I$  is the illuminant gamut for the light source  $I$ , and is the zonohedron generated by the spectrum locus vectors for  $I$  (the set  $\mathcal{L}_I$ ), as was to be shown. As an example, Figure 1 shows the zonohedral illuminant gamut for the response curves in Figure 5, when  $I$  has constant power 1 at each wavelength (such an SPD is an instance of Illuminant E).

A zonohedron's terminal point is the sum of all its generating vectors. An illuminant gamut's generating vectors consist of that illuminant's spectrum locus vectors, each of which corresponds to a reflectance spectrum with 100% reflectance at exactly one wavelength. The terminal point therefore corresponds to the spectrum with 100% reflectance at *all* wavelengths, which is an ideal white. The origin occurs when every  $\alpha_i$  is zero, so it corresponds to an ideal black. By convexity and linearity, the line joining the origin to the terminal point contains all the ideal greys, that is, object colours which reflect the same percentage of light at each wavelength.

While knowing the terminal point for a particular illuminant is helpful, surprisingly that knowledge does not define the illuminant, because the loci of two different illuminants can sum up to the same terminal point, even though the individual locus vectors have different lengths. One could therefore visualize a set of zonohedra, all fixed at the origin, all having the same terminal point, and all containing the same grey axis, but having different polyhedral shapes. The variability in shape depends of course on the sensors' response functions.

As discussed earlier, SPDs can be seen as varying independently in chromaticity and power level. Multiplying a light source by a positive constant  $k$  will dilate (if  $k > 1$ ) or contract (if  $k < 1$ ) its zonohedral gamut. If the power is doubled, for example, each vector

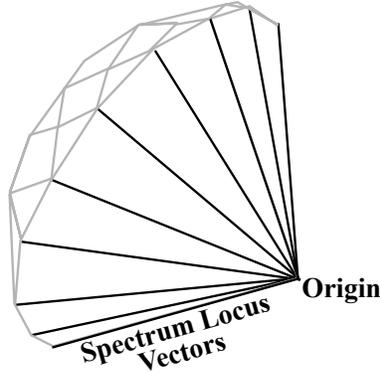


Figure 10: The Structure of a Zonohedral Illuminant Gamut at the Origin of  $RGB$  Space

in the zonohedron will double in length, but without any change in direction. The origin will remain unchanged, but the grey axis will stretch to twice its length, moving the terminal point twice as far from the origin. These considerations of power level versus chromaticity will be relevant later, for a proposed GBIE algorithm.

An important—and very counterintuitive—result is that a zonohedral gamut, for any non-zero illuminant, fits perfectly into the spectrum cone at the origin.<sup>8</sup> Figure 10 shows that a brilliant-cut diamond, in which many edges converge at a sharp point, has the same structure as a zonohedral gamut at the origin, where the spectrum locus vectors converge. The surface of the spectrum cone is constructed by extending the locus vectors to rays, so a zonohedral gamut must fit perfectly inside the spectrum cone. The spectrum cone is unique, and depends only on the sensor responses, not on the illuminant, so *every* zonohedral illuminant gamut (at least for a nowhere-zero illuminant) fits perfectly inside the cone.

A corollary is that any two (nowhere-zero) illuminant gamuts must fit perfectly inside each other, and thus have the same shape and location, in some small neighborhood of the origin. Figure 2 shows two such gamuts, for Illuminants A and E, both using the sensor responses in Figure 5. The spectrum locus vectors for the two illuminants all have the same orientation, but can differ in length. Away from the origin, the length differences give the illuminant gamuts different shapes, and can significantly separate their terminal points, as seen in the figure. At the origin, however, their shapes must be identical. In the next section, the common shape at the origin will imply that one illuminant gamut is almost never a linear transformation of another illuminant gamut, as many GBIE algorithms assume.

## 4 Applications to Illuminant Estimation

A GBIE algorithm starts with an image of a scene, and uses illuminant gamuts to make inferences about the illumination for that scene. The previous section’s geometric constructions suggest some improvements to current algorithms, and also show that some common techniques are theoretically unsound. The main positive improvement is the use of zonohedral gamuts, which produces a complete illuminant gamut by direct calculation; current methods involve training sets, and do not fill out the complete gamut.

There are also two negative results. First, linear transformations are often used to trans-

fer the  $RGB$ 's in a canonical gamut to the  $RGB$ 's in some other illuminant gamut; we will show a linear image of one illuminant gamut is almost never another illuminant gamut. Second, Finlayson recommends recasting gamut containment problems as chromaticity containment problems<sup>4,11,12</sup> Since the chromaticity diagram depends only on sensor response, and not on illumination, however, chromaticity containment tests cannot distinguish between illuminants.

To take advantage of the underlying geometry, and to avoid techniques contraindicated by that geometry, a new GBIE algorithm is suggested. The new algorithm uses zonohedral gamuts, decomposes illumination into chromaticity and power level, and finds the illumination of minimum power that contains the image gamut.

## 4.1 Constructing Zonohedral Illuminant Gamuts

Many implementations of Forsyth's GBIE algorithm require constructing the illuminant gamut for one or more known illuminants. Previous constructions have relied on training sets of object colours, taken from, for instance, Munsell atlases or spectral databases.<sup>12</sup> Typically, the convex hull of the  $RGB$ 's of the training set is taken as the illuminant gamut. Since physical examples of near-optimal reflectance spectra, especially dark ones, are rare, however, the resulting illuminant gamut is not complete. In addition, the training set method provides no geometric insight about the gamut.

This paper suggests zonohedral construction instead, which automatically incorporates all theoretically possible reflectance spectra. The comprehensiveness insures that the zonohedral gamut is complete, containing even colours that are very dark, light, or saturated. In addition, the gamut's polyhedral structure, with explicit vertices, edges, and faces, can be calculated directly, and provides some geometric intuition for further use. These two advantages recommend zonohedral construction methods over training set methods.

## 4.2 Linear Feasible Mappings

Forsyth's original GBIE algorithm<sup>2</sup> used the *canonical gamut*  $G_{\text{can}}$ , for a known illuminant  $I_{\text{can}}$ , and the image gamut  $G_{\text{im}}$ , which was the convex hull of all the  $RGB$ 's in an image. Likely  $G_{\text{can}}$  did not contain  $G_{\text{im}}$ , so  $I_{\text{can}}$  was not the illuminant used for the image. The problem then is to find an illuminant gamut  $G_I$ , such that

$$G_I \supset G_{\text{im}}. \quad (29)$$

Forsyth did not require that the illuminant  $I$  itself be found, as long as the  $RGB$ 's that would have occurred under the canonical illuminant could be calculated. Therefore, his algorithm only looked for a non-canonical gamut  $G_I$  that satisfied Equation (29), rather than an explicit  $I$ .

The original implementation estimated non-canonical illuminant gamuts from the canonical gamut. After some analysis, the relationship between  $G_I$  and  $G_{\text{can}}$  was written as a linear operator  $M$ , plus a residual error term (Equation (3) in the 1990 paper). Dropping the residual term left just  $M$  for the transformation:

$$G_I = M(G_{\text{can}}). \quad (30)$$

Conversely, if  $M$  satisfied certain conditions, such as being full-rank and preserving the positive octant, then  $M(G_{\text{can}})$  would be the gamut for some illuminant. Later implementations often used the much simpler condition, based on von Kries’s coefficient rule,<sup>1</sup> that  $M$  was diagonal with positive entries. Any  $M$  such that  $M(G_{\text{can}}) \supset G_{\text{im}}$  is called a *feasible mapping*. Many implementations of Forsyth’s algorithm find a set of feasible mappings, and then use a minimizing criterion to select one that gives the best estimate of  $G_I$ .

Unfortunately, our geometric constructions show that linear feasible mappings such as  $M$  are theoretically untenable. Section 3.5 shows that the canonical gamut is a zonohedron,  $Z_{\text{can}}$ , that fits perfectly into the spectrum cone,  $\mathcal{C}$ , at the origin. Since linear operators such as  $M$  preserve convexity relationships,  $M(Z_{\text{can}})$  is another zonohedron, and  $M(\mathcal{C})$  is another convex cone. Furthermore,  $M(Z_{\text{can}})$  fits perfectly inside  $M(\mathcal{C})$  at the origin. To be an illuminant gamut, though,  $M(Z_{\text{can}})$  must fit perfectly inside  $\mathcal{C}$  at the origin, which is only possible if  $\mathcal{C}$  and  $M(\mathcal{C})$  are the same set. Since a spectrum cone’s profile is usually irregular, however, it is very unlikely that the image under a linear  $M$  of an illuminant gamut is again an illuminant gamut.

As a consequence, the linear method of estimating further illuminant gamuts produces actual illuminant gamuts only by accident. In implementations, a null set of feasible mappings is a likely result. Finlayson *et al.* (p. 8 of Ref. 12) prevent such null sets by not using linear transformations at all: their Gamut Constrained Illuminant Estimation (GCIE) algorithm estimates new illuminant gamuts directly from a training set. In general, the underlying geometry contraindicates looking for linear feasible mappings; direct illuminant gamut calculations are preferable, especially when zonohedral calculations are used in place of training sets.

### 4.3 Chromaticity Containment Tests

Forsyth’s 1990 GBIE algorithm works in three-dimensional  $RGB$  space, checking whether an image gamut is contained in various illuminant gamuts. In 1996, Finlayson<sup>4</sup> suggested a containment check in two-dimensional chromaticity space: replace the set of all  $RGB$ s that an illuminant  $I$  can produce with the set of all chromaticities,  $\text{Chrom}(I)$ , that  $I$  can produce; similarly, use the image’s chromaticities, rather than its  $RGB$  outputs. Then  $I$  is a possible illuminant only if the image’s chromaticities are a subset of  $\text{Chrom}(I)$ .

Like linear feasible mappings, unfortunately, the chromaticity approach is also theoretically untenable. Sect. 3.3 shows that *any* nowhere-zero illuminant can produce *any* chromaticity in the device’s chromaticity diagram. Formally, given two nowhere-zero illuminants,  $I_1$  and  $I_2$ , we always have  $\text{Chrom}(I_1) = \text{Chrom}(I_2)$ —and both of them equal the chromaticity diagram, which only depends on the sensors’ response functions. As a consequence, chromaticity containment tests are vacuous, because they are always satisfied. Mathematically, Sect. III of Ref. 11 derived chromaticity criteria from linear feasible mappings—which we have also seen to be untenable. Likely, the incomplete gamuts resulting from limited training sets obscured the fact that different illuminants produce the same chromaticities. The geometrical constructions, however, make it clear.

In fact, we can see geometrically that any nowhere-zero illuminant can produce any  $RGB$  in the spectrum cone. Suppose we have a light source  $I$ , with zonohedral gamut  $G_I$ , and a pixel in a scene has an output vector  $(R, G, B)$ , that is not contained in  $G_I$ . Then it would

seem that  $I$  is not the illuminant for that scene—but that conclusion fails to distinguish an illuminant (in the sense of a relative SPD), from a power level. Since  $G_I$  fits perfectly inside the spectrum cone at the origin, the chromaticity ray through  $(R, G, B)$  must intersect  $G_I$ . If we multiply  $I$  by a sufficiently large positive factor  $k$ , however, then  $G_{kI}$  *would* contain  $(R, G, B)$ . Geometrically, multiplication by  $k$  dilates the zonohedral gamut, changing its size but not its shape.  $I$  and  $kI$  have the same relative SPD, but different power levels. A powerful enough light source, then, of any nowhere-zero relative SPD, can produce any pixel output that is possible for a particular sensor device. The new GBIE algorithm suggested in the next section explicitly accounts for this fact.

#### 4.4 A New GBIE Algorithm

The last two sections have been rather negative, arguing that Forsyth’s illuminant gamut algorithm is not actually working with illuminant gamuts, and that Finlayson’s chromaticity containment criterion is vacuous. This section takes a more positive approach, by suggesting a new GBIE algorithm that takes advantage of this paper’s geometric constructions.

A light source can be decomposed into a power level and a relative SPD. The relative SPD has a well-defined chromaticity in the sensor’s chromaticity diagram. The new algorithm takes as input the set  $\mathcal{P}$  of all pixel  $RGB$ ’s in an image, and estimates the chromaticity and power level of the light source under which the image was taken.

Suppose we hypothesize a certain chromaticity,  $(r, g)$ , and a certain power level. Then there is a set of metameric light sources of that power and chromaticity. Choose one  $I$  from this set, and find its zonohedral gamut  $Z_I$ . If  $Z_I$  does not contain  $\mathcal{P}$ , then multiply the power level of  $I$  by the smallest positive factor  $k > 1$  such that  $Z_{kI}$  does contain  $\mathcal{P}$ . If, on the other hand,  $Z_I$  already contains  $\mathcal{P}$ , then multiply by the smallest positive factor  $k < 1$  such that  $Z_{kI}$  still contains  $\mathcal{P}$ . The result either way is the minimum power needed for a light source of chromaticity  $(r, g)$  to produce that image. Define the function

$$f : \text{chromaticity diagram} \longrightarrow \mathbb{R}, \tag{31}$$

which assigns the minimum power to any set of chromaticity coordinates. Now minimize  $f$ . That is, find chromaticity coordinates  $(r_{\min}, g_{\min})$  such that  $f(r_{\min}, g_{\min})$  is as small as possible. Take  $(r_{\min}, g_{\min})$  to be the chromaticity we are seeking, and  $f(r_{\min}, g_{\min})$  to be the power level.

Geometrically,  $f$  shrinks or expands zonohedral gamuts, without changing their shapes, until they are just large enough to contain  $\mathcal{P}$ , but no larger. Physically,  $f$  finds the minimum power needed for a light source of a particular chromaticity to produce all the colours seen in the image. Perceptually,  $f$  implements an heuristic: a scene’s illumination level is the minimum needed to make the scene appear as it does.

While many details are still to be resolved, this description should be a helpful starting point for a new GBIE algorithm that explicitly incorporates zonohedral gamuts and the relevant geometric constructions. Other researchers are encouraged to develop it further.

## 5 Summary

This paper has rigorously analyzed the geometric setting of Forsyth’s 1990 GBIE algorithm. A series of constructions (spectrum locus, spectrum cone, chromaticity diagram) has been presented, culminating in zonohedral illuminant gamuts. The geometric results have suggested improvements to algorithm implementations. The major positive improvement is a zonohedral method for calculating illuminant gamuts directly, rather than estimating them from training sets. Two negative improvements have been recommendations against treating the relationship between two illuminant gamuts as linear, and against using chromaticity containment tests as a replacement for *RGB* containment tests. A new GBIE algorithm has been suggested that avoids both these practices, uses zonohedral gamuts, and draws on the three-dimensional geometrical setting. While much further work remains, it is hoped that the geometric entities presented here will provide a firm foundation for further progress.

## References

1. Mark D. Fairchild. *Color Appearance Models*, 2<sup>nd</sup> ed., John Wiley & Sons, Ltd, 2005.
2. D. A. Forsyth, “A Novel Algorithm for Color Constancy,” *International Journal of Computer Vision*, 5:1, pp. 5-36, 1990.
3. Arjan Gijsenij, Theo Gevers, & Joost van de Weijer, “Computational Color Constancy: Survey and Experiments,” *IEEE Transactions on Image Processing*, Vol. 20, No. 9, September 2011.
4. G. D. Finlayson, “Color in Perspective,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 18, No. 10, pp. 1034-1038, October 1996.
5. Paul Centore, “Shadow Series in the Munsell System,” *Color Research & Application*, Vol. 38, No. 1, February 2013, pp. 58-64.
6. 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.
7. Paul Centore, “A Zonohedral Approach to Optimal Colours,” *Color Research & Application*, Vol. 38, No. 2, pp. 110-119, April 2013.
8. Paul Centore, “Geometric Invariants Under Illuminant Transformations,” *Color Research & Application*, Vol. 39, No. 2, pp. 179-187, April 2014.
9. R. V. Benson, *Euclidean Geometry and Convexity*, McGraw-Hill Book Company, New York, 1966.
10. Paul Heckbert, “An Efficient Algorithm for Generating Zonohedra,” 3-D Technical Memo 11, Computer Graphics Lab of the New York Institute of Technology, 24 February 1985.
11. Graham Finlayson & Steven Hordley, “Improving Gamut Mapping Color Constancy,” *IEEE Transactions on Image Processing*, Vol. 9, Issue 10, pp. 1774-1783, October 2000.
12. G. D. Finlayson, S. D. Hordley, & I. Tastl, “Gamut Constrained Illuminant Estimation,” *International Journal of Computer Vision*, Vol. 67, No. 1, pp. 93-109, 2006.