19

Physics-Based Approaches to Modeling Surface Color Perception

Laurence T. Maloney

The diversity of color results from the diversity of surfaces which absorb the light
– Ulrich von Strassburg (1262)

Introduction

Surface Color Perception. The study of surface color perception is a proper subset of the study of color perception, and one way to highlight the difference between them is to consider the effective stimulus appropriate to each. The effective stimulus for the study of color perception, broadly construed, is the spectral power distribution of light arriving at each point of the left and right retinas. There is no assumption that the patterns of light correspond to any possible arrangement of surfaces, objects and illuminants in a three-dimensional scene.

In contrast, the study of surface color perception presupposes that the light reaching the retinas has a history. The effective stimulus is the result of the interaction of certain light sources (*the illuminant*) with the surfaces of objects in an environment. It is clear that any stimulus appropriate for the study of surface color perception is also appropriate for the study of color perception but not *vice versa*.

Once we assume that the stimulus results from the interactions of lights and surfaces in an environment, it is natural to ask what information about the illuminants and surfaces in the environment is visually available to the observer. In particular, we can ask, to what extent does the color appearance assigned to a surface provide information concerning the physical, spectral properties of the surface? Visual systems whose estimates of surface color are determined by the spectral properties of surfaces exhibit a constancy, *color constancy.* Visual systems with color constancy have an objective capability: they remotely sense surface spectral information and represent it through color. This objective capability can be assessed in other species (Ingle, 1985; Werner, 1990; Neumeyer, 1981; Jacobs, 1990; See also Jacobs, 1981; 1993) as well as in humans.

Accordingly, while surface color perception is only a part of the study of color perception, it has long been recognized as an important part: "Colours have their greatest significance for us in so far as they are properties of bodies and can be used as marks of identification of bodies." (von Helmholtz, 1896/1962, Vol. II, p. 286).

Surface color perception is intimately linked with the precise physical, spectral properties of illuminants and surfaces in a scene. Indeed, without restrictions on illuminants, no degree of color constancy is possible: "The ... problem ... of constant color appearance is met by just one condition, namely restriction to one light source of constant spectral character " (Ives, 1912b, p. 70). In studying surface color perception, we cannot ignore the physical constraints that make it possible.

Last of all, scene layout in three dimensions can profoundly affect lightness perception (Gilchrist, 1977; 1980; Gilchrist, Delman & Jacobsen, 1983; See also Gilchrist, 1994). Consideration of surface color-perception leads us to consider possible connections between depth and shape perception and color perception.

Scope of the Review. This chapter is primarily concerned with recent algorithms for surface color perception based on explicit physical models of the environment. The first, and largest, part of the chapter describes these models and algorithms. Many of these algorithms are drawn from the computer vision literature. Even when there is no explicit claim by the authors that an algorithm could serve as a model of any aspect of human vision (or, more generally, biological vision), I have included it if it has implications for human vision.

The description of the algorithms is followed by a discussion of the relation between them and traditional models of color vision based on hypothesized color channels and the transformation of color information through successive stages (See: Hurvich, 1981; Wyszecki & Stiles, 1982a; Kaiser & Boynton, 1996). I briefly discuss models of human surface perception that are not based on explicit physical models such as Land and McCann's retinex theory (Land & McCann, 1971; Land, 1983; 1986) in *von Kries Algorithms*.

The algorithms presented differ primarily in how they obtain information about the illuminant in a given scene. The last part of the review proposes that the problem of *illuminant estimation* can be formally treated as a cue combination or fusion problem, analogous to cue combination in depth or shape vision (Landy et al., 1995).

Difficulties. The study of surface color perception is beset by three difficulties, each of which I will return to below. The first is methodological, and will be discussed in *Methodological Issues*.

The second is our current lack of knowledge concerning the physics of light and surface interactions in three-dimensional scenes. Research in the last 30 years has led to a better representation of light-surface interaction. As a consequence, a number of new approaches to recovering surface properties have arisen. The bulk of this review concerns these algorithms and the models of light-surface interaction that underly them. A central theme in this review is the link between recovery of surface color information and information concerning object shape and scene layout. It will become clear, however, that our understanding of the interactions of light and surface in real environments is far from complete.

The last difficulty, for lack of a better term, might be referred to as "conceptual clutter". The study of color constancy is beset by certain imprecisions in terminology. Over the course of the review, beginning with *Terminology*, I will attempt to clarify some of them.

Related Work. Previous reviews of the material discussed here include Hurlbert (1986; 1998), and Maloney (1992). Wandell (1995) is a useful introduction to both the empirical issues surrounding surface color perception and the necessary mathematical tools. Kaiser and Boynton (1996, 570ff) contains a partial review of some of the material described here. Healey, Shafer and Wolfe (1992) provides an overview of work in physics-based vision, and books by Hilbert (1987) and Thompson (1995) describe work in philosophy closely related to the material reviewed here. The two volumes edited by Byrne and Hilbert (1997a;b) provide an interesting introduction to both color science and philosophy.

Terminology

Color Constancy. The term 'color constancy' is employed in different ways in the literature. For some authors (Jameson & Hurvich, 1989; Kaiser & Boynton, 1996; Hurlbert, 1998), the term describes human perceptual *performance*: "the tendency to see colors as unchanging even under changing illumination conditions" (Hurlbert, 1998). The emphasis here is on whatever it is observers do achieve in any particular circumstance.

A second use of the term 'color constancy' is to treat it as a synonym for discounting changes in illumination of a scene (A sampling of authors: Beck, 1972; Arend, 1993; Foster & Nascimento, 1994; Bäuml, 1995). Brown and MacLeod (1997) have correctly criticized this use of the term 'color constancy' as neglecting other factors that might affect surface color perception: the presence of other surfaces, atmospheric haze, etc. We might reasonably refer to stability of color appearance despite changed in the illuminant as 'illuminant color constancy', defining as many new constancies as there are factors that can potentially influence perceived color: 'haze color constancy', etc.

Nevertheless, it is awkward to define something by listing the many things it does not depend on. Consequently, I adopt the following definition: *an observer has (perfect) color constancy precisely when the color appearance assigned to a small surface patch by the visual system is completely determined by that surface's local spectral properties.*[1] It should be clear that, if the color appearance of a surface patch is determined by its surface spectral reflectance, then it is not affected by changes in the illuminant, surrounding patches, haze, etc.

To summarize: the term 'color constancy' will be used in this review to describe a *goal*, one that is not necessarily achieved by any observer. Of central interest is the degree to which, and the circumstances under which, a human observer approximates 'perfect color constancy' in his or her judgments of surface color.

Various authors (Brill, 1978; 1979; Craven & Foster, 1992; Foster & Nascimento, 1994; Foster et al., 1997). have considered alternative color invariances, strictly weaker than color constancy, notably *relational color constancy* (Craven & Foster, 1992; Foster & Nascimento, 1994; Foster et al., 1997). Many of the issues raised here with respect to color constancy could also be raised with respect to these alternative invariances.

1. By "local spectral properties" is meant the bidirectional reflectance density function at a point of the surface, defined later in the chapter.

The Environment. It is often said that human color vision is 'approximately color constant' (Hurvich, 1981, p. 199; Brainard, Brunt & Speigle, 1997). This claim is misleading, if not further qualified, for the degree of color constancy exhibited can be very slight: "If changes in illumination are sufficiently great, surface colors may become radically altered ... weakly or moderately selective illuminants with respect to wave length leave surface colors relatively unchanged ... but a highly selective illuminant may make two surfaces which appear different in daylight indistinguishable, and surfaces of the same daylight color widely different" (Helson & Judd, 1936). If one does not restrict the range of lights and surfaces used in an experiment, then, human color constancy can be close to non-existent.

What is usually meant by the claim that human color vision is approximately color constant is that there are circumstances, such as a range of of possible illuminants and surfaces, similar to those encountered in everyday life, where human color vision approximates perfect color constancy: "With moderate departures from daylight in the spectral distribution of energy in the illuminant, external objects are seen ... nearly in their natural daylight colors." (Judd, 1940). The term *environment* will be used throughout to specify a set of assumptions concerning possible illuminants, surfaces, spatial layouts, etc. The central issues of surface color perception include (1) the determination of the degree of color constancy exhibited by a visual system for any given environment, and (2) the determination of the environments under which human color vision exhibits a given degree of color constancy.

Intrinsic Colors. The estimates of surface color appearance produced by a visual system that is even approximately color constant must depend on certain physical spectral properties of the surface. These properties will be referred to as 'intrinsic colors' (Shepard, 1992) for convenience in describing the algorithms below. The term 'intrinsic colors' will typically refer to whatever it is we are trying to estimate by means of a particular algorithm.

Methodological Issues

Experimental Methods. There are several methods commonly used to operationalize and measure the degree of 'color constancy' exhibited by human observers: First, there is *simultaneous asymmetric color matching* (Arend & Reeves, 1986; Arend et al., 1991; Brainard, Brunt & Speigle, 1997) in which the observer sees two regions of a scene comprising colored patches that are illuminated with two different illuminants. The observer adjusts a colored patch in one region (under one illuminant) to 'be the same color as' a fixed test patch in the second. Brainard, Brunt and Speigle (1997) argue that, if we wish to study the performance to be expected of an observer in a natural setting, he should be free to move his eyes around in the stimulus display. Nevertheless, the use of two illuminants in a single scene raises questions concerning the observer's adaptational state if he is allowed freely to look back and forth from one region to the other.

Second, there is *successive asymmetric color matching* (Brainard & Wandell, 1991; 1992; Bäuml, 1995) in which the observer views a single scene under first one illuminant and then a second. The observer adjusts a colored patch under the second illuminant to "be the same color as" a test patch seen under the first. The adjustable patch is typically in the same location as the test patch. Obviously, the observer's match depends on the observer's ability to remember colors accurately.

There is a potential confound with this method if the scene remains unchanged while the illuminant changes and the observer is aware that it does. Suppose, for example, that the observer noted that the fixed test patch is identical in appearance to a specific patch π under the first illuminant. He could then set the adjustable patch to match π in the second scene. This strategy could result in a good approximation to color constancy for visual systems that, in reality, have none. The same confound is present in simultaneous asymmetric color matching if the two scenes presented are the same, and the observer knows that they are. In general, the scene containing the test patch and the scene containing the adjustable patch should be different.

Third, there is *achromatic matching* (Helson & Michels, 1948; Werner & Walraven, 1982; Fairchild & Lennie, 1992; Arend, 1993; Bäuml, 1994; Brainard, 1998) where an observer adjusts a specified patch to appear achromatic. This method provides less information concerning the remapping of colors induced by changes in illumination than the previous two methods. We know only that the observer's setting corresponds to some point in his *achromatic locus*, but not which point. Nevertheless, the task is apparently very easy to explain to naive observers and very easy to carry out.

Andres and Mausfeld (described in Mausfeld, 1998) require observers to set a test patch to be in an alternative *color locus*, comprising those colors that are neither 'reddish' nor 'greenish'. It is plausible that certain colors or color loci such as the achromatic are easier to remember, or to communicate to observers, and that color appearance measures based on these loci will be more stable.

Task and Instruction Dependence. Arend and Reeves (1986; Arend, 1993) report that observers are capable of reliably performing two different tasks in asymmetric color matching of surfaces under reference illuminants and can be instructed to perform either. They may equate the chromaticity of the *light* (*color signal* in the terminology developed below) radiating from the two patches, or they may choose the setting consistent with the same surface viewed under two different illuminants. Troost and De Weert (1991) report that the effect of illuminant changes depends on the task the observer undertakes. Such task dependence and sensitivity to instructions only complicate the interpretation of experiments on surface color perception. Speigle and Brainard (1996, p. 171), however, consider three tasks involving surface color judgment (simultaneous asymmetric matches, achromatic matches, color naming). They report that "... all three tasks reveal similar and perhaps identical affects of the illumination ... " . Speigle (1998) also examined performance in color naming and color scaling tasks and reaches similar conclusions. While this issue is far from resolved, we may tentatively conclude that the same 'color percept' mediates performance in many tasks.

Simulations and Reality. In recent work, the stimulus arrangement ('scene') is most often simulated on a CRT display. The settings of the CRT's guns are chosen so as to produce the precise stimulation of the observer's photoreceptors that would result from a particular scene composed of specified physical surfaces under specified illuminants.

Researchers working with physical surfaces and lights in real scenes (Berns & Gorzynski, 1991; Brainard, 1998) typically report greater degrees of color constancy than do researchers working with simulated scenes displayed on CRTs. Most recently, Brainard, Rutherford and Kraft (1997) reported that they were unable to reproduce on a CRT display the degree of color constancy exhibited by observers in a real scene despite every attempt to make the simulated and real scenes identical. The degree of color constancy exhibited depends on the realism of the simulation in ways we do not yet understand, suggesting that there are cues present in real scenes that we also do not understand.

We know that it possible to choose viewing conditions (an 'environment') where human observers exhibit little or no color constancy (Helson & Judd, 1936). However, it is difficult, given the methodological problems outlined above, to draw any firm conclusions about the upper limit of performance in realistic environments. The results of Brainard and colleagues suggest that the *upper limit* to color constancy performance is quite high. We do not yet know how high, or what sorts of 'environments' produce optimal color constancy performance.

Environment I: Flat World

In this section I describe a model ('an environment') for surface color perception that abstracts away the three-dimensional layout of surfaces in a normal scene. The observer, in effect, views a scene painted on a planar surface, or perhaps the inside of a large sphere centered on him or her. The scene is illuminated uniformly by a single illuminant. There is no inter-reflection ('mutual illumination') among surfaces nor any specularity. I'll refer to the collection of assumptions that make up this environment as *Flat World*. It is an idealization of typical experimental arrangements that have been used to measure human surface color perception, and it differs in many respects from the 'realistic scenes' discussed in the previous section. Its importance stems from both the close relation between Flat World and previous experimental work and also its importance to many of the models of color-constant color perception or models of observed human performance that we will review.

Illuminant. Light from a single, distant, punctate light source (the *illuminant*) is absorbed by surfaces within a scene and re-emitted. $E(\lambda)$ will be used to denote the spectral power distribution of the incident illuminant at each wavelength λ in the electromagnetic spectrum. The re-emitted light that reaches the observer will be referred to as the *color signal*. Its spectral power distribution is denoted $L(\lambda)$. It contains the information about illuminant and surface at each point in the scene that is available to the observer.

The color signals reaching the observer (Fig. 19.1) are imaged onto a two-dimensional *sensor array*. The sensor array can be thought of as simplified model of a retina. We assign coordinates xy to each point in the array. The exact choice of a coordinate frame is not very important so long as each point in the sensor array has a unique coordinate xy. The color signal arriving at point xy on the sensor array is then denoted $L^{xy}(\lambda)$.

Surface. Consider now a small patch of the surface plane shown in Fig. 19.1 that is imaged at location xy in the sensor array. Part of the light that is absorbed by the surface patch is re-emitted and radiates in various directions. We characterize the effect of the surface patch on the resulting color signal by defining its *surface spectral reflectance* which is denoted $S^{xy}(\lambda)$:

$$L^{xy}(\lambda) = E(\lambda)S^{xy}(\lambda) \quad (1)$$

The above equation is assumed to hold for all λ.

In environments more complex than Flat World, the function $S(\lambda)$ for a surface patch depends on the *viewing geometry*: the location in three dimensions of the surface patch, the locations of other surfaces, the location of the observer, and that of the illuminant. We will return to this point below when we consider a second environment, *Shape World*.

Photoreceptor classes. The sensor array contains multiple classes of sensors (photoreceptors). Let $R_K(\lambda)$, $k = 1, 2, \ldots, P$ denote the spectral sensitivity of P distinct classes of photoreceptors. For a trichromatic human observer, P is taken to be 3.

We assume that the initial information available to a color vision system at a single retinal location is the excitation of each of the P classes of receptor,

$$(2)\quad \rho_1^{xy} = \int L^{xy}(\lambda)R_1(\lambda)d\lambda = \int E(\lambda)S^{xy}(\lambda)R_1(\lambda)d\lambda$$

$$\rho_2^{xy} = \int L^{xy}(\lambda)R_2(\lambda)d\lambda = \int E(\lambda)S^{xy}(\lambda)R_2(\lambda)d\lambda$$

$$\ldots$$

$$\rho_P^{xy} = \int L^{xy}(\lambda)R_P(\lambda)d\lambda = \int E(\lambda)S^{xy}(\lambda)R_P(\lambda)d\lambda .$$

The P numbers at each location form a vector $\rho^{xy} = \left[\rho_1^{xy}, \ldots, \rho_P^{xy}\right]$. A glance at Eqn 2 discloses that the entries of the vector ρ^{xy} at each location xy depend on both the light $E(\lambda)$ and the surface reflectance $S^{xy}(\lambda)$.

The RGB Heuristic. There is a persistent belief that crops up from time to time in the study of surface color perception that might be called the 'RGB Heuristic'. Define the 'color of the illuminant' $E(\lambda)$ to be the vector ρ^E whose components, $k = 1, 2, \ldots, P$, are given by,

$$(3)\qquad \rho_k^E = \int E(\lambda)R_k(\lambda)d\lambda .$$

These would be the excitations of the photoreceptors while looking directly at the illuminant $E(\lambda)$, or when viewing a perfectly reflecting surface, $S(\lambda) = 1$, under $E(\lambda)$. Define the 'color of the surface' $S(\lambda)$ to be ρ^S with components,

$$(4)\qquad \rho_k^S = \int S(\lambda)R_k(\lambda)d\lambda .$$

These photoreceptor excitations correspond to those for the surface $S(\lambda)$ viewed under a light with constant, unit spectral power density.

Now, to what extent do the 'color of the surface' ρ^S and the 'color of the illuminant' ρ^E determine the actual excitations corresponding to the surface viewed under the illuminant (denoted ρ^{ES})? It might seem plausible that

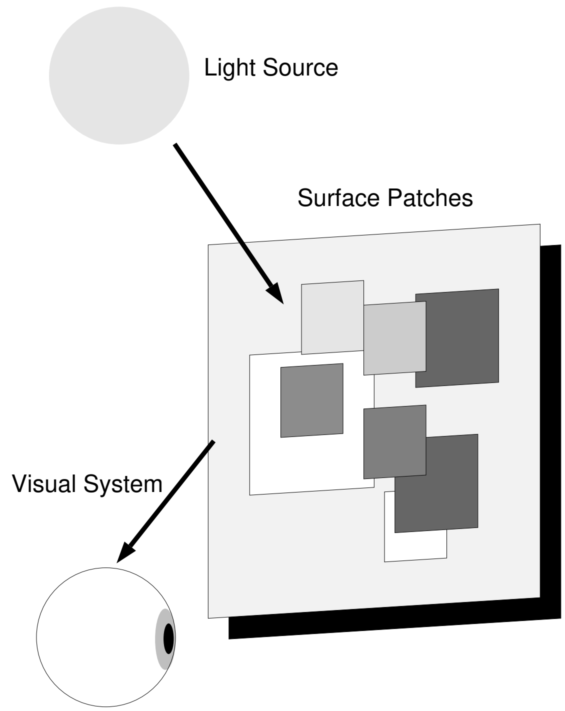

Figure 19.1: Flat World. A model environment for surface color perception that abstracts away the three-dimensional layout of surfaces in a normal scene. The observer views a scene painted on a planar surface. The scene is illuminated uniformly by a single, distant, punctate illuminant. There are no shadows, no inter-reflection ('mutual illumination') among surfaces nor any specularity.

$$(5) \qquad \rho_k^{ES} \approx \rho_k^E \rho_k^S$$

and this sort of 'approximation' is widely used in computer graphics for modeling the interaction of light and surface. However, there is no *mathematical* reason to expect Eqn 5 to hold, even approximately.

Evans (1948) gives a delightful counterexample in the form of two color plates showing objects illuminated under two lights of identical color appearance (ρ^E) but decidedly different spectral power distributions. Ives (1912b, p. 70) noted that an artificial daylight that matches daylight in color appearance need not render colors correctly: "A white surface under this [light] would look as it does under daylight but hardly a single other color would."

The accuracy (or inaccuracy) of the approximation in Eqn 5 depends on the possible surfaces and illuminants present in the environment. It is natural, then, to ask what kinds of constraints on lights and surfaces are needed in order for the approximation of Eqn 5 to achieve a specified degree of accuracy, and also to ask, whether physical surfaces in our environment satisfy those constraints. We will continue this discussion in *von Kries Algorithms* below.

Linear Models: Mathematical Notation

The Geometry of L^2. The functions $E(\lambda)$, $S^{xy}(\lambda)$, and $R_k(\lambda)$, introduced above, are assumed to be square-integrable functions, defined next. This assumption imposes no empirically significant constraint on the possible illuminants, surfaces, and receptors present in a scene.

The space of square-integrable functions L^2 is defined as follows. The magnitude or norm of a function $f(\lambda)$, is defined to be,

$$(6) \qquad \|f\| = \left[\int_{-\infty}^{\infty} f(\lambda)^2 d\lambda\right]^{1/2} .$$

The set of square-integrable real functions are those real functions which have finite magnitude, that is, the integral in Eqn 6 converges to a finite limit. The set of square-integrable functions form a vector space L^2, which has an inner product. The inner product of two square-integrable functions $f(\lambda)$, $g(\lambda)$ is defined to be,

$$(7) \qquad \langle f, g \rangle = \int_{-\infty}^{\infty} f(\lambda) g(\lambda) d\lambda \quad .$$

The two functions are orthogonal precisely when their inner product is 0.

L^2 is an example of a linear function space, a vector space whose elements are functions. Apostol (1969, Chaps. 1 and 2) is an elementary introduction to linear function subspaces. Strang (1988) is a standard introduction to finite-dimensional linear algebra. See Maddox (1970) and Young (1988) for more advanced treatments. Wandell (1995) contains an introduction to the use of linear function subspaces in vision.

Basis Functions. It is possible to show (Young, 1988) that we can find square-integrable functions $E_i(\lambda), i = 1, 2, 3, \ldots$, such that for any light $E(\lambda)$, there are unique real numbers ε_i such that

$$(8) \qquad E(\lambda) = \sum_{i=1}^{\infty} \varepsilon_i E_i(\lambda) \quad .$$

The functions $E_i(\lambda)$ form a *basis* of the linear function space L^2. Just as in the finite dimensional case, there are infinitely many possible choices of a basis for L^2. Our choice of basis for the lights $E(\lambda)$ will be guided by empirical considerations described below.

Coordinates. The real numbers ε_i are the *coordinates* corresponding to the light $E(\lambda)$ in the space L^2. The infinite vector $\varepsilon = [\varepsilon_1, \ldots, \varepsilon_i, \ldots]$ determines $E(\lambda)$. Not every function in L^2 corresponds to a physically possible light $E(\lambda)$; some represent spectral power densities with negative power at some point of the visible spectrum. The 'physically-realizable' lights form a convex set in L^2 analogous to the region within the spectral locus in CIE coordinates (Wyszecki & Stiles, 1982a).

The basis $E_i(\lambda), i = 1, 2, 3, \ldots$ could also serve to express the coordinates of the surface reflectances. I will instead choose a second basis $S_j(\lambda)$, $i =$

1, 2, 3, ... to express the coordinates of the surface reflectances $S^{xy}(\lambda)$:

$$(9) \qquad S^{xy}(\lambda) = \sum_{j=1}^{\infty} \sigma_j^{xy} S_j(\lambda) \quad .$$

The infinite vector $\sigma = [\sigma_1^{xy}, \ldots, \sigma_j^{xy}, \ldots]$ determines $S^{xy}(\lambda)$. The choice of this second basis will also be guided by empirical considerations. Note that while the coordinates σ_j^{xy} of $S^{xy}(\lambda)$ vary with location, the fixed basis elements do not. The same basis is used to model surface reflectance at every location in the scene.

History. Yilmaz (1962) first used truncated Fourier series expansions to model illuminants and surface reflectances across the visible spectrum, and Sällström (1973) first framed the problem in terms of truncated expansions using arbitrary bases. Brill (1978; 1979), Buchsbaum (1980), and Maloney and Wandell (1986; Maloney, 1984) independently developed this same representation of illuminants and surface spectral reflectance, and their interactions. Maloney and Wandell (1986) termed these constraints on light and surface, *linear models*.

Light-Surface Interaction. Substituting Eqs 8 and 9 into Eqn 2, we get,

$$(10) \qquad \rho_k^{xy} = \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} \varepsilon_i \sigma_j^{xy} \int E_i(\lambda) S_j(\lambda) R_k(\lambda) d\lambda \; .$$

The integrals $\int E_i(\lambda) S_j(\lambda) R_k(\lambda) d\lambda$ contain only fixed elements independent of the particular scene viewed. Setting $\gamma_{ijk} = \int E_i(\lambda) S_j(\lambda) R_k(\lambda) d\lambda$, Eqn 10 becomes,

$$(11) \qquad \rho_k^{xy} = \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} \varepsilon_i \sigma_j^{xy} \gamma_{ijk} \quad .$$

The values γ_{ijk} are fixed, known, and independent of any particular scene, once the bases have been selected. It should be clear that the specific illuminant enters into the visual process only through its coordinates ε_i, i=1, 2, ..., and the specific spectral reflectance functions only through their coordinates at each location, σ_j^{xy}, j=1, 2, ... The equation above is exact within the framework of assumptions adopted so far. Eqn 11 merely restates Eqn 2 with respect to two infinite-dimensional coordinate systems; it can no more be solved for information about the surface reflectance (about the σ^{xy}) independent of the light than could Eqn 2. Any of the coordinates, σ_j^{xy} could serve as an intrinsic color (see Terminology above) – if we could reliably estimate it despite changes in the illuminant. But, in the equation above, information about light and surface is irreversibly tangled.

We next approximate the infinite summations above (that perfectly capture light and surface reflectance) by truncated, finite summations:

$$(12) \qquad E_\varepsilon^{xy}(\lambda) = \sum_{i=1}^{M} \varepsilon_i E_i(\lambda)$$

$$(13) \qquad S_\sigma^{xy}(\lambda) = \sum_{j=1}^{N} \sigma_j^{xy} S_j(\lambda)$$

The class of lights that can be represented in this way for a fixed value of M, and fixed basis elements $E_1(\lambda)$, $E_2(\lambda)$, ... $E_M(\lambda)$ is a finite-dimensional linear function subspace ('linear model') of $\boldsymbol{L}^2$ that has dimension M. A linear function subspace of surface reflectances ('linear model') with dimension N is defined analogously.

The *finite*-dimensional vectors $\varepsilon=[\varepsilon_1, \ldots, \varepsilon_M]$ and $\sigma^{xy} = [\sigma_1^{xy}, \ldots, \sigma_N^{xy}]$ will be referred to as the 'coordinates' of light and surface within their respective linear subspaces. The subscripted variables $E_\varepsilon(\lambda)$ and $S_\sigma^{xy}(\lambda)$ will be used throughout to denote lights and surfaces constrained to lie in finite-dimensional linear models.

Linear Models: Fits To Empirical Data

Surfaces. The algorithms described in the next section depend crucially on the assumption that we can approximate real illuminants and surfaces by linear models with low values of M and N. In this section we review work concerning the realism of such models as descriptions of empirical surfaces and illuminants. An expanded version of this section may be found in Maloney (1998).

Fitting Methods. I describe here how to fit an optimal least-squares linear model to the set of empirical surface reflectances described in Vrhel, Gershon and Iwan (1994). More sophisticated methods permit simultaneous choice of optimal linear models for any sets of empirical illuminants and surfaces (Vrhel & Trussel, 1992; Marimont & Wandell, 1992). I will refer to the data set of Vrhel, Gershon and Iwan (1994) as the *Kodak Data.*[2]

Suppose that we have a set of empirically-measured surface spectral reflectances, $S^{v}(\lambda)$, $\nu = 1, 2, \ldots, V$,. We will treat a sampled surface reflectance function, sampled at wavelengths λ_i, as a step function that is constant between λ_i and λ_{i+1} , and has as its constant value, the sampled value at λ_i. These step functions have values defined at all λ and their use allows me to use the same notation and conventions for empirical and theoretical surface spectral reflectance functions.

For any fixed value of N, and any choice of basis functions $S_1(\lambda), \ldots, S_N(\lambda)$ we can compute, by linear regression (Maloney, 1986), the weights $\hat{\sigma}_j^v$ in the truncated series of Eqn 13, that minimize the least square error,

$$(14) \qquad \Xi_v = \left\| S_v - S_{\hat{\sigma}}^v \right\|^2 = \int (S_v(\lambda) - S_{\hat{\sigma}}^v(\lambda))^2 d\lambda$$

.

We now wish to choose the *basis functions* $S_j(\lambda), j = 1, 2, \ldots, N$ to minimize the overall error,

$$(15) \qquad \Xi = \Xi^1 + \ldots + \Xi^V \quad .$$

The optimal basis functions will be each orthogonal to all of the others. It is convenient to assume that they are scaled to be of unit magnitude. Several authors fit empirical data using an alternative linear model to that of Eqn 13:

$$(16) \qquad \hat{S}^{v}(\lambda) = \bar{S}(\lambda) + \sum_{j=1}^{N} \sigma_j^v S_j(\lambda)$$

where $\bar{S}(\lambda)$ is the mean of the $S^v(\lambda)$. This model is presupposed by *principle component analysis* (Mardia, Kent & Bibby, 1979) and authors who report using principle components analysis to fit their data will likely assume the model of Eqn 16, not that of Eqn 13.

The models are different. Note, in particular, that the mean in Eqn 16 will not, in general, be orthogonal to the remaining basis elements or independent of them. It is even possible that $\bar{S}(\lambda)$ be identical to $S_1(\lambda)$, an undesirable outcome. Further, with Eqn 13, the scaled copies of any model surface reflectance $aS_{\sigma}(\lambda)$ are automatically in the space of model surface reflectances, a scaling property that does not, in general, hold for Eqn 16. Many authors (e.g. Judd, MacAdam & Wyszecki, 1964) normalize the vectors in their data set before fitting them. They may normalize the data vectors to have $\boldsymbol{L}^2$ norm 1, or to have value 1 at a specified wavelength. They then fit the normalized data, using principle components, and, in reconstructing the original data, scale the mean by a value σ_0,

$$(17) \qquad \hat{S}^v(\lambda) = \sigma_0 \bar{S}(\lambda) + \sum_{j=1}^{N} \sigma_j^v S_j(\lambda) \quad .$$

which is simply the inverse of the constant by which the original measured surface reflectance was scaled in normalizing it. Eqn 17 and Eqn 13 appear to be identical (except for a change in the index), but, again, the vectors in Eqn 17 need not be orthogonal. Accordingly, I use the models of Eqs 12 and 13 in the sequel.

For empirically measured sets of surface reflectance functions, the optimal basis consistent with Eqs 12 or 13 can be readily computed by standard linear algebra methods using the singular value decomposition (Mardia, Kent & Bibby, 1979). Once we have computed the optimal basis $S_1(\lambda), \ldots, S_N(\lambda)$, we

2. The Kodak Data and other sets of surface spectral reflectance functions collected by various authors are available from ftp.cns.nyu.edu:pub/ltm/SSR via anonymous ftp. There are currently no large set sof illuminants publically available.

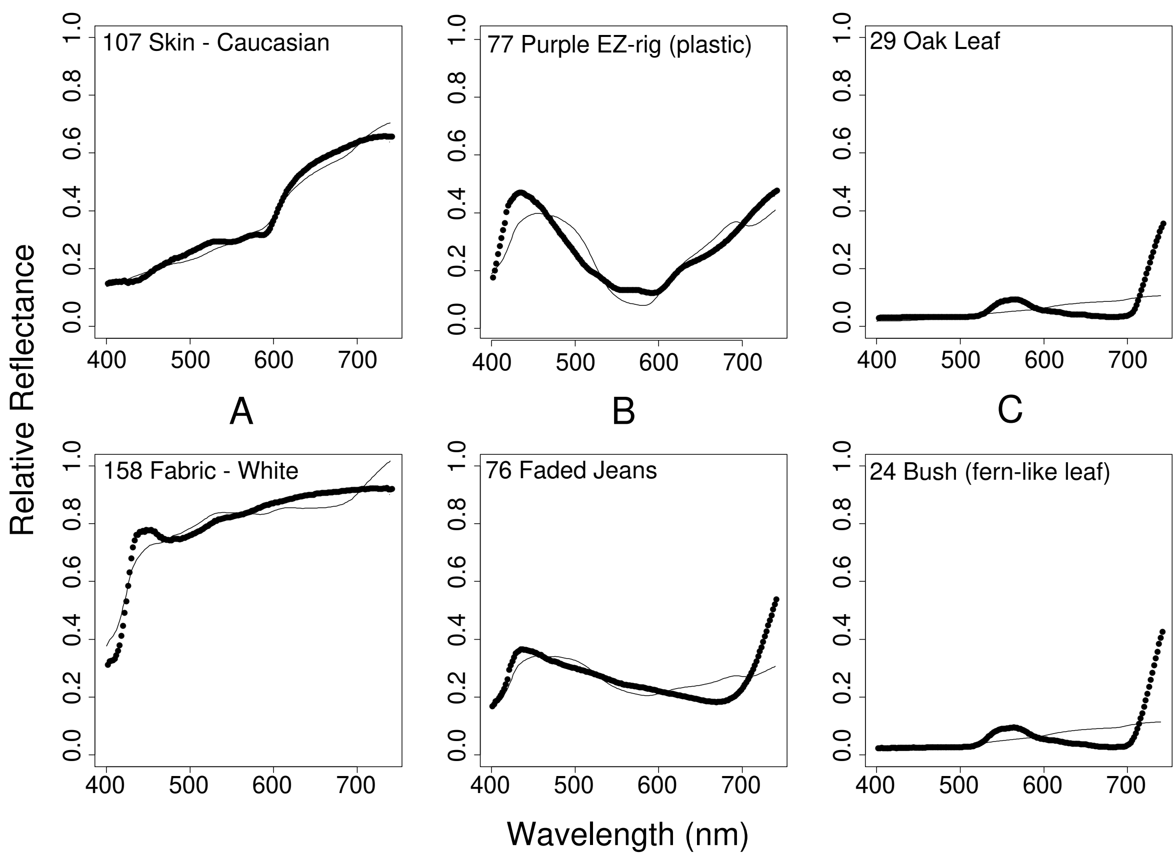

Figure 19.2: Surface spectral reflectance data from Vrhel, Gershon and Iwan (1994) (plotted with filled circles) and reconstructions with three basis elements. The horizontal axes are wavelength, in nanometers, the vertical, relative reflectance. A: The two fits at the first quartile of Ξ_ν. B: The two fits at the third quartile. C: The worst two fits. See text for details.

can compute the optimal approximation $\hat{S}^\nu(\lambda)$ by linear regression for each of the empirical functions $S^\nu(\lambda)$ and compare it to the original.

The Kodak Data Set. Fig. 19.2 contains plots of six of the surfaces in a collection of 170 surfaces measured by Vrhel, Gershon and Iwan (1994) and approximations[3] to those surfaces with $N = 3$. The two plots in Fig. 19.2A correspond to the two surfaces whose Ξ^ν's fell at the first quartile of the 170 values of Ξ^ν. That is, about 25% of the surfaces in the sample have smaller (better) values of Ξ^ν. The two plots in Fig. 19.2B correspond to the two surfaces whose Ξ^ν's fell at the third quartile of the 170 values of Ξ^ν. About 75% of the surfaces in the sample were better fit. The two plots in Fig. 19.2C correspond to the two surfaces with the largest values of Ξ^ν ('the worst fits').

Fig. 19.3 contains analogous plots for $N = 8$. These results illustrate the conclusion of Vrhel, Gershon and Iwan (1994), that linear models with $N = 3$ provide poor approximations to the measured surface reflectances in their data set while, with $N = 8$, the approximations are good.[4]

Fig. 19.4 shows the 'variance accounted for' (the Ξ

3. Vrhel, Gershon and Iwan (1994) used principle component analysis to analyze this data. I refit their data by the least-square model for reasons described in the main text and so that it could be more readily compared to earlier work.

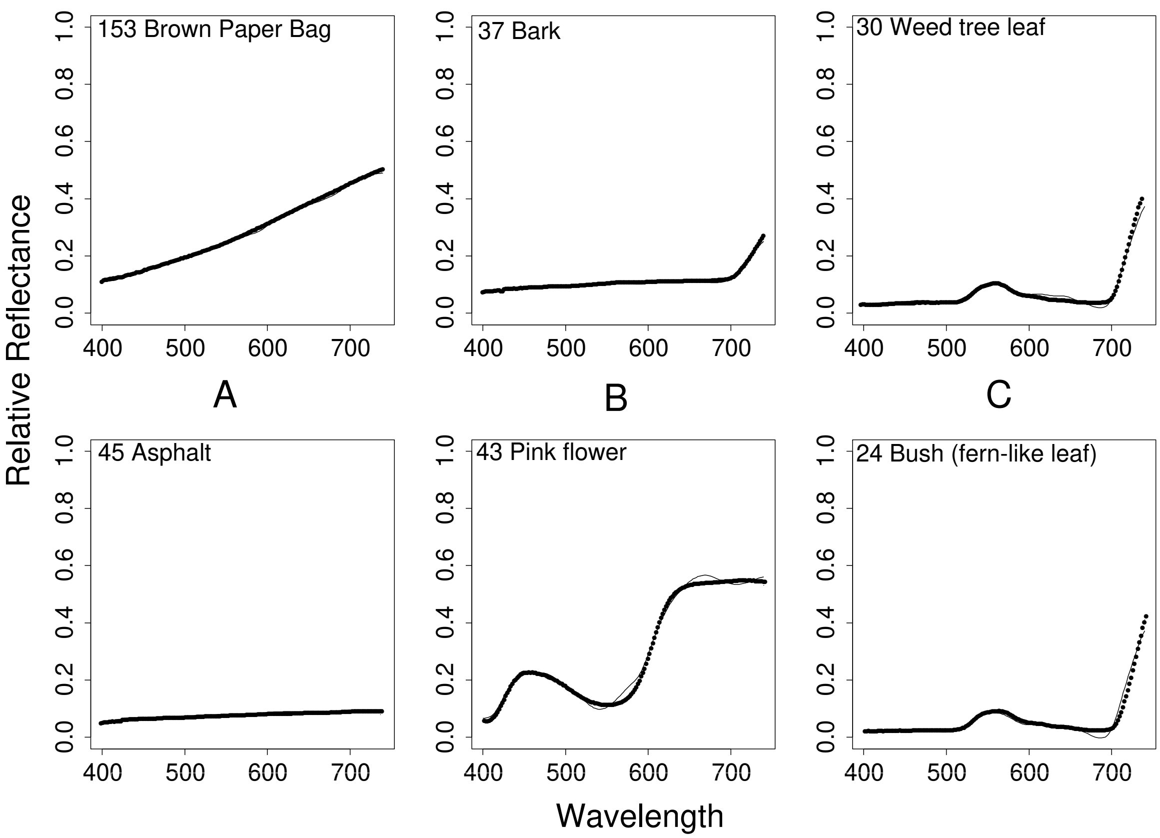

Figure 19.3: Surface spectral reflectance data from Vrhel, Gershon and Iwan (1994) (plotted with filled circles) and reconstructions with eight basis elements. The horizontal axes are wavelength, in nanometers, the vertical, relative reflectance. A: The two fits at the first quartile of Ξ_v. B: The two fits at the third quartile. C: The worst two fits. See text for details.

term normalized) as a function of N for the data of Vrhel, Gershon and Iwan (1994). The results of Maloney (1986) for the Krinov (1947) data set are also plotted in Fig. 19.4. It is clear that the fits to the Krinov data set reported by Maloney understate the difficulty of modeling empirical surface reflectances with low-dimensional linear models (if we take the Vrhel, Gershon & Iwan data as representative of empirical surface reflectances). The results, however, do confirm the conclusions drawn in Maloney (1986): "... the number of parameters required to model ... spectral reflectances is five to seven, not three" (Vrhel, Gershon & Iwan, 1994, p. 1674). The failure to find highly accurate approximations to real surfaces with $N = 3$ will have implications for the models and algorithms reviewed below which we will return to in *Model Failure and Approximate Color Constancy.*

History. Cohen (1964) used principle components analysis (Mardia, Kent & Bibby, 1979) to fit the surface spectral reflectances of a subset of the Munsell color chips. He concluded that the mean surface reflectance and as few as two additional components pro-

4. Vrhel, Gershon and Iwan (1994) report that roughly seven principle components suffice. The choice of 7 or 8 or 9 is somewhat arbitrary. See Fig. 19.4.

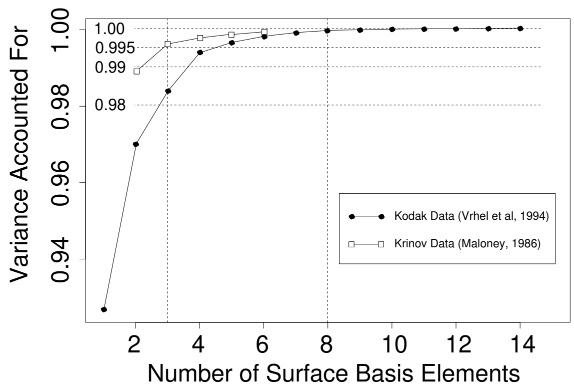

Figure 19.4: Variance accounted for versus number of basis elements for fits of the Vrhel, Gershon and Iwan (1994) data (filled circles) and for fits of the Krinov data from (Maloney, 1986) plotted as open squares. The dashed horizontal and vertical lines mark cutoffs of interest.

vided good fits to the Munsell surface reflectances. Maloney (1986) refit the same data set using Eqn 13 and the criterion in Eqs 14 and 15, and measured how well the 'Munsell basis' fit a a large set of surface reflectance functions collected by Krinov (1947/1953). The results of this work were described above. Parkkinen, Hallikainen and Jaaskelainen (1989) measured the surface spectral reflectance of a large collection of Munsell chips and fit them using principle components analysis. As many as 8 basis elements were needed to permit accurate fits to the surfaces.

Illuminants. The fitting methods and models for illuminants are identical to those for surface reflectance discussed above. Judd, MacAdam and Wyszecki (1964) report summary results for a large set of measured spectral power distributions of daylight (principle components fit). Their results and later, more extensive work by other researchers (Das & Sastri, 1965; Sastri & Das, 1966; 1968; Dixon, 1978) indicate that sampled daylight may be well described by a small number of basis elements (possibly as small as $M = 3$). Romero, Garciá-Beltrán and Hernández-Andrés (1997) sampled the spectral power distributions of daylight from 400 nm to 700 nm over a period of four days in Granada, Spain. They performed a principle components analysis on the resulting 99 spectral power distributions, each normalized to magnitude 1 (in $\boldsymbol{L}^2$). They found that approximations to the measured spectral power distributions using three basis elements accounted for 0.9997 of the variance. The number of samples collected was not large, the time period over which they were collected was short, and it is not clear how representative the climate of Granada is of climates in other regions of the world. Nevertheless, the fit is remarkable, and their results, together with the results of earlier research, indicate that low-dimensional linear models provide very good approximations to daylight spectral power distributions.

Issues in Fitting Empirical Data. The results of this section suggest that surface reflectance functions and illuminants in 'natural environments' are constrained. This idea is not new. Several authors (Stiles, Wyszecki & Ohta, 1977; Lythgoe, 1979; MacAdam, 1981) have expressed the opinion that empirical surface reflectances are smooth, constrained curves. Land (1959/1961) asserts that 'Pigments in our world have broad reflection characteristics.' Still, there are many open questions concerning the nature and importance of the constraints on 'natural' surfaces and light and how they might best be modeled. In the remainder of this section I raise some of them.

Non-Linear Models. Only linear models are considered here as candidate representations for natural surfaces and reflectances. It is certainly possible that a non-linear model with N parameters such as,

$$(18) \qquad S_{\sigma}^{xy}(\lambda) = \prod_{j=1}^{N} S_j^{\sigma_j^{xy}}(\lambda) \quad .$$

might provide better approximations to empirical surface reflectance functions than any linear model with N parameters. There seems to have been no systematic attempt to find non-linear models that provide better fits than linear models. Many of the algorithms below could be readily altered to take advantage of a nonlinear constraint such as that embodied in Eqn 18.

Loss Functions. The use of the least-square error measure Ξ in fitting models to data is questionable. An advantage of the least-square error measure is that it is independent of the properties of any particular visual system. Any conclusions drawn are statements about the empirically measured surface spectral reflectances themselves, and, in attempting to understand the physical bases for the empirically-observed constraints on surfaces, this is desirable (Maloney, 1986).

Yet it is also desirable to employ error measures that reflect the sensitivity of specific visual systems. Human vision, for example, is likely to be very insensitive to failures in approximation near the extremes of the (human) visible spectrum. Maloney (1986) refit the Krinov data with error weighted by human photopic visual sensitivity $V(\lambda)$ (See Wyszecki & Stiles, 1982a). He compared the variance accounted for by the weighted fit and the unweighted fit and concluded that the weighted fit accounted for markedly more of the variance. That is, the least-square approximations to the Krinov data were enhanced by the spectral properties of the human visual system. A better approach (Dannemiller, 1992) is to measure the ability of an ideal-observer to discriminate approximations from real surface functions. Marimont and Wandell (1992), developed fitting techniques that took into account human visual sensitivity, and derived basis functions that , as expected, approached zero at the ends of the visible spectrum.

It is important to consider both the nature of the physical constraint on surfaces (independent of any visual system) and also the impact of the constraint on visual performance for particular visual systems.

Theoretical Approaches. It is not clear what constitutes a 'natural environment' for human vision or how to sample it, what surfaces should be included, and what weight each should be given. When we consider other biological visual systems, the problem is scarcely less difficult. It is, therefore, desirable to consider why at least some classes of physical surfaces exhibit physical constraints and what these constraints might be. If we understood the theoretical bases for these constraints, we need not wonder whether they might vanish with the next collection of empirical surface reflectance functions.

Stiles, Wyszecki and Ohta (1977) and, later, Buchsbaum and Gottschalk (1984) suggested that many surface reflectance functions are 'low-pass': their surface spectral reflectance functions, as a function of wavelength, are approximately low-pass. Maloney (1984; 1986) tested this 'low-pass hypothesis' for the Krinov data and concluded that the Krinov reflectances contained little spectral energy above a band-limit corresponding to three samples. He suggested specific physical processes responsible for this observed 'low-pass' constraint for organic colorants.

Mollon, Estévez and Cavonius (1990) argue that the low-pass constraint observed in the Krinov data is simply an artifact of the measurement process employed by Krinov. Krinov measured not isolated, homogeneous surfaces, but natural formations, each comprising many distinct materials. The spatial mixing of distinct surface spectral reflectances, they argue, leads to measured surface spectral reflectances that have been effectively passed through a low-pass filter and, in contrast, "[i]f the measurement is confined to part of an individual leaf or individual fruit, then fine detail is readily apparent in the spectra of the world of plants.(p. 129)". This issue remains to be resolved. However, the first eight basis elements of the Vrhel, Gershon and Iwan (1994) data (shown in Fig. 19.5) exhibit the same band-pass form with increasing peak frequencies asfound by Maloney (1984; 1986) with the Krinov data. Further discussion of the low-pass hypothesis may be found in Maloney (1998).

Flat World Algorithms

This section describes algorithms that presuppose a Flat World environment with surfaces and illuminants perfectly described by low-dimensional linear models ($M = N = 3$).

We substitute $E_{\varepsilon}(\lambda)$ from Eqn 12 for $E(\lambda)$ in Eqn 2, and $S_{\sigma}(\lambda)$ from Eqn 13 for $S(\lambda)$ in the same Eqn 2. We then rewrite Eqn 2, expressing the basic relations among the light, surface reflectances, and receptor responses, as the matrix equation

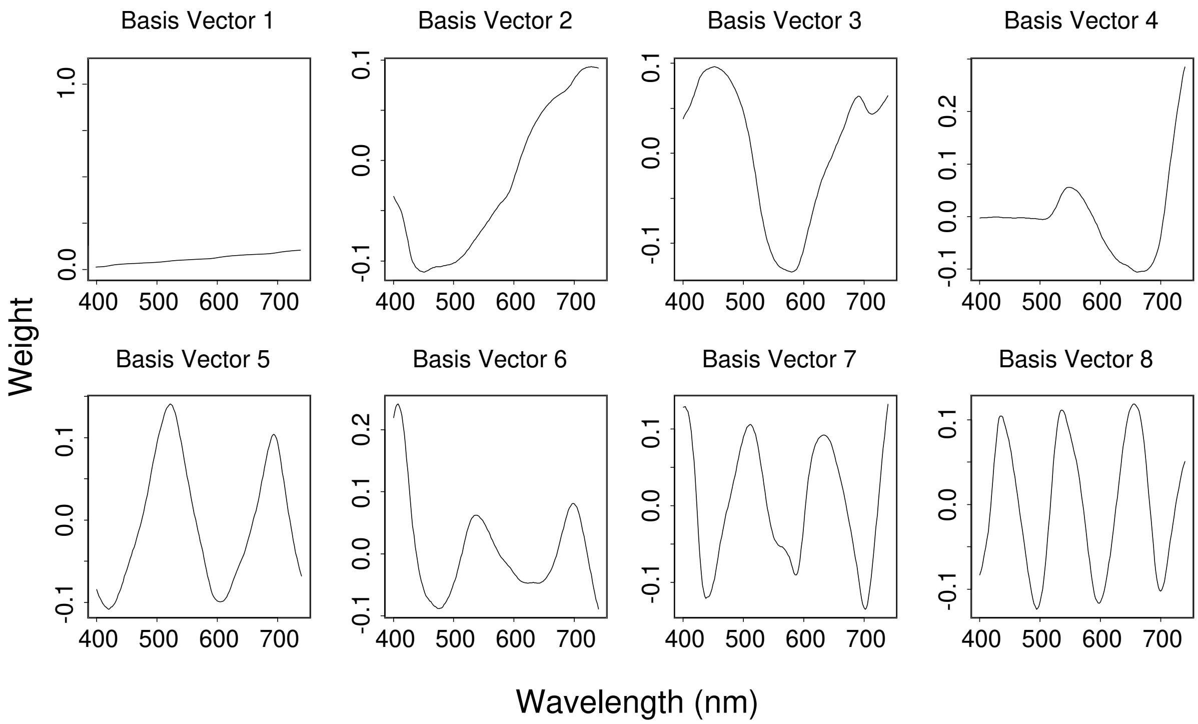

Figure 19.5: Figure 5: The first eight basis functions for the Vrhel, Gershon and Iwan (1994) data. The units of the horizontal axes are wavelength in nanometers, those of the the vertical axes, relative reflectance.

(19) $\rho^{xy} = \Lambda_{\varepsilon}\sigma^{xy}$

where $\rho^{xy} = [\rho_1^{xy}, \ldots, \rho_P^{xy}]$ is, as above, the vector formed from the receptor excitations of the P receptors at location xy. The matrix Λ_{ε} is P by N, and its kj'th entry is of the form $\int E_{\varepsilon}(\lambda)S_j(\lambda)R_k(\lambda)d\lambda$. Note that the matrix Λ_{ε} depends only on the light (as the basis elements S_j and receptor spectral sensitivities R_k are fixed, independent of any particular scene). This matrix captures the role of the light in transforming surface reflectances at each location xy into receptor excitations.

The coordinates $\sigma = [\sigma_1^{xy}, \ldots, \sigma_N^{xy}]$ (or any convenient transform of them) could serve as intrinsic colors, i.e., the quantities we seek to estimate given the ρ^{xy} at each location. Various limits on recovery of the σ^{xy} are dictated by Eqn 19. We consider the limits on recovery when the light on the scene is assumed to be known, and when the light on the scene is unknown.

In the simple case in which the ambient light, and (therefore) the lighting matrix Λ_{ε}, is known, we see that to recover the N weights that determine the surface reflectance we need merely solve a set of simultaneous linear equations. The recovery procedure reduces to matrix inversion when $P = N$. That is,

(20) $\Lambda_{\varepsilon}^{-1}\rho^{xy} = \sigma^{xy}$

where the quantities on the left hand side are all known or computable from known quantities. Recovery is also possible when $P > N$ whenever Λ_{ε} corresponds to an injective (1-1) linear transformation. If P is less than N, Eqn 19 is under-determined and there is no (unique) solution. (See Maloney (1984) for a discussion of the invertibility of the various matrices referenced above and in the sequel.)

If the ambient light is unknown then it is easy to show that we cannot do as well: we cannot, in general, recover the ambient light vector ε or the spectral reflectances even when $P = N$, so that matrix Λ_ε is square. Given any collection of photoreceptor excitations ρ^{xy} across a scene, there are, in general, many possible choices of surface reflectances σ^{xy} and illuminant ε that satisfy Eqn 19. Since any such choice of a light vector ε and corresponding surface reflectances σ^{xy} could have produced the observed receptor excitations, we cannot determine which, in fact, did. Even if we restrict attention to the convex subset of 'physically-realizable' lights, such confusions are still possible for many choices of lights and surfaces.

Each of the linear models algorithms, described next, takes a distinct approach to estimating the σ^{xy}.

The Reference Surface Algorithms of Brill and Buchsbaum. Brill (1978; 1979) considers the case where $P = N = 3$ and there are three reference surfaces available in the scene at known locations xy_1, xy_2 and xy_3. (Brill's algorithm requires no restriction on the illuminant: $M = \infty$.) Then the receptor excitations ρ^{xy} at these locations (among others) are known and we have the simultaneous matrix equations:

$$(21) \qquad \begin{aligned} \rho^{xy_1} &= \Lambda_\varepsilon \sigma^{xy_1} \\ \rho^{xy_2} &= \Lambda_\varepsilon \sigma^{xy_2} \\ \rho^{xy_3} &= \Lambda_\varepsilon \sigma^{xy_3} \end{aligned}$$

Both sides of each equation are known. If the reference surfaces $\sigma^{xy_l}, l = 1, 2, 3$ are linearly independent, then it is possible to solve for Λ_ε (if the σ^{xy_i} are taken as the basis of the space of surfaces, then the matrix Ω whose columns are $\rho^{xy_i}, i = 1, 2, 3$ is the desired matrix Λ_ε). If another basis is used, then Λ_ε is simply the inverse of the matrix Ω premultiplied by the matrix that changes from the σ^{xy_i} basis to the second basis (See Strang, 1988). Once Λ_ε is known, we invert Eqn 19 and solve for the coordinates (intrinsic colors) of all sources in the scene.

Buchsbaum (1980) assumes that $M = N = 3$, and requires that the location of one reference surface σ^{xy_0} be known. To explain his algorithm, we first define

$$(22) \qquad \begin{aligned} \Lambda_1 &= \Lambda_{[1,0,0]} \\ \Lambda_2 &= \Lambda_{[0,1,0]} \\ \Lambda_3 &= \Lambda_{[0,0,1]} \end{aligned},$$

the light matrix Λ_ε for each of the known basis lights. Then, the quantities

$$(23) \qquad \begin{aligned} \rho^{xy_1} &= \Lambda_1 \sigma^{xy_0} \\ \rho^{xy_2} &= \Lambda_2 \sigma^{xy_0} \\ \rho^{xy_3} &= \Lambda_3 \sigma^{xy_0} \end{aligned}$$

are all known once the reference surface σ^{xy_0} is given. These are the receptor excitations corresponding to the reference surface under each of the basis lights in turn. It can be shown that (See Maloney, 1984, Chap. 3),

$$(24) \qquad \Lambda_\varepsilon = \varepsilon_1 \Lambda_1 + \varepsilon_2 \Lambda_2 + \varepsilon_3 \Lambda_3,$$

giving the following expression for the receptor excitation of the reference surface under an unknown light ε:

$$(25) \qquad \rho^{xy_0} = \Lambda_\varepsilon \sigma^{xy_0} = \varepsilon_1 \rho^{xy_1} + \varepsilon_2 \rho^{xy_2} + \varepsilon_3 \rho^{xy_3}$$

.

If the fixed vectors $\rho^{xy_l}, l = 1, 2, 3$ are linearly independent, then the above equation can be solved for ε given ρ^{xy_0}. (The coordinates of the light ε are precisely the coordinates of the reference surfaces' receptor excitations ρ^{xy_0} with respect to the basis $\{\rho^{xy_1}, \rho^{xy_2}, \rho^{xy_3}\}$.) In summary, a single reference surface permits estimating the light when $P = N = M = 3$. Once ε is known, Eqn 19 may be solved for the intrinsic colors of surfaces.

The algorithm of Brill can be generalized to the case where $P = N$ takes on any arbitrary value; N

linearly independent reference surfaces are then required. Buchsbaum's algorithm can be generalized to arbitrary values $M = N = P$ and it still requires only a single reference surface. Both algorithms can be applied to a visual system with any number of types of receptors P.

Note that Brill's algorithm makes no assumptions about the dimensionality of the linear model describing the illuminant (M). It can be used to estimate the matrix Λ_ε in environments where no restrictions are placed upon the illuminant $E(\lambda)$. Brill's algorithm can, therefore, be used in environments where little is known about the illuminant. Of course, if the estimated Λ_ε is not invertible, then the surface parameters σ cannot be recovered by any method.

In Buchsbaum's algorithm, the three fixed receptor excitations $\rho^{xy_l}, l = 1, 2, 3$ generated by the single reference surface under the known basis lights serve much the same role as Brill's three reference surfaces: they 'pin down' the light matrix.

Buchsbaum's Gray World Algorithm. In either Brill's or Buchsbaum's algorithm, the receptor excitation corresponding to one reference surface may be replaced by the average of the receptor excitations catches ρ^{xy} across a portion of the scene W, provided we know the true mean of the intrinsic colors σ^{xy} across that part of the scene.

This substitution is an obvious consequence of the linearity of Eqn 19: if, for all $xy \in W$,

$$\rho^{xy} = \Lambda_\varepsilon \sigma^{xy} \tag{26}$$

then

$$\sum_{xy \in W} \rho^{xy} = \Lambda_\varepsilon \sum_{xy \in W} \sigma^{xy} \tag{27}$$

and, dividing both sides by the number of locations in the set W,

$$\rho^{mean} = \Lambda_\varepsilon \sigma^{mean} \tag{28}$$

In particular we can use the mean across the entire scene. If, for example, we assume that the mean intrinsic color of a scene is a specific 'gray', then this 'gray' σ^{xy_0}, paired with, ρ^{xy_0}, the mean of the observed receptor excitations the scene, permit estimation of the illuminant using Buchbaum's algorithm. The assumption concerning the mean of the intrinsic colors in a scene is sometimes termed the *Gray World* assumption (Buchsbaum, 1980, p. 24): "It seems that arbitrary natural everyday scenes composed of dozens of colour subfields, usually none highly saturated, will have a certain, almost fixed spatial reflectance average. It is reasonable that this average will be some medium gray..." D'Zmura and Lennie (1986, p. 1667) make a similar claim "... we expect that the space-averaged light from most natural scenes will bear a chromaticity that closely approximates that of the illuminant." The 'Gray World' assumption is close in spirit to Helson's 'adaptation level' (Helson, 1934).

The term 'Gray World Assumption' is misleading in two senses. First, for Buchsbaum's algorithm to work, the known average of the intrinsic surfaces need only be *known*; it need not be 'gray'. If the average of the intrinsic colors across the scene corresponds to a 'green' surface, the algorithm of Buchsbaum can as easily make use of a 'Green World Assumption'. Second, the average need not include the entire scene: any portion of the scene (e.g. the ground plane) can serve as an average reference. From this point on, I will refer to the Gray World Assumption in this wider sense as the 'Stable Mean Assumption'.

The 'Stable Mean Assumption' is, fundamentally, a claim about the physical environment as seen through a particular set of photoreceptors, and such a claim is testable. It is obvious that a Stable Mean algorithm such as Buchsbaum's will erroneously 'correct' any deviation of the mean intrinsic color away from the known reference color as readily as it corrects an imbalance induced by the illuminant. The algorithm is of value to the extent that typical excursions of the mean intrinsic color away from the reference are small and/or infrequent compared to the magnitude of changes in the illuminant. Empirical tests of these claims are lacking.

It should be noted that the Stable Mean Assumption is specific to a particular set of photoreceptor types. It is entirely possible that a Stable Mean Algorithm may fail dramatically for one species and succeed for another in the same environment. Discussion of the status of the 'Stable Mean Assumption' in human vision will be postponed until *von Kries Algorithms*.

The use of reference surfaces in two of the algorithms above make them implausible candidates for a model of human color vision. The remaining algorithms described in this section illustrates ways of dispensing with reference surfaces altogether. Some of the algorithms first estimate the coordinates of the illuminant ε and then employ Eqn 20 to solve for intrinsic colors. Others estimate ε and the σ^{xy} simultaneously. Only Brill's algorithm avoids estimating ε and instead estimates Λ_ε directly. The remaining algorithms differ mainly in how they go about determining the illuminant.

The Subspace Algorithm of Maloney and Wandell. Maloney and Wandell (1986) assumed that there are more classes of receptors than dimensions in the linear model of surface reflectances: $P > N$. Suppose that there are $P = N + 1$ linearly independent receptors available to spectrally sample the image at each location. They proposed a method for computing the light coordinates ε and the N dimensional surface reflectance vectors σ^{xy} given the $N + 1$ dimensional receptor response vector ρ^{xy} at each location. The matrix Λ_ε is then a linear transformation from the N-dimensional space of surface reflectances σ^{xy} into the $N + 1$ -dimensional space of receptor excitations ρ^{xy}. As Λ_ε is a linear transformation, the receptor responses must fall in a proper subspace of the receptor space. In the case $P = 3$, $N = 2$, the vectors ρ^{xy} must lie on a plane in the three-dimensional receptor space In the case $P = 4$, $N = 3$, the vectors ρ^{xy} must lie in a three-dimensional subspace (a '3-space') of the four-dimensional receptor space. The particular subspace is determined by Λ_ε and therefore by the lighting parameter ε. That is, as ε is varied for a fixed set of surfaces, the subspace spanned by ρ^{xy} moves about in the receptor space.

. Maloney and Wandell (1986) proposed a two-step procedure to estimate normalized light and surface reflectance properties. First, they determine the plane spanning the receptor excitations. Second, knowledge of the plane permits recovery of the normalized ambient light vector $\hat{\varepsilon}$ by computations described by Maloney (1984). These computations use the vector at the origin normal to the subspace spanned by the receptor receptor excitations. This vector serves a role analogous to the single reference surface in Buchbaum's algorithm: it is a recoverable 'landmark' that moves around as a function of the illuminant. By requiring $P > N$, Maloney and Wandell are able to replace Buchsbaum's reference surface with a geometrical landmark that serves the same purpose.

Note outcome of the algorithm consists of estimates of the intrinsic colors of surfaces known up to a single common 'lightness' scaling factor C. That is, if ε is the true light and $\sigma^{xy_1}, \sigma^{xy_2}, \ldots, \sigma^{xy_Z}$ are the true intrinsic colors at Z locations, the algorithm returns a normalized estimate of the light $\bar{\varepsilon} = \frac{1}{C}\varepsilon$ and corresponding estimates of the surface properties $\bar{\sigma}^{xy_1} = C\sigma^{xy_1}, \ldots, \bar{\sigma}^{xy_Z} = C\sigma^{xy_Z}$ where C is a unknown scaling factor common to all the estimates. (In contrast, the 'reference surface' algorithms above can use the reference surface to estimate the absolute power output of the illuminant.)

The Chromatic Aberration Algorithm of Funt and Ho. Ho (Ho, 1988; Funt & Ho, 1989; Ho, Funt & Drew, 1990) used the chromatic aberration inherent in lens systems to derive an estimate of the illuminant parameters ε up to an unknown scale factor, and showed that a visual system with only one receptor class can estimate the difference $\Delta(\lambda)$ between the color signals radiating from two adjoining surfaces separated by a sharp edge. Let S_{σ^1} and S_{σ^2} denote the two surface reflectance functions. Then,

$$\Delta(\lambda) = E_\varepsilon(\lambda)S_{\sigma^1}(\lambda) - E_\varepsilon(\lambda)S_{\sigma^2}(\lambda) \quad , \tag{29}$$

and the estimated color signal difference can be used to solve for ε up to an unknown scale factor.

Bayesian Algorithms. D'Zmura, Iverson and Singer (1995) and Freeman and Brainard (1995;

Brainard & Freeman, 1997) reformulate the problem of estimating illuminant and surface within the linear models framework as a problem in Bayesian statistical decision theory. This sort of approach presupposes, first of all, that the *prior distribution* of possible scenes composed of illuminants and surfaces, and the *likelihood function*, the likelihood of any possible retinal excitation patterns conditional on a particular scene, are both known.

We need not review the details of their work except to note, first of all, that essentially any algorithm that can be formulated as a maximum likelihood estimate can be trivially reformulated as a Bayesian estimator that can take advantage of knowledge of the prior distribution (Blackwell & Girshick, 1954; Ferguson, 1967). That is, each of the algorithms reviewed here has a Bayesian counterpart that makes use of the known prior distribution. Second, the Bayesian estimator can do no worse than its maximum likelihood counterpart (Blackwell & Girshick, 1954; Ferguson, 1967). That is, there is no point in evaluating them computationally except to determine *how much* better they perform. Last, the expected advantage in performance depends on the prior distribution of illuminants and surfaces. The true prior distribution of illuminants and surfaces in realistic environments is, currently, unknown. Both D'Zmura, Iverson and Singer and Freeman and Brainard simply assume prior distributions of lights and surfaces. Both, in particular, make the unrealistic assumption that surface patches at different locations in a scene are statistically independent.

Although there is little to be learned from their computational results, the *approach* they propose is sound and their computational experiments potentially important if redone with accurate estimates of prior distributions in realistic environments.

Multiple View Algorithms. Tsukada and Ohta (1990) and, independently, D'Zmura (1992) examined the information made available by viewing the same scene under multiple successive illuminants. Let the coordinates of the two illuminants be ε and ε'. D'Zmura first sets up the equations $\Lambda_{\varepsilon}\sigma^{xy} = \rho^{xy}$ and $\Lambda_{\varepsilon'}\sigma^{xy} = \rho'^{xy}$, describing the photoreceptor excitations corresponding to the surfaces at each location illuminated by each the two illuminants. He then shows how to solve these simultaneous equations for the σ^{xy}. In particular, for $P = N = M = 3$, two views permit recovery of the intrinsic colors σ^{xy} up to an overall unknown scaling factor. Thus, if a trichromatic visual system can arrange to view a scene under two successive illuminants, it can recover the surface descriptors up to an overall scaling. In a later section I describe a second application of these results (*The Shadow Algorithm of D'Zmura*).

The method of D'Zmura (1992) is a generalization of the subspace method of Maloney and Wandell. In an impressive series of articles, D'Zmura and Iverson (1993a;b; 1994; Iverson & D'Zmura, 1995a;b) study visual systems with P photoreceptors in the Flat World environment with surface reflectance modeled by an N-dimensional linear model and illuminants modeled by an M-dimensional linear model. They derive necessary and sufficient conditions for recovery of surface and light parameters when the same scene is seen under Q distinct illuminants including the case $Q = 1$ treated by Maloney and Wandell.

Illuminant Estimation. As discussed above, several of the algorithms above so far share a common form. Each seeks an explicit estimate of the illuminant parameters ε using the constraints imposed by Flat World and Linear Models. The results of Gilchrist and colleagues (Gilchrist, 1977; 1980; Gilchrist, Delman & Jacobsen, 1983; See also Gilchrist, 1994). demonstrate that lightness perception depends on the visual system's representation of the three-dimensional layout of scenes. In the next section, we will examine more complex, three-dimensional environments where there are additional cues available concerning the illuminant.

Environment II: Shape World

At this point, we must expand the mathematical models of illuminant and surface reflectance developed in the Flat World environment to include more realistic models of light/surface interaction in a three-dimensional scene. The coordinate xy in $S^{xy}(\lambda)$ now specifies both a location xy on the sensor array (retina) and the surface point imaged on the sensor array at that location. We allow only one surface point corresponding to each retinal location (ignoring the possibility of transparency). In most of this section, it is assumed that there is a single, punctate illuminant distant from the observer and the scene (for some of the discussion, there will be a small number of illuminants).

The algorithms in this section share the same common goal as the algorithms described in the previous. All describe a particular computation designed to capture information about the illuminant ε. Some of the algorithms use this information to compute Λ_ε and invert it, just as in the algorithms described above. Others simultaneously estimate illuminant and surface descriptors. The algorithms share a second common feature: each goes beyond the assumptions of a simple Flat World environment, basing its estimate of the illuminant on information only available from surfaces arranged in three dimensions.

Fig. 19.6 indicates some of the additional structure introduced into the environment: shading, specularity, mutual illumination, etc. The shape-from-shading literature is a source of models for illuminant-surface interaction in three-dimensional scenes (see Horn & Brooks, 1989; in particular, Horn & Sjoberg, 1989). This new environment will be referred to as *Shape World*.

Bidirectional Reflectance Density Functions. The surface spectral reflectance function of Flat World $S^{xy}(\lambda)$ is replaced by a *bidirectional reflectance density function* (BRDF), $S^{xy}(\lambda, V^{xy}, N^{xy}, I^{xy})$. N^{xy} is a unit normal vector to the surface at location xy. V^{xy} is a unit vector from the surface at xy in the direction of the visual system (the sensor array) and I^{xy} is the unit vector from the surface toward the (punctate) illuminant. The inset to Fig. 19.6 illustrates the interrelations among these vectors. Horn and Sjoberg (1989) provide a more detailed

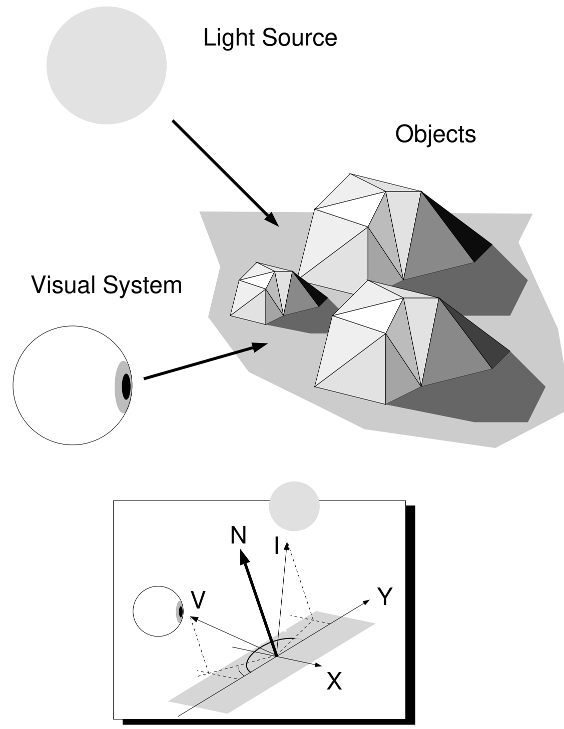

Figure 19.6: Shape World. A model environment for surface color perception that includes explicit representation of light/surface interaction in a three-dimensional scene, including shadows, inter-reflection ('mutual illumination') among surfaces, and specularities. See text for an explanation. The inset illustrates the coordinate system used in expressing the relation between the illuminant, surface, and visual system.

description of the BRDF and further references.

The possible interactions between light and surface are complex (Nassau, 1983; Weisskopf, 1968). Lee, Breneman and Schulte (1990) summarize some of the common models of BRDFs. Hurlbert (1998) discusses many of the models of BRDFs in use in computer vision and computer graphics. Oren and Nayar (1995; Nayar & Oren, 1995) review more recent work and propose a new model of the BRDF that takes into account the roughness of the surface at fine scales. Lee, Breneman and Schulte (1990) report measurements of BRDFs of a small number of surfaces, which we will return to below. There is currently no large collection of empirically-measured BDRFs available to test existing or new models.

I will not attempt to summarize the literature on models of BRDFs, which is extensive. Instead, I will describe only the properties of BRDFs that are needed in order to understand the Shape World algorithms reviewed below.

Geometry-Reflectance Separability. The first property we consider is *geometry-reflectance separability.*

$$(30) \quad S^{xy}(\lambda, V^{xy}, N^{xy}, I^{xy}) = S^{xy}(\lambda) G(V^{xy}, N^{xy}, I^{xy})$$

It states that the effect of a change in viewing conditions V^{xy}, N^{xy}, I^{xy} is simply to scale the surface reflectance function (Shafer, 1985). $S^{xy}(\lambda)$ will be referred to as the *spectral component* of the BRDF. The surface spectral reflectance $S^{xy}(\lambda)$ of Flat World is, of course, the spectral component $S^{xy}(\lambda)$ in Shape World, scaled by the geometric factors. It will cause no confusion to use the same notation for the Flat World spectral sensitivity function and the spectral component of a BDRF in Shape World (when geometric-reflectance separability holds).

If there are multiple illuminants or non-punctate illuminants, the color signal $L^{xy}(\lambda)$ is the superposition of the color signals corresponding to each illuminant point. It is possible that light from an illuminant can be successively absorbed and re-emitted by more than one surface before reaching the sensor array. In such cases, the light emitted from one surface acts as an illuminant to a second and is treated accordingly.

Diffuse-Specular Superposition. Surfaces often do not satisfy geometry-reflectance separability. Shafer (1985; Klinker, Shafer & Kanade, 1988) suggested that many surface BRDFs (corresponding to dielectric (non-conducting) surfaces such as plastics) may be represented as the sum of two surface BRDFs each of which satisfies geometric-reflectance separability:

$$(31) \quad S^{xy}(\lambda, V^{xy}, N^{xy}, I^{xy}) = m_d S^{xy}(\lambda) G(V^{xy}, N^{xy}, I^{xy}) + m_s Spec(\lambda) G'(V^{xy}, N^{xy}, I^{xy}).$$

The first term in the summation is termed the *diffuse component*, the second, the *specular component*. $Spec(\lambda) = 1$ is the surface spectral sensitivity function of a perfect reflector, G (.,.,.,) is the geometric function for a diffuse surface, and G' (.,.,.,) the geometric function for a specular (mirror-like) surface. The weights m_d and m_s control the diffuse-specular balance. Note that $Spec(\lambda)$ is the same at every location while the $S^{xy}(\lambda)$ may vary. The constraint on surfaces embodied in Eqn 31 will be referred to as the *diffuse-specular superposition* property. The Neutral Interface Model of Lee, Breneman and Schulte (1990) exhibits this property. I will refer to $S^{xy}(\lambda)$ in Eqn 31 as the spectral component of the BDRF.

Lee, Breneman and Schulte (1990) tested whether surfaces satisfied the diffuse-specular superposition property. They measured the spectral reflectance functions of nine surface materials for different viewing geometries. They found that the property was satisfied for some of the surface materials (including ‘yellow plastic cylinder’, ‘green leaf’, and ‘orange peel’) but not all (e.g. ‘blue paper’, ‘maroon bowl’). Tominaga and Wandell (1989; 1990) report empirical tests of the property as well.

In the description of the algorithms, I will assume that $N = M = P = 3$. That is, the possible illuminants are drawn from a linear model with three parameters $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$, the visual system has three classes of photoreceptor, and the possible spectral components of surfaces are also drawn from a linear model with three parameters $(\sigma_1, \sigma_2, \sigma_3)$. I will generally assume that either geometric-reflectance separability or diffuse-specular condition holds. The term *viewing geometry* refers to the relative positions of surfaces, illuminants, and the observer in Shape World.

Viewing the Light/White Surface. The first model we consider is one where the observer is able to look around and somehow identify the illuminant itself. Viewing the illuminant is equivalent to having a single white reference surface in Buchsbaum’s algorithm. We can solve directly for the light. Alternatively, the observer may attempt to identify a perfectly reflective (‘white’) patch in the scene and use it just as a glimpse of the illuminant. These approaches fail, of course, when the scene contains no white patch or directly observable illuminant, or when the observer cannot identify one or the other.

Specularity Algorithms. Lee (1986) and D'Zmura and Lennie (1986) proposed algorithms based on the diffuse-specular superposition property. Fig. 19.7 illustrates the key idea. Suppose we have two *objects*, each with a uniform diffuse-specular surface and with sufficient variation in the viewing geometry so that the color signals from the object are not uniform. Further, assume that the objects have distinct diffuse components $S(\lambda)$ and $S'(\lambda)$.

The receptor excitations from the first object will be weighted mixtures of a diffuse color signal,

$$D_k = \int E(\lambda)S(\lambda)R_k(\lambda)d\lambda,\ k = 1, 2, 3 \tag{32}$$

and a specular color signal

$$Spec_k = \int E(\lambda)Spec(\lambda)R_k(\lambda)d\lambda,\ k = 1, 2, 3 \tag{33}$$

where the weights are determined by the surface mixture parameters m_d and m_s and the viewing geometry. The receptor excitations for the second object will also be weighted mixture of a diffuse color signal,

$$D'_k = \int E(\lambda)S'(\lambda)R_k(\lambda)d\lambda,\ k = 1, 2, 3 \tag{34}$$

and the same specular color signal

$$Spec_k = \int E(\lambda)Spec(\lambda)R_k(\lambda)d\lambda,\ k = 1, 2, 3. \tag{35}$$

Then the photoreceptor excitation ρ^{xy} for the first object at location xy is a weighted mixture,

$$\rho^{xy} = \alpha^{xy}D + \beta^{xy}Spec \quad , \tag{36}$$

where $D = [D_1, D_2, D_3]'$ and $Spec = [Spec_1, Spec_2, Spec_3]$. *Spec* is the color of the illuminant (as defined in the *RGB Heuristic*). As shown in Fig. 19.7, all the photoreceptor excitations must lie in a plane through the origin spanned by D and $Spec$.

By a similar argument, all the photoreceptor excitations for the second object, must lie in the plane through the origin spanned by D' and $Spec$:

$$\rho'^{xy} = \alpha^{xy}D' + \beta^{xy}Spec \quad . \tag{37}$$

If there are enough different points from each of the objects to permit detection and estimation of the two planes, and if the planes are distinct, the intersection of the planes determines $Spec$ up to an unknown scaling. From this we can learn ε up to an overall scaling.

One strength of the method is that it does not assume that bright, specular highlights visible on the surface reflect precisely the spectral power distribution of the illuminant. Instead, the light from each point on the surface is modeled as a weighted mixture of a specular component and a diffuse component. The specular component has the spectral power distribution of the illuminant. The spectral power distribution of the diffuse component is characteristic of the object and is assumed to be independent of viewing angle.

Klinker, Shafer and Kanade (1988) analyze the diffuse-specular superposition constraint further and demonstrate that the diffuse-specular mixtures are not simply confined to a plane in Fig. 19.6 but will typically form a characteristic skewed 'T' shape within the plane. The skewed 'T' is a color space 'feature' that can be extracted and used in estimating the illuminant. Healey (1991) shows that it is possible to recover the illuminant using one diffuse specular object if the viewing geometry guarantees that some point on the object will exhibit a pure diffuse reflectance, and another point a pure specular reflectance with 100% reflectance.

The Mutual Reflection Algorithm of Drew and Funt. Drew and Funt (1990) describe how it is possible to estimate the illuminant parameters ε using the mutual reflection between adjacent surfaces in Shape World. Light emitted from one surface with spectral component $S^1(\lambda)$ may reach a second nearby surface

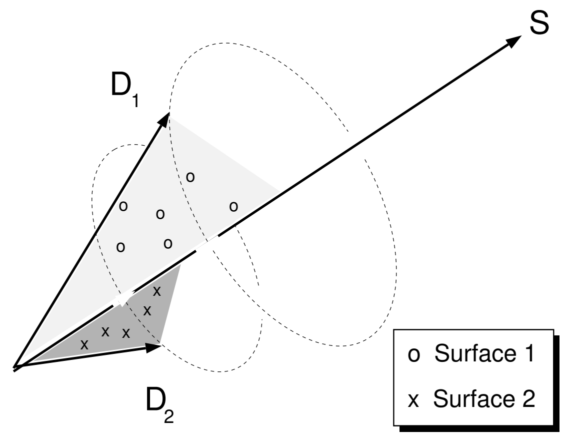

Figure 19.7: The Specular-Diffuse Constraint. The photoreceptor excitations for an object with a homogeneous surface that satisfies the Specular-Diffuse constraint (see text) must lie in a plane. The intersection of two such planes, corresponding to two different specular-diffuse objects, marks the chromaticity of the illuminant.

with spectral component $S^2(\lambda)$ and be absorbed and re-emitted in the direction of the observer. The color signal reaching the observer is proportional to $E(\lambda)S^1(\lambda)S^2(\lambda)$, where the constant of proportionality is determined by the viewing geometry.

Drew and Funt (1990) restrict attention to the case where each ray of light from the illuminant encounters either one or two surfaces in its transit to the sensor array (a 'single-bounce' or 'one-bounce' model of mutual reflection). They express the color signal as a sum of 'zero-bounce' and 'one-bounce' components for each surface, substitute Eqs 12 and 13 into their expressions, and show that the resulting third-degree equations in ε and σ^1, σ^2 can be solved by iterative methods. The method is 'reasonably robust (p. 399)' and can be used even when the two mutually-illuminating surfaces have the same spectral components, e.g., at a corner in a room.

The Shadow Algorithm of D'Zmura. D'Zmura (1992) points out that his multiple illuminant algorithm can be applied across shadow boundaries. At a shadow boundary that does not coincide with the boundary between two different surfaces, the same surface is illuminated by two different illuminants. If the shadow boundary intersects three or more surfaces, the multiple illuminant algorithm of D'Zmura (1992) can be applied to recover estimates of three-dimensional illuminants and surface parameters with only three photoreceptor classes.

Other Work. We have, so far, not made use of an evident constraint on *S(λ)*, the surface spectral reflectance function, that $0 \le S(\lambda) \le 1$. Forsyth (1990) develops an approach to color constancy that makes use of such physical-realizability constraints. Rubner and Schultern (1989) apply a regularization approach similar in spirit to the Bayesian. Other work, not further discussed here, includes that of Gershon and Jepson (1988; 1989).

Other Environments. The Shape World Environment is not a realistic model of natural scenes. One glaring deficiency is the restriction to one punctate illuminant or a small number of punctate illuminants that illuminate the scene successively or illuminate non-overlapping parts of the scene. Research is needed to develop realistic models of light-surface interaction in scenes in a true *Real World.*

Model Failure and Approximate Constancy

The Consequences of Model Failure. As noted above, linear models of illuminant and surface reflectance with three dimensions do not provide perfect fits to empirical data sets. Accordingly, if any of the algorithms described above for 'Flat World' or 'Shape World' are to be applied in an environment composed of illuminants and surfaces drawn from such empirical collections, the consequences of discrepancies between idealization and application must be considered.

We can model the effect of linear model failure within the linear model framework as follows (Maloney, 1984). We represent the true spectral power distribution of the illuminant as the sum of the linear model approximation to it and a residual illumination term $e(\lambda)$;

(38) $$E(\lambda) = E_{\varepsilon}(\lambda) + e(\lambda) \quad .$$

A similar decomposition is adopted for the surface reflectance at location xy:

(39) $$S^{xy}(\lambda) = S^{xy}_{\sigma}(\lambda) + s^{xy}(\lambda) \quad ,$$

where $s^{xy}(\lambda)$ is the residual surface spectral reflectance.

The color signal $L^{xy}(\lambda)$ at location xy is then,

(40) $$L^{xy}(\lambda) = E_{\varepsilon}(\lambda)S^{xy}_{\sigma}(\lambda) + e(\lambda)S^{xy}_{\sigma}(\lambda) + E_{\varepsilon}(\lambda)s^{xy}(\lambda) + e(\lambda)s^{xy}(\lambda).$$

Let

(41) $$L^{xy}_{\varepsilon\sigma}(\lambda) = E_{\varepsilon}(\lambda)S^{xy}_{\sigma}(\lambda) \quad ,$$

the color signal from location xy if the linear models had accurately captured the illuminant spectral power distribution and the surface spectral reflectance. The discrepancy between $L^{xy}(\lambda)$ and $L^{xy}_{\varepsilon\sigma}(\lambda)$ is the sum of three terms,

(42) $$e(\lambda)S_{\sigma}(\lambda) + E_{\varepsilon}(\lambda)s(\lambda) + e(\lambda)s(\lambda) \quad ,$$

corresponding to (a) the model light $E_{\varepsilon}(\lambda)$ shining on the residual surface $s^{xy}(\lambda)$, (b) the residual light $e(\lambda)$ shining on the model surface $S^{xy}_{\sigma}(\lambda)$ and (c) the residual light $e(\lambda)$ shining on the residual surface $s^{xy}(\lambda)$. We can view the scene, then, as the superposition of two scenes: the *model scene* $L^{xy}_{\varepsilon\sigma}(\lambda)$, whose parameters we wish to estimate, and the *error scene* in Eqn 42, whose presence perturbs our estimates. The magnitude of the perturbation at location xy is $\Delta\rho^{xy} = [\Delta\rho^{xy}_1, \Delta\rho^{xy}_2, \Delta\rho^{xy}_3]$, where

$$\Delta\rho^{xy}_k = \int e(\lambda)S_{\sigma}(\lambda)R_k(\lambda)d\lambda + \int E_{\varepsilon}(\lambda)s(\lambda)R_k(\lambda)d\lambda + \int e(\lambda)s(\lambda)R_k(\lambda)d\lambda.$$

The error term $\Delta\rho^{xy}$ depends, first of all, on the bases E_i and S_j and, of course, we would like to choose bases to minimize the impact of the error scene. Marimont and Wandell (1992) describe how to do this for any empirical collection of lights, surfaces, and photoreceptor sensitivities. Brill (1978; 1979), Buchsbaum (1980), Maloney (1984), and Wandell (1987) describe perturbation analyses of some of the Flat World algorithms. A few results can be found in Maloney (1984, Chap. 4).

Linear Models Algorithms and Human Vision

In this section, we consider the relation between the work reviewed above and classical models of color vision framed in terms of hypothesized color channels and transformation of color information through successive stages. Earlier forms of this discussion may be found in Maloney (1984; 1992).

Adaptational State and Adaptational Control. Let ρ^{xy} denote, as above, the excitations of the three photoreceptor classes in a small retinal patch near xy. It is well known that ρ^{xy} alone does not determine the perceived color appearance of the patch at xy which may be profoundly influenced by photoreceptor excitations in other retinal areas, in one or both eyes. I assume that there are *mechanisms of color appearance* whose excitations, written in vector form $\mu^{xy} = [\mu^{xy}_1, \ldots, \mu^{xy}_W]$, determine performance in tasks measuring color appearance in a small patch at location xy.

The number W of mechanisms of color appearance is unknown but, for the normal trichromatic observer, it is plausible that there are three independent mechanisms, and possibly others. For simplicity, I assume W=3 in the following. The relationship between μ^{xy} and ρ^{xy} can be written as,

(43) $$\mu^{xy} = \tau_{\Sigma}(\rho^{xy}) \quad ,$$

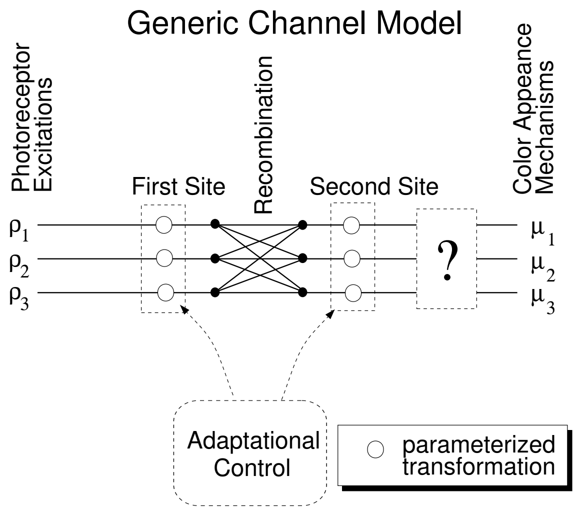

Figure 19.8: The 'Generic Channels' Framework. The open circles mark points in channels where variable gain controls or additive biases may be applied. See text for an explanation.

where τ_Σ: $\boldsymbol{R}^3 \rightarrow \boldsymbol{R}^3$ is a transformation from one three-dimensional vector space to a second, and where Σ denotes the influence of the surround: the photoreceptor excitations of previous retinal stimulation, and current retinal stimulation at locations other than *xy*.

The *color transformations*, denoted $\tau_\Sigma(.)$, represent the net effect of all the transformations (at various 'stages') performed by the visual system on the initial retinal information corresponding to location *xy*. The precise form of $\tau_\Sigma(.)$ has received considerable attention. There is widespread agreement that initial information in three color pathways passes through a 'First-Site' (Fig. 19.8) where it it scaled multiplicatively (Chichilnisky & Wandell, 1995), and shifted additively (Burnham, Evans & Newhall, 1957; Jameson & Hurvich, 1964; Walraven, 1976; Shevell, 1978), followed by an opponent recombination (Hurvich & Jameson, 1957; See Hurvich, 1981; Wandell, 1995; Kaiser & Boynton, 1996), and by 'Second Site' multiplicative attenuation (Hurvich & Jameson, 1957; Webster & Mollon, 1995; Brown & MacLeod, 1997).

The transformation $\tau_\Sigma(.)$ is likely non-linear if examined across a wide range of overall scene intensities, but, "[f]or modest signals under a constant adaptation state, single-cell responses and psychophysical sensitivity are consistent with mechanisms that respond to simple sums or differences of the cone contrasts" (Webster, 1996, p. 595). Under such circumstances, we can assume that the $\tau_\Sigma(.)$ are approximately affine transformations of the form,

$$\tau_\Sigma(\rho) = D_2 H D_1(\rho - O_1) \quad , \tag{44}$$

where D_1 and D_2 are 3x3 diagonal matrices corresponding to the multiplicative attenuations at first and second sites, respectively, O_1 is an offset 3-vector, corresponding to a first site additive shift, and H is a 3x3 opponent transformation matrix.

Of interest to us here are the components of $\tau_\Sigma(.)$ that may change in response to Σ: the multiplicative attenuations at first and second site, and the additive shifts. I refer to these variable components of models of $\tau_\Sigma(.)$ as *parameters*. These parameters depend on the surround Σ. If we regard the opponent transformation as fixed (non-parameterized), the parameters above include the first-site offsets and attenuations, $O = [\pi_1, \pi_2, \pi_3]$, and $D_1 = Diag[\pi_4, \pi_5, \pi_6]$, and the second site attenuations $D_2 = Diag[\pi_7, \pi_8, \pi_9]$. The state vector of nine parameters $\pi = [\pi_1, ..., \pi_9]$ determines $\tau_\Sigma(.)$ which may be written $\tau_{\pi(\Sigma)}(.)$ or $\tau_\pi(.)$ to emphasize that the influence of the surround/scene reduces to selecting the settings of a handful of parameters, π. This formulation of the study of color appearance is essentially that employed by Krantz (1968) in modeling the effect of 'context' on color appearance.

Stiles (1961, p. 264) proposed that the study of color vision be considered as the study of two processes, *local retinal adaptational state* and the process that selects the local state, *the control of adaptation*: "we anticipate that a small number of variables -- adaptation variables -- will define the condition of a particular visual area at a given time, instead of the indefinitely many that would be required to specify the conditioning stimuli. The adaptation concept -- if it works -- divides the original problem into two: what are the values of the adaptation variables corresponding to different sets of conditioning stimuli, and how does adaptation, so defined, modify the visual response to given test stimuli."

The exact number and nature of the local retinal adaptational state parameters π is not important to the discussion below. Various authors have suggested models with from 3 to 9 parameters. See the discussion in Brainard, Brunt and Speigle (1997). What is important is that π determines the affine transformation $\tau_{\pi}(\Sigma)$ and that the surround/scene Σ determines π:

$$(45) \qquad \Sigma \to \pi(\Sigma) \to \tau_{\pi(\Sigma)} \quad .$$

Again, the advantage of this approach 'if it works', is that is reduces the study of color vision to two questions:

1. *What are the possible local retinal adaptation states?*[5] That is, What are the local retinal state parameters and what possible settings can they take on? Let the set of all possible parameter settings π be denoted Π.

2. *How are possible conditioning stimuli mapped to a choice of retinal adaptational state?* In our terminology, the set Π of possible parameter settings corresponds to the possible retinal adaptational states, and the control of adaptation is the selection process $\pi(\Sigma)$ that sets the retinal state, selecting the color transformation as a function of the surround Σ.

As discussed above, we know something about the structure of the $\tau_{\pi(\Sigma)}$, the affine transformations of Eqn 44. Until recently, the second question has received considerably less attention than the first. I will argue below in Linear Models and Adaptational Control that the linear model algorithms make their most original and significant contributions as candidate theories of adaptational control, addressing the second of the two questions. In the intervening sections, I describe a small number of earlier theories of adaptational control necessary to the argument. The following is neither a comprehensive nor a representative discussion of models of adaptational control, omitting important work from Judd (1940) to Hunt (1991).

5. The term `local retinal adaptational state' will be used to refer to the state of the pathways that carry chromatic information corresponding to a small region of the visual scene. I do not mean to imply that chromatic adaptation is purely retinal but rather to emphasize that the local retinal color information is transformed at successive stages of the visual system in ways that are not fixed.

Recent work of Zaidi, Spehar and DeBonet (1997; 1998) is not directly relevant to this review, but certainly deserves mention. Zaidi and colleagues model the effect of illuminant changes at both first and second sites in the color channels. They then derive simple transformations that are intended to cancel the illumination-induced changes at the second site. They propose mechanisms of adaptational control that are sensitive to spatial variation in the scene viewed. The resulting 'low-level' (their terminology) approach to color constancy is an elegant alternative to the sort of physics-based approach described here. It would be particularly interesting to analyze the behavior of their model in response to specularity, shadow, etc. See Poirson and Maloney (1998) for additional discussion. Poirson and Maloney (1998) and Webster (1996) proposes contrasting models of chromatic adaptation and adaptational control.

von Kries Algorithms. Von Kries (1902/1970; 1905/1970) proposed that retinal adaptation states be identified with scalings of photoreceptor excitations. Translated into the notation above, this is equivalent to restricting the $\tau(.)$ to be diagonal matrices,

$$(46) \qquad \mu^{xy} = D\rho^{xy}$$

where $D = Diag[\pi_1, \pi_2, \pi_3]$. Von Kries' discussion of the Coefficient Law for simple center-surround configurations of stimuli includes the assumption that $\pi_k^{xy} = 1/\rho_k^{\Sigma}$ where ρ_k^{Σ} is the excitation of the kth class of receptor in response to the large conditioning Surround field. It is not made clear how the visual system might choose the coefficients π_k in more complex scenes, lacking a center surround structure.

Helson (1934; 1938) proposed a model of color adaptation in which the coefficient π_k was the inverse of the average of photoreceptor excitations within the kth class across a region of the retina. This von Kries-Helson model coincides with von Kries' proposal for the simple center surround configuration, and is one possible extension of von Kries adaptation to arbitrary scenes. The von Kries-Helson model of adaptation is closely related to 'Gray World' (Buchsbaum, 1980, pp. 24) and Buchsbaum's algorithm is intended as a natural generalization of von Kries-Helson that takes into account the linear models constraints on light and surface, making explicit the dependence of the algorithm's performance on environmental constraints.

The retinex algorithm of Land and McCann (1971) assumes Eqn 46 and propose a complicated computation of the parameters (π_1, π_2, π_3). In particular, the value of the parameter π_k depends only on the photoreceptor excitations of the kth class of photoreceptors, $k = 1,2,3$. For the purposes of adaptational control, then, the visual system is divided into three, isolated *retinexes*, pathways driven by photoreceptor excitations from a single class of photoreceptor but from widely separated areas in the retina.

Brainard and Wandell (1986) analyze this retinex algorithm and conclude that the computation, in effect, selects, as π_k, the inverse of the largest photoreceptor excitation in the kth class of photoreceptors across a region of the retina. A heuristic motivation for this choice of π is that, if there is a perfectly reflecting white surface visible in the scene, the retinex (Land & McCann, 1971) algorithm will 'lock on' to it and, in effect, scale by coefficients inversely proportional to the color of the illuminant (See *The RGB Heuristic*).

Brainard and Wandell (1986) analyze the later retinex algorithms of Land (1983; 1986) and conclude that they, in effect, tend to choose the inverse of the spatially-weighted *geometric mean* of the photoreceptors in the kth class of photoreceptors as the π_k. In contrast, the von Kries-Helson algorithm described above used the inverse of the spatially-weighted arithmetic mean.

These algorithms are all examples of so-called 'Lightness Algorithms' (see Hurlbert, 1986; 1998; Maloney, 1992 for reviews). The term 'lightness algorithm' is unfortunate, for there is no special link between such algorithms and the study of lightness (Gilchrist, 1994). I refer to them as *von Kries Algorithms*.

Brill and West (1981; West & Brill, 1982) study necessary and sufficient conditions under which von Kries adaptation can result in color constancy in response to changes in the illuminant.

How well does von Kries describe human vision? It is believed by many that von Kries adaptation, possibly with an additive bias, is an accurate model for human observers, at least in simple scenes. See Brainard (1998) for a discussion and Chichilnisky and Wandell (1995) for representative results. Foster and his colleagues (Craven and Foster, 1992; Foster & Nascimento, 1994) found that simulated von Kries transformations on images were judged illuminant changes more often than simulated illuminant changes. Note, however, that these results imply that a von Kries transform accounts for much of visual adaptation, not all. The analyses typically performed do not look for small systematic patterns of deviation consistent with a contribution from second site adaptation.

Recent experimental work is consistent with von Kries adaptation but reject von Kries-Helson adaptation in even modestly complex scenes (Jenness & Shevell, 1995; Mausfeld, 1998; Brainard, 1998; Brown & MacLeod, 1997). For example, Brainard (1998; discussed in Kaiser & Boynton, 1996, p. 519) draped part of a room in red vinyl cloth. This change in average chromaticity of a large part of the scene should have resulted, according to von Kries-Helson, in a substantial change in the von Kries adaptational coefficients and a marked change in the color appearances of objects in the scene. It did not. In more recent experiments, reported at the Optical Society Meetings, Kraft and Brainard (1997) described experiments in which the mean chromaticity of the color signals across the scene was held fixed while the illuminant was varied. The observer looked into a three foot cubical chamber. He viewed one of two scenes, one created with a neutral illuminant on a variety of mostly gray surfaces and one created with a red illuminant on a variety of mostly blue surfaces. The mean color signal across the observer's field of view was the same. Even with this mean held fixed, the visual system still adjusted substantially to the illuminant. Their results suggest that, under these experimental circumstances, the mean color signal is influential but not decisive

(D.H. Brainard, personal communication).

von Kries-Ives Adaptation. Clearly, von Kries adaptation is not co-extensive with von Kries-Helson adaptation. The retinex algorithms assume the same local retinal adaptation states (von Kries) but propose different rules for controlling the choice of adaptational state.

We can distinguish one particular class of von Kries algorithm in which the coefficients π_k are set equal to the chromaticity of the illuminant defined above in *The RGB Heuristic*. I will refer to this form of adaptation as *von Kries-Ives* adaptation. Again, without constraints on illuminants and surfaces in the scene, such an algorithm will not, in general, 'discount the illuminant' or result in approximate color constancy.[6]

A crucial difference between von Kries-Ives adaptation and the other von Kries algorithms, described above, is that it is not obvious how to compute the chromaticity of the illuminant from retinal excitations. We can, however, view this limitation in another way. Given any source of information about the chromaticity of the illumination, we can implement a von Kries-Ives algorithm based on it. Knowledge of the illumination parameters ε is enough to determine the chromaticity of the illuminant. Hence any of the linear model algorithms that estimate ε can be used to implement a von Kries-Ives algorithm. We next consider the relation between linear models algorithms and adaptational control in biological vision.

Linear Models and Adaptational Control. Taken as models of chromatic adaptation the linear models algorithms highlight the control mechanisms of adaptation, the second of Stiles' processes. Each of the linear models algorithms described above can be taken as a theory of the sources of information in a scene that affect local adaptational state (Maloney & Varner, 1986).

Many but not all of the linear models algorithms derive an estimate $\hat{\varepsilon}$ of the illuminant vector ε, and then compute estimates of the intrinsic surface colors, σ^{xy} by applying the inverse matrix to the photoreceptor receptor excitations,

$$(47) \qquad \hat{\sigma}^{xy} = \Lambda_{\hat{\varepsilon}}^{-1} \rho^{xy} \quad .$$

Any one-to-one transformation of the $\hat{\sigma}^{xy}$ would equally-well serve to determine color appearance, at least in the absence of further assumptions concerning the behavior of the mechanisms. If we let, $F\colon \boldsymbol{R}^3 \to \boldsymbol{R}^3$ denote any one-to-one function, we can hypothesize a link between the linear models algorithms and Eqn 45 by writing,

$$(48) \qquad \mu^{xy} = F(\hat{\sigma}^{xy}) = F(\Lambda_{\hat{\varepsilon}}^{-1} \rho^{xy}) \quad .$$

Again, comparing the two equations, we see that the linear model algorithms reduce the information present in the surround (the remainder of the scene) to the handful of numbers $\hat{\varepsilon}$. That is, the control processes implicit in each of the algorithms above may use information based on specularity cues, mutual illumination cues, shading cues, and so forth, but their contribution to the equation above is reduced to their effect on the handful of parameters $\hat{\varepsilon}$.

The algorithms described above puts the possible *local retinal adaptational states* into one-to-one correspondence with the matrices $\Lambda_{\varepsilon}^{-1}$, the matrices that compensate for each of the possible illuminants ε. Recall that these matrices form a three-parameter family,

$$(49) \qquad \Lambda_{\varepsilon}^{-1} = [\varepsilon_1 \Lambda_1 + \varepsilon_2 \Lambda_2 + \varepsilon_3 \Lambda_3]^{-1}$$

where the Λ_i are fixed, determined by the linear models. Within the linear models framework, this family of matrices, parameterized by ε corresponds to the family of transformations $\tau_{\pi}(.)$ parameterized by π. It is interesting to ask, how close the two matrix families are to one another for realistic models of illuminant

6. It is sometimes claimed that von Kries algorithms do not require strong assumptions about possible illuminants and surface reflectances such as are presupposed by 'linear model algorithms'. On the contrary, they do. The relation between performance and environment is often left unstated, disguised as some form of the *RGB Heuristic* and Eqn 5.

and surface and various channel-stage models of early vision including the von Kries model and the full affine model of Eqn 44. This is an open question.

If the two sets of linear transformations are not identical, then we could propose a constrained form of any of the linear model algorithms where the illuminant estimate is computed as prescribed by the algorithm, but the matrix $\tau_{\pi}(.)$ chosen is not the correct matrix $\Lambda_{\hat{\varepsilon}}^{-1}$ but the best approximation to it among the possible choices of $\tau_{\pi}(.)$. Linear model algorithms constrained to use the transformations available in human vision are, in effect, models of adaptation control based on a particular cue to the illuminant.

Privileged Environments. Many of the algorithms described above, if used across a range where their environmental assumptions are satisfied, would provide accurate estimates of surface properties. This range, described by low-dimensional linear models, may include illuminants that markedly differ in color and surfaces spectral reflectances that also span a wide chromatic gamut, and still permit perfect color constancy. *Large* changes in the physical light can be consistent with essentially perfect color constancy.

Yet, once lights or surfaces are drawn from outside the linear models to which the algorithm is attuned, the estimates of surface properties will fail to be constant for almost all changes in the light. . Consequently, if one were to select lights and surfaces haphazardly, one would almost certainly conclude that a visual system which, in fact, embodied one of the algorithms above was at best approximately color constant, failing less on some occasions than on others. It would be easy to overlook the class of illuminants and surfaces where the algorithms operate flawlessly. The existence of such *Privileged Environments* of lights and surfaces is perhaps the most significant prediction for the study of human color vision.

The preceding discussion leads to the following considerations. *What, precisely, are the linear models for surface reflectances and for illuminants that lead to optimal human performance in surface color perception experiments?* This question remains open. Researchers typically pick plausible simulated illuminants and surface reflectances drawn from Judd, MacAdam and Wyszecki (1964)'s linear model of daylight and a linear model derived from Munsell sample measurements (e.g. Brainard & Wandell, 1991). The resulting measures of performance likely understate optimal human performance as noted in the section on methodology.

Light Estimation as Cue Combination

Consideration of the algorithms above suggests that there are several possible cues to the illuminant. It is natural to consider the estimation of the illuminant as a *cue combination problem*, analogous to cue combination in depth/shape vision. Kaiser and Boynton (1996, p. 521) have previously suggested that illuminant estimation is best thought of as combination of information from multiple cues.

Consider the following simple model of illuminant estimation: each of several cues (specularity, etc) is used to estimate the illuminant parameters ε. One cue could correspond to the Stable Mean Assumption (also known as 'Gray World'), deriving an estimate of the illuminant parameters ε from a weighted average of the photoreceptor excitations across the scene. This estimate is denoted $\hat{\varepsilon}_{SM}$. One or more cues might be based on specular information, leading to the estimate, $\hat{\varepsilon}_{Spec}$. The weighted average of estimates based on these cues is taken to be a single illuminant estimate $\hat{\varepsilon}$ that controls the color transformation as discussed in the previous section.

$$\hat{\varepsilon} = \alpha_{SM}\hat{\varepsilon}_{SM} + \alpha_{Spec}\hat{\varepsilon}_{Spec} + \dots \tag{50}$$

The weights $\alpha_{.}$ are non-negative and sum to 1, and each value $\alpha_{.}$ reflect the relative importance assigned to the corresponding cue.

Landy et al. (1995) report empirical tests which imply that depth cue weights change. The implication for surface color perception is that the relative weight assigned to different estimates of the illuminant from different cue types may also change. In particular, consider the sort of experiment where almost all cues to the illuminant are missing: the observer views a large, uniform surround with a small number of test regions superimposed. It is plausible that the only cue to the illuminant remaining is the Stable Mean cue, the chromaticity of the surround. Then the visual system might set α_{SM} to essentially 1, effectively basing its light estimate on this single cue. The result would be the sort of 'von Kries-Helson' behavior observed experimentally. In more complex scenes, the weight assigned to the Stable Mean cue might be reduced as other cues become available. This would explain the puzzling observation discussed previously above, that the von Kries-Helson pattern of response found in many laboratories does not generalize to the outside world (or even David Brainard's laboratory).

A second, and surprising analogy between depth cue combination and illuminant estimation, is that not all cues to the illuminant provide full information about the illuminant parameters ε. Some of the methods (such as the Subspace Algorithm) lead to estimates of ε that include an unknown multiplicative scale factor. It would be meaningless to average such an estimate of 'relative' ε together with a full estimate of ε.

An analogous problem arises in depth cue combination since certain depth cues (such as relative size) provide depth information up to an unknown multiplicative scale factor. The problem of combining depth estimates some of which are in meters, others in only relative units, is termed *cue promotion* by Maloney and Landy (1989) and is treated further in Landy et al. (1995).

In summary, I advance the proposal that a central step in the control of adaptation is computation of illuminant estimates based on illuminant cues. In scenes rich in accurate illuminant cues, the estimate will be not too far from correct, and approximate color constancy results. In simple center-surround scenes, the single cue would seem to be the large surround, the Stable Mean cue, resulting in approximate von Kries-Helson behavior. What the cues to the illuminant employed in human vision are, and how they are combined, remain open questions. Many of the algorithms above can be identified with potential cues to the illuminant, but it is likely that some of them will prove to be irrelevant to human vision. It is also plausible that better models of illuminant and surface interactions in complex scenes will lead to discovery of other candidate cues to the illuminant.

Summary

Physics-based models of surface color perception has implications for the study of color vision as a whole.These include, first of all, the importance of studying adaptational control as well as the structure of single color channels. A second implication is that we need to know more about how light and surface interact in scenes if we are to understand how color vision proceeds in complex, natural environments.

In addition, the investigation of explicit models of light and surface in scenes has rled us to a number of candidate cues to the illumination in scenes and has permitted the development of precise models of how this information may be recovered. It remains to be seen whether any of these cues to the illuminant is used by biological vision systems.

Last, once the possibility that there are many possible cues to the illuminant is considered, it is natural to ask how a visual system, ideal or biological might pick and choose among them, and this leads to consideration of the similarities between surface color perception and depth/shape perception.

Not addressed in this review is the issue of the relationship between shape perception and surface color perception. If the spatial layout of the scene is known, including the precise shapes and positions of objects, it is entirely plausible that more accurate estimates of the illuminant parameters can be computed. Also not addressed is the issue of complex illuminant environments with multiple, non-punctate sources of light.

Acknowledgments

This work was supported by a grant from the National Eye Institute EY08266 and a fellowship from the Deutscher Akademischer Austauschdienst. The data from Vrhel, Gershon and Iwan (1994) were provided by Ron Gershon. I thank Ron Gershon, his co-authors, and the Kodak company for making them available. Several people were kind enough to read and comment on earlier drafts: David Brainard, Michael D'Zmura, Karl Gegenfurtner, Susan Hodge, Michael Landy, and Joong Nam Yang. I am very grateful to them all. Special thanks to Rainer Mausfeld for sharing the quote from Ulrich von Strassburg, and to Michael Brill for pointing out the enduring importance of Ives (1912b).

References

Apostol, T.M. (1969) *Calculus, 2nd Edition, Volume II.* Xerox, Waltham, Massachusetts.

Arend, L.E. Jr. (1993) how much does illuminant color affect unattributed colors? *Journal of the Optical Society of America A*, 10, 2134-2147.

Arend, L.E. & Reeves, A. (1986) Simultaneous color constancy. *Journal of the Optical Society of America A*, 3, 1743-1751.

Arend, L.E., Reeves, A., Schirillo, J. & Goldstein, R. (1991) Simultaneous color constancy: Papers with diverse Munsell values. *Journal of the Optical Society of America A*, 8, 661-672

Bäuml, K.H. (1994) Color appearance: effects of illuminant changes under different surface collections. *Journal of the Optical Society of America A*, 12, 531-542.

Bäuml, K.H. (1995) Illuminant changes under different surface collections: Examining some principles of color appearance. *Journal of the Optical Society of America A*, 12, 261-271.

Beck, J. (1972) *Surface Color Perception*. Cornell University Press, Ithaca, New York.

Berns, R.S. & Gorzynski, M.E. (1991) Simulating surface colors on CRT displays: the importance of cognitive clues. *AIC Conference: Colour and Light*, 21-24.

Blackwell, D. & Girshick, M.A. (1954) *Theory of Games and Statistical Decisions*. Wiley, New York.

Brainard, D.H. (1998) Color constancy in the nearly natural image. 2. achromatic loci. *Journal of the Optical Society of America A*, 15, 307-325.

Brainard, D.H., Brunt, W.A. & Speigle, J.M. (1997) Color constancy in the nearly natural image. 1. Asymmetric matches. *Journal of the Optical Society of America A*, 14, 2091-2110.

Brainard, D.H. & Freeman, W.T. (1997) Bayesian color constancy. *Journal of the Optical Society A*, 14, 1393-1411.

Brainard, D.H., Rutherford, M.D. & Kraft, J.M. (1997) Color constancy compared: Experiments with real images and color monitors. *Investigative Ophthalmology & Visual Science* (Suppl.), 38, 476.

Brainard, D.H. & Wandell, B.A. (1986) An analysis of the retinex theory of color vision. *Journal of the Optical Society of America A*, 3, 1651-1661.

Brainard, D.H. & Wandell, B.A. (1991) A bilinear model of the illuminant's effect on color appearance. In *Computational Models of Visual Processing* (eds. Movshon, J.A. & Landy, M.S.) MIT Press, Cambridge, MA.

Brainard, D.H. & Wandell, B.A. (1992) Asymmetric color-matching: How color appearance depends on the illuminant. *Journal of the Optical Society of America A*, 9, 1433-1448.

Brill, M.H. (1978) A device performing illuminant-invariant assessment of chromatic relations. *Journal of Theoretical Biology*, 71, 473.

Brill, M.H. (1979) Further features of the illuminant-invariant trichromatic photosensor. *Journal of Theoretical Biology*, 78, 305.

Brill, M.H. & West, G. (1981) Contributions to the theory of invariance of color under the condition of varying illumination. *Journal of Mathematical Biology*, 11, 337-350

Brown, R. & MacLeod, D.I.A. (1997) Color appearance depends on the variance of surround colors. *Current Biology*, 7, 844-849.

Buchsbaum, G. (1980) A spatial processor model for object colour perception. *Journal of the Franklin Institute*, 310, 1-26.

Buchsbaum, G. & Gottschalk, A. (1984) Chromaticity coordinates of frequency-limited functions. *Journal of the Optical Society of America*, 1, 885-887.

Burnham, R.W., Evans, R.M. & Newhall, S.M. (1957) Prediction of color appearance with different adaptation illuminations. *Journal of the Optical Society of America*, 47, 35-42.

Byrne, A. & Hilbert, D.R. (1997a) *Readings on Color; Volume 1*: *The Philosophy of Color*. MIT Press, Cambridge, MA.

Byrne, A. & Hilbert, D.R. (1997b) *Readings on Color; Volume 2*: *The Science of Color*. MIT Press, Cambridge, MA.

Chichilnisky, E.J. & Wandell, B.A. (1995) Photoreceptor sensitivity changes explain color appearance shifts induced by large uniform backgrounds in dichoptic matching. *Vision Research*, 35, 239-254.

Cohen, J. (1964) Dependency of the spectral reflectance curves of the Munsell color chips. *Psychonomic Science*, l, 369.

Craven, B.J. & Foster, D.H. (1992) An operational approach to color constancy. *Vision Research*, 32, 1359-1366.

Dannemiller, J.L. (1992) Spectral reflectance of natural objects: how many basis functions are necessary? *Journal of the Optical Society of America A*, 9, 507-515.

Das, S.R. & Sastri, V.D.P. (1965) Spectral distribution and color of tropical daylight. *Journal of the Optical Society of America*, 55, 319.

Dixon, E.R. (1978) Spectral distribution of Autralian daylight. *Journal of the Optical Society of America*, 68, 437-450.

Drew, M.S. & Funt, B.V. (1990) Calculating surface reflectance using a single-bounce model of mutual reflection. *Proceedings of the Third International Conference on Computer Vision*, *Osaka, Japan,* December 4-7, 1990. IEEE Computer Society, Washington.

D'Zmura, M. (1992) Color constancy: Surface color from changing illumination. *Journal of the Optical Society of America A*, 9, 490-493.

D'Zmura, M. & Iverson, G. (1993a) Color constancy: I. Basic theory of two-stage linear recovery of spectral descriptions for lights and surfaces. *Journal of the Optical Society of America A*, 10, 2148-2165.

D'Zmura, M. & Iverson, G. (1993b) Color Constancy: II. Results for two-stage linear recovery of spectral descriptions for lights and surfaces. *Journal of the Optical Society of America A*, 10, 2166-2180.

D'Zmura, M. & Iverson, G. (1994) Color Constancy: III. General linear recovery of spectral descriptions for lights and surfaces. *Journal of the Optical Society of America A*, 11, 2389-2400.

D'Zmura, M., Iverson, G. & Singer, B. (1995) Probabilistic color constancy. In *Geometric Representations of Perceptual Phenomena; Papers in Honor of Tarow Indow on His 70th Birthday* (eds. Luce, R.D., D'Zmura, M., Hoffman, D., Iverson, G.J. & Romney, A.K.) pp. 187-202. Lawrence Erlbaum Associates, Mahwah, New Jersey.

D'Zmura, M. & Lennie, P. (1986) Mechanisms of color constancy. *Journal of the Optical Society of America A*, 3, 1662-1672.

Evans, R.M. (1948) *An Introduction to Color*. Wiley, New York.

Fairchild, M.D. & Lennie, P. (1992) Chromatic adaptation to natural and incandescent illuminants. *Vision Research*, 32, 2077-2085.

Ferguson, T.S. (1967) *Mathematical Statistics; A Decision Theoretic Approach*. Academic Press, New York.

Forsyth, D. (1990) A novel algorithm for color constancy. *International Journal of Computer Vision*, 5, 5-36.

Foster, D.H. & Nascimento, S.M.C. (1994) Relational colour constancy from invariant cone-excitation ratios. *Proceedings of the Royal Society of London B*, 257, 115-121.

Foster, D.H., Nascimento, S.M.C., Craven, B.J., Linnell, K.J., Cornelissen, F.W. & Brenner, E. (1997) Four issues concerning colour constancy and relational colour constancy. *Vision Research*, 37, 1341-1345.

Freeman, W.T. & Brainard, D.H. (1995) Bayesian decision theory, the maximum local mass estimate, and color constancy. *Proceedings of the 5th International Conference on Computer Vision*, Cambridge University Press, MA, 210-217.

Funt, B.V. & Ho, J. (1989) Color from black and white. *International Journal of Computer Vision*, 3, 109-117.

Gershon, R. & Jepson, A.D. (1988) Discounting illuminants beyond the sensor level. *Proceedings of the*

SPIE Conference on Intelligent Robots and Computer Vision VII, 1002, 250-257.

Gershon, R. & Jepson, A.D. (1989) The computation of color constant descriptors in chromatic images. *Color Research and Application*, 14, 325-334.

Gilchrist, A.L. (1977) Perceived lightness depends on perceived spatial arrangement. *Science*, 195, 185.

Gilchrist, A.L. (1980) When does perceived lightness depend on perceived spatial arrangement? *Perception & Psychophysics*, 28, 527-538.

Gilchrist, A.L. (1994) *Lightness, Brightness & Transparency*. Lawrence Erlbaum Associates, Hillsdale, New Jersey.

Gilchrist, A.L., Delman, S. & Jacobsen, A. (1983) The classification and integration of edges as critical to the perception of reflectance and illumination. *Perception & Psychophysics*, 33, 425-436.

Healey, G. (1991) Estimating spectral reflectances using highlights. *Image and Vision Computing*, 9, 333-337.

Healey, G., Shafer, S. & Wolfe, L. (1992) *Physics-Based Vision: Principles and Practice*. Jones & Bartlett, London.

von Helmholtz, H. (1896/1962) *Helmholtz's Treatise on Physiological Optics* (ed. Southall, J.P.C.). Dover, New York.

Helson, H. (1934) Some factors and implications of color constancy. *Journal of the Optical Society of America*, 33, 555-567.

Helson, H. (1938) Fundamental problems in color vision. I. The principle governing changes in hue saturation and lightness of non-selective samples in chromatic illumination. *Journal of Experimental Psychology*, 23, 439.

Helson, H. & Judd, D.B. (1936) An experimental and theoretical study of changes in surface colors under changing illuminations. *Psychological Bulletin*, 33, 740-741.

Helson, H. & Michels, W.C. (1948) The effect of chromatic adaptation on achromaticity. *Journal of the Optical Society of America*, 38, 1025-1032.

Hilbert, D.R. (1987) Color and color perception; A study in anthropocentric realism. *CSLI Lecture Notes Number 9*. Center for the Study of Language and Information, Stanford, CA.

Ho, J. (1988) *Chromatic aberration: A new tool for colour constancy.* Master's Thesis, School of Computer Science, Simon Fraser University, Vancouver, Canada.

Ho, J., Funt, B.V. & Drew, M.S. (1990) Separating a color signal into illumination and surface reflectance components: Theory and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12, 966-977.

Horn, B.K.P. & Brooks, M.J. (1989) *Shape from Shading*. The MIT Press, Cambridge, MA.

Horn, B.K.P. & Sjoberg, R.W. (1989) Calculating the reflectance map. In *Shape from Shading* (eds. Horn, B.K.P. & Brooks, M.J.) pp. 215-244. The MIT Press, Cambridge, MA.

Hunt, R.W.G. (1991) Revised colour-appearance model for related and unrelated colours. *Colour Research an Applications*, 16, 146-165.

Hurlbert, A. (1986) Formal connections between lightness algorithms. *Journal of the Optical Society of America A*, 3, 1684-1693.

Hurlbert, A. (1998) Computational models of color constancy. In *Perceptual constancies* (eds. Walsh, V. & Kulikowski, J.) Cambridge University Press, Cambridge (in Press).

Hurvich, L.M. (1981) *Color Vision*. Sinauer, Sunderland, Massachusetts.

Hurvich, L.M. & Jameson, D. (1957) An opponent process theory of color vision. *Psychological Review*, 64, 384-404.

Ingle, D. (1985) The goldfish as a retinex animal. *Science*, 227, 651-654.

Iverson, G. & D'Zmura, M. (1995a) Criteria for color constancy in trichromatic bilinear models. *Journal of the Optical Society of America A*, 11, 1970-1975.

Iverson, G. & D'Zmura, M. (1995b) Color Constancy: Spectral Recovery Using Trichromatic Bilinear Models. In *Geometric Representations of Perceptual Phenomena, Papers in Honor of Tarow Indow on His 70th Birthday* (eds. Luce, R.D., D'Zmura, M., Hoffman, D., Iverson, G.J. & Romney, A.K.)

pp. 169-185. Lawrence Erlbaum Associates, Mahwah, New Jersey.

Ives, H.E. (1912b) the relation between the color of the illuminant and the color of the illuminated object. *Transactions of the Illuminating Engineering Society*, 62-72.

Jacobs, G.H. (1981) *Comparative Color Vision*. Academic Press, New York.

Jacobs, G.H. (1990) Evolution of mechanisms for color vision. *Proceedings of the SPIE*, 1250, 287-292.

Jacobs, G.H. (1993) The distribution and nature of colour vision among the mammals. *Biological Review*, 68, 413-471.

Jameson, D. & Hurvich, L.M. (1964) Theory of brightness and color contrast in human vision. *Vision Research*, 4, 135-164.

Jameson, D. & Hurvich, L.M. (1989) Essay concerning color constancy. *Annual Review of Psychology*, 40, 1-22.

Jenness, J.W. & Shevell, S.K. (1995) Color appearance with sparse chromatic context. *Vision Research*, 35, 797-805.

Judd, D.B. (1940) Hue saturation and lightness of surface colors with chromatic illumination. *Journal of the Optical Society of America*, 30, 2.

Judd, D.B., MacAdam, D.L. & Wyszecki, G.W. (1964) Spectral distribution of typical daylight as a function of correlated color temperature. *Journal of the Optical Society of America*, 54, 1031.

Kaiser, P.K. & Boynton, R.M. (1996) *Human Color Vision, 2nd Edition*. Optical Society of America, Washington, D.C..

Klinker, G.J., Shafer, S.A. & Kanade, T. (1988) The measurement of highlight in color images. *International Journal of Computer Vision*, 2, 7-32.

Kraft, J.M. & Brainard, D.H. (1997) An analysis of cues contributing to color constancy. *Program of the Optical Society of Americal Annual Meeting*, Long Beach, CA, October 12-17, 1997, p. 110.

Krantz, D. (1968) A theory of context effects based on cross-context matching. *Journal of Mathematical Psychology*, 5, 1-48.

von Kries, J. (1902/1970) Chromatic adaptation. Selection translated and reprinted in *Sources of Color Science* (ed. MacAdam, D.L.) pp. 109-119. The MIT Press, Cambridge, MA.

von Kries, J. (1905/1970) Influence of adaptation on the effects produced by luminous stimuli, Selection translated and reprinted in *Sources of Color Science* (ed. MacAdam, D.L.) pp. 120-126. The MIT Press, Cambridge, MA.

Krinov, E.L. (1947/1953) *Spectral'naye otrazhatel'naya sposobnost'prirodnykh obrazovanii*. Izd. Akad. Nauk USSR (Proc. Acad. Sci. USSR); translated by G. Belkov, *Spectral reflectance properties of natural formations*; Technical translation: TT-439. Ottawa, Canada: National Research Council of Canada, 1953.

Land, E.H. (1959/1961) Experiments in color vision. *Scientific American*, 201, 84-99, reprinted in *Color Vision; An Enduring Problem in Psychology* (eds. Teevan, R.C. & Birney, R.C.). Van Nostrand, Toronte.

Land, E.H. (1983) Recent advances in retinex theory and some implications for cortical computations: color vision and the natural image. *Proceedings of the National Academy of Sciences*, 80, 5163-5169.

Land, E.H. (1986) Recent advances in retinex theory. *Vision Research*, 26, 7-22.

Land, E.H. & McCann, J.J. (1971) Lightness and retinex theory. *Journal of the Optical Society of America*, 61, 1-11.

Landy, M.S., Maloney, L.T., Johnston, E.B. & Young, M. (1995) Measurement and modeling of depth cue combination: In defense of weak fusion. *Vision Research*, 35, 389-412.

Lee, H.-C. (1986) Method for computing the scene-illuminant chromaticity from specular highlights. *Journal of the Optical Society of America A*, 3, 1694-1699.

Lee, H.-C., Breneman, E.J. & Schulte, C.P. (1990) Modeling light reflection for computer color vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12, 402-409.

Lythgoe, J.N. (1979) *The Ecology of Vision*. Clarendon, Oxford.

MacAdam, D.L. (1981) *Color Measurement. Theme and Variations*. Springer-Verlag, Berlin.

Maddox, I.J. (1970) *Elements of Functional Analysis*. Cambridge University Press, Cambridge.

Maloney, L.T. (1984) Computational approaches to color constancy. Dissertation: Stanford University. Reprinted as (1985) *Stanford Applied Psychology Laboratory Report* 1985-01.

Maloney, L.T. (1986) Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *Journal of the Optical Society of America A*, 3, 1673-1683.

Maloney, L.T. (1992) Color constancy and color perception: The linear-models framework. In *Attention and Performance XIV: Synergies in Experimental Psychology, Artificial Intelligence, and Cognitive Neuroscience - a Silver Jubilee* (eds. Meyey, D.E. & Kornblum, S.) pp. 59-78. MIT Press, Cambridge, MA.

Maloney, L.T. (1998) Surface spectral reflectance: Models and evaluation. In *Colour Vision: From Light to Object* (eds. Mausfeld, R. & Heyer, D.) (in Preparation).

Maloney, L.T. & Landy, M.S. (1989) A statistical framework for robust fusion of depth information. In *Visual Communications and Image Processing, IV Proceedings of the SPIE* (ed. Pearlman, W.A.) p. 1199, 1154-1163.

Maloney, L.T. & Varner, D.C. (1986) Chromatic adaptation, the control of chromatic adaptation, and color constancy (abstract). *Optics News*, 12, 134.

Maloney, L.T. & Wandell, B.A. (1986) Color constancy: A method for recovering surface spectral reflectance. *Journal of the Optical Society of America A*, 3, 29.

Mardia, K.V., Kent, J.T. & Bibby, J.M. (1979) *Multivariate Analysis*. Academic Press, London.

Marimont, D. & Wandell, B.A. (1992) Linear models of surface and illuminant spectra. *Journal of the Optical Society of America A*, 9, 1905-1913.

Mausfeld, R. (1997) Colour perception: From Grassman codes to a dual code for object and illuminant colours. In *Color Vision* (eds. Backhaus, W., Kliegl, R. & Werner, J.). De Gruyter, Berlin.

Mollon, J.D., Estévez, O. & Cavonius, C.R. (1990) The two subsystems of colour vision and their roles in wavelength discrimination. In *Vision, Coding and Efficiency* (ed. Blakemore, C.). Cambridge University Press, Cambridge.

Nayar, S.K. & Oren, M. (1995) Visual appearance of matte surfaces. *Science*, 267, 1153-1156.

Nassau, K. (1983) *The Physics and Chemistry of Color: The Fifteen Causes of Color*. Wiley, New York.

Neumeyer, C. (1981) Chromatic adaptation in the honey bee: Successive color contrast and color constancy. *Journal of Comparative Physiology*, 144, 543-553.

Oren, M. & Nayar, S.K. (1995) Generalization of the Lambertian model and implications for machine vision. *International Journal of Computer Vision*, 14, 227-251.

Parkkinen, J.P.S., Hallikainen, J. & Jaaskelainen, T. (1989) Characteristic spectra of Munsell colors. *Journal of the Optical Society of America A*, 6, 318-322.

Poirson, A.B. & Maloney, L.T. (1998) Surface color appearance in simple and complex scenes. In *Colour Vision: From Light to Object* (eds. Mausfeld, R. & Heyer, D.) (in Preparation).

Romero, J., Garciá-Beltrán, A. & Hernández-Andrés, J. (1997) Linear bases for representation of natural and artificial illuminants. *Journal of the Optical Society of America A* (in Press).

Rubner, J. & Schultern, K. (1989) A regularized approach to color constancy. *Biological Cybernetics*, 61, 29-36.

Sällström, P. (1973) Colour and physics: Some remarks concerning the physical aspects of human colour vision. University of Stockholm: Institute of Physics Report, 73-09.

Sastri, V.D.P. & Das, S.R. (1966) Spectral distribution and color of north sky at Delhi. *Journal of the Optical Society of America*, 56, 829.

Sastri, V.D.P. & Das, S.R. (1968) Typical spectra distributions and color for tropical daylight. *Journal of the Optical Society of America*, 58, 391.

Shafer, S.A. (1985) Using color to separate reflectance components. *Color Research and Applications*, 10, 210-218.

Shepard, R.N. (1992) The perceptual organization of colors: An adaptation to regularities of the terrestrial world? In *The adapted mind; Evolutionary psychology and the generation of culture* (eds. Barkow, J.H., Cosmides, L. & Tooby, J.) pp. 495-531. Oxford University Press, New York.

Shevell, S.K. (1978) The dual role of chromatic backgrounds in color perception. *Vision Research*, 18, 1649-1661.

Speigle, J.M. (1998) Testing whether a common representation mediates the effects of viewing context on color appearance, Unpublished Ph.D. Thesis, University of California, Santa Barbara.

Speigle, J.M. & Brainard, D.H. (1996) Is color constancy task independent? *Proceedings of the 4th IS&T/SID Color Imaging Conference*, 167-172.

Stiles, W.S. (1961) Adaptation, chromatic adaptation, colour transformation. *Anales Real Soc. Espan. Fis. Quim.*, Series A, 57, 149-175.

Stiles, W.S., Wyszecki, G. & Ohta, N. (1977) Counting metameric object-color stimuli using frequency-limited spectral reflectance functions. *Journal of the Optical Society of America*, 67, 779.

Strang, G. (1988) *Linear Algebra and its Applications.* Harcourt, Brace, Jovanovich, New York.

Thompson, E. (1995) *Colour Vision; A Study in Cognitive Science and the Philosophy of Perception*. Routledge, London.

Tominaga, S. & Wandell, B.A. (1989) The standard surface reflectance model and illuminant estimation. *Journal of the Optical Society of America A*, 6, 576-584.

Tominaga, S. & Wandell, B.A. (1990) Component estimation of surface spectral reflectance. *Journal of the Optical Society of America A*, 7, 312-317.

Troost, J.M. & de Weert, C.M. (1991) Naming versus matching in color constancy. *Perception & Psychophysics*, 50, 591-602.

Tsukada, M. & Ohta, Y. (1990) An approach to color constancy using multiple images. *Proceedings of the Third International Conference on Computer Vision, Vol. 3*. 385-393.

Ulrich von Strassburg (1262) De Pulchro, In Grabmann, M. (1926), *Des Ulrich Engelberti O. Pr. (1277) Abhandlung de Pulchro: Untersuchung und Texte*. München: Sitzungsberichte der Bayerischen Akademie der Wissenschaften, Phil.-hist. Klasse, Jg. 1925.

Vrhel, M.J., Gershon, R. & Iwan, L.S. (1994) Measurement and analysis of object reflectance spectra. *Color Research and Applications*, 19, 4-9.

Vrhel, M.J. & Trussel, H.J. (1992) Color correction using principal components. *Color Research and Applications*, 17, 328-338.

Walraven, J. (1976) Discounting the background: The missing link in the explanation of chromatic induction. *Vision Research*, 16, 289-295.

Wandell, B.A. (1987) The synthesis and analysis of color images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-9, 2-13.

Wandell, B.A. (1995) *Foundations of Vision*. Sinauer & Associates, Sunderland, MA.

Webster, M.A. (1996) Human colour perception and its adaptation: topical review. *Network: Computation in Neural Systems*, 7, 587-634.

Webster, M.A. & Mollon, J.D. (1995) Colour constancy influenced by contrast adaptation. *Nature*, 373, 694-698.

Weisskopf, V.F. (1968) How light interacts with matter. *Scientific American*, 219, 59-71.

Werner, A. (1990) Farbkonstanz bei der Honigbiene, Apis Mellifera. Doctoral dissertation, Fachbereich Biologie, Freie Universität Berlin.

Werner, J.S. & Walraven, J. (1982) effect of chromatic adaptation on the achromatic locus: The role of contrast, luminance, and background color. *Vision Research*, 22, 929-944.

West, G. & Brill, M.H. (1982) Necessary and sufficient conditions for von Kries chromatic adaptation to give color constancy. *Journal of Mathematical Biology*, 15, 249-258.

Wyszecki, G. & Stiles, W.S. (1982) *Color Science; Concepts and Methods, Quantitative Data and Formulas. 2nd Edition*. Wiley, New York.

Yilmaz, H. (1962) Color vision and a new approach to color perception. In *Biological Prototypes and Synthetic Systems, Vol. 1*. Plenum, New York.

Young, N. (1988) *An Introduction to Hilbert Space*. Cambridge University Press, Cambridge.

Zaidi, Q., Spehar, B. & DeBonet, J. (1997) Color constancy in variegated scenes: Role of low-level mechanisms in discounting illumination changes. *Journal of the Optical Society of America A*, 14, 2608-2621.

Zaidi, Q., Spehar, B. & DeBonet, J. (1998) Adaptation to textured chromatic fields. *Journal of the Optical Society of America A*, 15, 23-31.