



Illuminant cues in surface color perception: tests of three candidate cues

Joong Nam Yang^{a,*}, Laurence T. Maloney^{a,b}

^a Department of Psychology, New York University, New York, NY, USA

^b Center for Neural Science, New York University, New York, NY, USA

Received 17 May 2000; received in revised form 18 May 2001

Abstract

Many recent computational models of surface color perception presuppose information about illumination in scenes. The models differ primarily in the physical process each makes use of as a *cue to the illuminant*. We evaluated whether the human visual system makes use of any of three of the following candidate illuminant cues: (1) specular highlight, (2) full surface specularity [Lee, H. C. (1986). Method for computing the scene-illuminant chromaticity from specular highlights. *Journal of the Optical Society of America A*, 3(10), 1694–1699; D'Zmura, M., & Lennie, P. (1986). Mechanisms of color constancy. *Journal of the Optical Society of America A*, 3(10), 1662–1672], and (3) uniform background. Observers viewed simulated scenes binocularly in a computer-controlled Wheatstone stereoscope. All simulated scenes contained a uniform background plane perpendicular to the observer's line of sight and a small number of specular, colored spheres resting on the uniform background. Scenes were rendered under either standard illuminant D65 or standard illuminant A. Observers adjusted the color of a small, simulated test patch to appear achromatic. In a series of experiments we perturbed the illuminant color signaled by each candidate cue and looked for an influence of the changed cue on achromatic settings. We found that the specular highlight cue had a significant influence, but that the influence was asymmetric: greater when the base illuminant, CIE standard Illuminant A, was perturbed in the direction of Illuminant D65 than vice versa. Neither the full surface specularity cue nor the background cue had any observable influence. The lack of influence of the background cue is likely due to the placement of the test patch in front of the background rather than, as is typical, embedded in the background. © 2001 Elsevier Science Ltd. All rights reserved.

Keywords: Color; Color perception; Surface color; Color constancy; Illuminant cue

1. Introduction

[I]n our observations with the sense of vision, we always start out by forming a judgment about the colors of bodies, eliminating the differences of illumination by which a body is revealed to us.

von Helmholtz (1896/1962, Helmholtz's treatise on physiological optics, p. 287).

In this remarkable sentence, von Helmholtz (von Helmholtz, 1896/1962) proposes a theory of surface color perception: bodies have intrinsic surface colors,

and, while the initial visual information available to biological systems confounds light and surface, the visual system manages to arrive at surface color estimates that are invariant under changes in illumination, that depend only on the intrinsic properties of surfaces. Now, over a century later, we might want to qualify every part of the statement above. First of all, the degree of surface color constancy that we experience depends on viewing conditions: under some circumstances we have essentially no color constancy (Helson & Judd, 1936) and under others we show a remarkable, nearly perfect, degree of constancy (Brainard, Brunt, & Spiegle, 1997; Brainard, 1998). Von Helmholtz's assertion can only apply to the latter sort of viewing conditions.

A mathematical analysis of how surfaces and light interact and how spectral information is encoded in the retina leads to the provisional conclusion that von

* Corresponding author. Present address: Visual Science Laboratories, University of Chicago, 940 East 57th Street, Chicago, IL 60637, USA. Tel.: +1-773-702-9412; fax: +1-773-702-0939.

E-mail address: jyang2@uchicago.edu (J.N. Yang).

Helmholtz posed an impossible task for biological vision. The color signal that comes to the eye has two components, illuminant and surface reflectance, and the data available to the visual system are simply the excitations of photoreceptors at each location xy in the retina:

$$\rho_k^{xy} = \int E(\lambda) S^{xy}(\lambda) R_k(\lambda) d\lambda, \quad k = 1, 2, 3 \quad (1)$$

Here, $S^{xy}(\lambda)$ is used to denote the *surface spectral reflectance function* of a surface patch imaged on retinal location xy , $E(\lambda)$ is the *spectral power distribution* of the light incident on the surface patch, and $R_k(\lambda)$, $k = 1, 2, 3$ are the photoreceptor sensitivities, all indexed by wavelength λ in the electromagnetic spectrum.¹ The visual system is assumed to contain photoreceptors with three distinct sensitivities ($k = 1, 2, 3$), although, of course, at most one photoreceptor can be present at a single retinal location. $E(\lambda)$ and $S^{xy}(\lambda)$ are, in general, unknown, while the $R_k(\lambda)$, $k = 1, 2, 3$ are taken to be known. Any visual system that is color constant (Fig. 1) must effectively invert Eq. (1), transforming photoreceptor excitations into non-trivial surface color descriptors that depend only on $S^{xy}(\lambda)$. Yet, without further constraints on the problem, Eq. (1) cannot be inverted in this way, and the




Fig. 1. A simplified model of surface color perception. $S^{xy}(\lambda)$ is used to denote the surface spectral reflectance function of a surface patch imaged on retinal location xy , $E(\lambda)$ is the spectral power distribution of the light incident on the surface patch, and $R_k^{xy}(\lambda)$, $k = 1, 2, 3$ are photoreceptor sensitivities, all indexed by wavelength λ in the electromagnetic spectrum. Light is absorbed and re-emitted by the surface toward the eye where the retinal image is sampled spatially and spectrally. Under conditions of color constancy, the eventual perception of surface color must be determined primarily by the surface $S^{xy}(\lambda)$ and not by the illuminant $E(\lambda)$.

¹ Eq. (1) is a simplification of the physics of light-surface interaction. We are ignoring the effects of changes in the positions of light sources, the location and surface orientation of the surface patch, and the position of the eye. See Maloney (1999).

problem cannot be solved, even approximately (Ives, 1912).

1.1. Environments and algorithms

How, then, is color constancy, approximate or exact, ever possible for a visual system like ours? In the last 20 years, a number of researchers have attempted to develop models of biologically plausible, color constant visual systems (for reviews, see Hurlbert (1998) and Maloney (1999)). For our purposes, we can think of each model as comprising (1) a mathematical description of an idealized world (referred to as an *environment* by Maloney (1999)) and (2) an *algorithm* that can be used to compute invariant surface color descriptors within the specified environment. The statement of the environment comprises the constraints that make it possible to invert Eq. (1), and the algorithm is a recipe for doing just that.

Once removed from its environment, an algorithm may fail partially or completely (as we noted above, human color constancy also fails dramatically under some viewing conditions). An active area of research concerns the match or lack of match between mathematically-described environments, and particular subsets of the terrestrial environment where we suspect that human surface color perception is constant or nearly so (Maloney, 1986; Parkkinen, Hallikainen, & Jaaskelainen, 1989; van Hateren, 1993; Vrhel, Gershon, & Iwan, 1994; Romero, Garcia-Beltran, & Hernandez-Andres, 1997; Bonnardel & Maloney, 2000) (for a review, see Maloney (1999)).

Many recent algorithms have a common structure: first,² information concerning the illuminant spectral power distribution is estimated. This information is usually equivalent to knowing how photoreceptors would respond if directly stimulated by the illuminant without an intervening surface (Maloney, 1999). This illuminant estimate is then used in inverting Eq. (1) to obtain invariant surface color descriptors, typically by using the method of Buchsbaum (1980). The algorithms differ from one another primarily in how they get information about the illumination: there are currently algorithms that make use of surface specularities (Lee, 1986; D'Zmura & Lennie, 1986), shadows (D'Zmura, 1992), mutual illumination (Funt, Drew, & Ho, 1991), reference surfaces (Brill, 1978; Buchsbaum, 1980), subspace constraints (Maloney & Wandell, 1986; D'Zmura & Iverson, 1993), scene averages (Buchsbaum, 1980), and more. It is evident that there are many potential *cues to the illuminant* in everyday, three-dimensional scenes.

² Some of the algorithms compute estimates of illuminant information and surface color descriptors cooperatively, rather than successively. We describe the computation as sequential ('first the illuminant, then the surfaces') for convenience in presentation.

1.2. Three candidate illuminant cues

The question addressed in this article is: Do biological visual systems make use of any of the illuminant cues proposed in the computational literature? In this article we will examine three candidate cues, and test whether information about the illuminant encoded in any of the three influences surface color perception. The first two cues make use of surface specularities as, in effect, a mirror that can be used to view the illuminant.

The first cue, *specular highlight*, uses the photoreceptor excitations corresponding to one or more neutral specular highlights in a scene as an estimate of the photoreceptor excitations to be expected when the visual system directly views the illuminant:

$$\rho_k^E = \int E(\lambda) R(\lambda) d\lambda, \quad k = 1, 2, 3. \quad (2)$$

We will refer to these excitations as the *chromaticity* of the illuminant.

In order to use this cue, a visual system must decide which parts of a scene count as specular highlights and which do not. This *highlight classification task* (correctly detecting specular highlights) is similar in many respects to the problem of correctly classifying light sources in a scene (see Ullman, 1976) and has no obvious solution. Suppose, for example, that a visual system picked the ‘brightest point’ (by some definition of ‘brightness’) in the scene or in a region of the scene and classified it as a specular highlight. This ‘bright point’ may correspond to a distant light source that does not contribute significantly to the illumination of the scene; it may be a reflection in a surface (gold, copper) that is not spectrally neutral; or the ‘bright point’ may signal the correct chromaticity of the illuminant for one part of the scene, but not for another. Ullman (1976) noted, in a sort of converse, that a true light source need not be the ‘brightest’ point in a scene when there is a sharp illuminant gradient across space. The same may be said of specular highlights. Any of these misclassifications could lead to significant errors in estimating illuminant chromaticity by means of the specular highlight cue.

If the human visual system makes use of the specular highlight cue at all, it is of interest to examine whether it can discriminate between true and false highlights in scenes and ignore the latter. We will return to this point in the discussion.

The second cue considered, *full-surface specularity*, was independently proposed by Lee (1986) and by D’Zmura and Lennie (1986). It makes use of surface specularity information concerning the illuminant, but does not restrict attention to specular highlights or require that specular highlights be perfect mirrors reflecting the illuminant. One of the environmental assumptions underlying the Lee–D’Zmura–Lennie cue

is that the spectral characteristics of surfaces are accurately described by a model due to Shafer (Cook & Torrance, 1982; Shafer, 1985). In the Shafer model, a surface reflectance is a superposition of an idealized matte surface (‘Lambertian’) and a neutral mirror (‘specular’):

$$S(\lambda) = \alpha S^*(\lambda) + \beta \quad (3)$$

where α and β are non-negative ‘geometric’ scale factors that vary with the relative position of the light source and the eye and $S^*(\lambda)$ is the surface spectral reflectance function of the Lambertian surface for some fixed choice of viewing geometry. The geometric scale factors are further constrained so that $S(\lambda)$ is a valid surface reflectance function with values between 0 and 1 inclusive at every wavelength. When α is large relative to β , the surface will look like a piece of colored blotting paper, when β is large relative to α , the surface will look like a mirror.

The key idea in the algorithm proposed by Lee and D’Zmura–Lennie is that, for any extended surface under near-punctate illumination, α and β will naturally vary as the angles from the eye and from the light source to different points on the surface patch vary. This variation is enough to allow estimation of the contribution of the specular component uncontaminated by the Lambertian component by, in effect, constructing a virtual mirror in which the eye may view the illuminant. Fig. 2 illustrates the key idea of the algorithm.

The full-surface specularity cue is available even for objects that are only slightly specular, such as human skin. The most specular point on a face may still be an




Fig. 2. The Lee–D’Zmura–Lennie model. The LMS coordinates for all of the points on a homogeneous Lambertian-specular surface fall on a plane that contains the LMS coordinates of the Lambertian component (‘Matte’) and the LMS values of the specular component. (‘Specular’). The intersection of the planes corresponding to two distinct surfaces is a line containing the LMS coordinates of the specular component which is then determined up to an unknown scaling factor.

evident mixture of the color of the illuminant and the color of the underlying matte component of the face. The Lee–D’Zmura–Lennie approach can be used to estimate photoreceptor excitations corresponding to the illuminant in conditions where the specular highlight cue would give a seriously misleading estimate. One peculiarity of the Lee–D’Zmura–Lennie algorithm is that there must be at least two surfaces available in the scene with Lambertian surface reflectance functions that are distinct (specifically, not proportional). In a scene with many Shafer objects, all with the same ‘color’ (Lambertian surface reflectance function), the Lee–D’Zmura–Lennie cue is not available. We will use this fact in designing the experiments reported below.

The third cue, the color of a *uniform background*, has been extensively studied in scenes containing little else besides the uniform background and a test patch. In such scenes, the photoreceptor excitations of the background have an evident effect on the apparent color of the test patch (see Whittle, 1973). When the background is no longer homogeneous (the Mondrian stimuli of Land & McCann (1971)) or other objects are placed in the scene, it is less clear that the background has much effect, if any, on perceived color. Helson (Helson, 1938, 1943) used the term *adaptation reflectance*, which was meant to be a weighted average of the surfaces in the scene, and advanced the hypothesis that the color of the test patch in such a complex scene would be the color seen in a scene where the test patch was surrounded by a uniform background that was everywhere set to the average, adaptation reflectance. His hypothesis is sometimes referred to as the *equivalent background hypothesis* or *gray world hypothesis*. The later Retinex algorithms (Land, 1983, 1986) effectively computed equivalent backgrounds for local regions of any scene, but used the geometric rather than the arithmetic average of the region (see discussion in Brainard & Wandell (1986)).

We consider the hypothesis that the photoreceptor excitations of a uniform background in a scene are interpreted as an estimate of the illuminant. The reader should be aware that, when the background surface is not close to achromatic, that this cue is a remarkably bad one. Thus, while psychophysical evidence indicates that it is used when it is arguably the only cue available (e.g. Whittle, 1992), we might expect that, in the presence of other cues, its influence may vanish.

1.3. Cue perturbation methods

We wanted to determine whether the visual system ever makes use of any of the candidate cues to the illuminant just described, and we set out to do so using the cue perturbation approach that Maloney and Landy (Maloney & Landy 1989; Landy, Maloney,

Johnston, & Young, 1995) applied to depth and shape vision: we first simulated binocular scenes where multiple candidate cues to the illuminant are available. We next measured the observer’s achromatic setting (described in Section 2) for a small test surface within the scene when the scene was illuminated under illuminant I_1 . We repeated this measurement under a second standard illuminant I_2 . The two achromatic settings, one for each of the illuminants, are plotted in a standard color space as shown in Fig. 3A (Lu’v’ space, see Benzschawel (1992)). The direction and magnitude of any observed change in achromatic setting, in response to changes in the illuminant, are useful measures of the observer’s degree of color constancy and whether the visual system discounted the change in illumination. However, so far, we can conclude nothing about the *relative* importance of any of the illuminant cues present, since all signal precisely the same illuminant in both rendered scenes.

We next ask the observer to make a third achromatic setting in the scene where the illuminant information for one cue is set to signal Illuminant I_2 , while all other cues are set to signal Illuminant I_1 . This sort of cue manipulation is not difficult with simulated scenes, but would be very difficult to do in a real scene. The experimental data we now have comprises three achromatic settings: under Illuminant I_1 , under Illuminant I_2 , and under Illuminant I_1 with one cue perturbed to signal Illuminant I_2 . We wish to determine whether the visual system is ‘paying attention’ to the perturbed cue, whether the perturbed cue has a measurable influence on color perception measured by achromatic adjustment.

What might happen? One possibility is that the observer’s setting in the scene with one cue perturbed to signal Illuminant I_2 is identical to the setting that he or she chose when all cues signaled Illuminant I_1 (point α in Fig. 3A). We would conclude that the perturbed cue had no effect whatsoever on surface color perception – it is not a cue to the illuminant, at least in the scene we are considering (see Fig. 3A).

Suppose, on the other hand, the observer’s achromatic setting in the scene with one cue perturbed to signal Illuminant I_2 (and all others are set to signal Illuminant I_1) is the same as it was when all cues signaled Illuminant I_2 (point β in Fig. 3A). This would suggest that the observer is *only* using the manipulated cue, ignoring the others.

A third possibility is that the observer chooses a setting somewhere between his or her settings for the two illuminants (see Fig. 3A), along the line joining them (point γ in Fig. 3A). Let δ be the change in setting when only the perturbed cue signals Illuminant I_1 and let Δ be the change in setting when all cues signal Illuminant I_2 (i.e. the illuminant is Illuminant I_2 and no




Fig. 3. Hypothetical data from a cue perturbation experiment. Each shape is a hypothetical achromatic setting for a single observer viewing a single scene under different illumination conditions. The white circle is the setting when the scene is illuminated by Illuminant I_1 , the black circle, by Illuminant I_2 . The diamonds correspond to the hypothetical achromatic settings when one illuminant cue signals I_2 and the remainder signal I_1 . (A) *Noise- and distortion-free settings*. Point α : the setting coincides with the settings for the unperturbed cues, indicating that the perturbed cue has no influence. Point β : the setting coincides with the setting for the perturbed cue, indicating that the perturbed cue determines the setting. Point γ : The setting falls halfway between the settings for the two illuminants, indicating that the influence of the perturbed cue is 0.5. See text. (B) *Realistic data*. The observer can place his achromatic setting anywhere in color space. The three settings (for I_1 , for I_2 , and the perturbed cue setting) need not be collinear and, in general, will not be. In computing the influence of an illuminant cue, we use only the magnitude of the projection onto the line joining the unperturbed settings. See text.

cues are perturbed). We define the *influence* of the perturbed cue to be:

$$I = \frac{\delta}{\Delta} = \frac{\|\gamma - I_1\|}{\|I_2 - I_1\|}. \quad (4)$$

The value I should fall between 0 and 1. A value of 0 implies that the perturbed cue is not used, a value of 1 implies that only the perturbed cue is used.

Of course, the idealized results shown in Fig. 3A are not what we expect to obtain experimentally. In the perturbed scenes, the observer is free to make achromatic settings that do not fall on the line joining the settings in the two unperturbed scenes.³ We expect such an outcome, if only as a consequence of measurement errors. The computation of influence we actually employ is illustrated in Fig. 3B. The value δ in Eq. (4) is taken to be the length of the *projection* of the observer's setting on the line passing through the two unperturbed settings. This is the value we report below. A final technical point: the idealized definition of influence in Eq. (4) is invariant under non-singular linear transformations of color space. The length of the projection in Fig. 3B, however, varies with the choice of color space. We use CIE Lu'v' space in computing influence.

³ We emphasize that observers are permitted two degrees of freedom in their achromatic settings: only luminance is held constant.

1.4. Realism

An evidently critical factor in studies of surface color perception using computer graphics is that the images that are displayed on a computer monitor must be rendered correctly. Human color constancy with simulated images (quantified by a commonly-used color constancy index defined later on) is markedly less than that obtained with real scenes (Arend & Reeves, 1986; Arend, Reeves, Schirillo, & Goldstein, 1991; Brainard, 1998; Kuriki & Uchikawa, 1996). With real scenes, the index reaches an average of 0.84 (Brainard, 1998) while typical results with rendered scenes lead to values of 0.5 or less. We have taken several steps to ensure that the scenes we present are rendered accurately. All stimuli are presented binocularly with correct rendering of disparity cues, and as described in the Appendix, we use a special rendering method, Step-Function Rendering, to ensure that spectral information is not distorted by the rendering process. The stereo image pairs have relatively high resolution, i.e. 500×500 in pixels, on a 1024×860 screen. We have chosen stimuli so that the resulting image pairs do not exceed the contrast range of the computer monitors we use.

2. General methods

2.1. Apparatus

The observer viewed a large, high-resolution stereoscopic display. The viewing area was a box, 1.24 m on each side, with one side open, as shown in Fig. 4. The interior of the box was lined with black pressure-sensitive flocked paper (Edmund Scientific, catalogue number CR70-621). The observer sat at the open side, positioned in a chin rest, gazing into the box. Two identical Hitachi Superscan 17 in. display screens were located to either side of the observer. Small mirrors directly in front of the observer's eyes reflected images of the left and right display screens to the observer's left and right eyes, respectively. The observer was able to fuse the left and right components of the stereoscopic stimuli displayed on the screens without difficulty.

Three computers were used to control the display of stereoscopic stimuli. A control program on the control computer (Gateway Pentium II PC) selected stereo image pairs on each trial and transmitted them to two Image Computers (TriStar 486 PC Computers). The two image computers contained SVGA graphics cards (Mach32) that were used to display the left and right images of a stereo pair on the left and right display screens, respectively. All software was written in the C programming language and used X Windows (Version 11R6) to control transmission and display through an




Fig. 4. The experimental apparatus. The observer viewed three-dimensional scenes presented in a stereoscopic display controlled by three computers. The apparent location of the scene is marked *fused image*. See text for details.

isolated local network linking only the control and image computers. Loading the image pairs at the start of an experimental session required about 3 s. Once an image was displayed, the observer pressed keys that altered the color of the test patch as described below. Each image of a stereo image pair occupied a display area of 500×500 pixels (5.6×5.6 cm, $20^\circ \times 20^\circ$ of visual angle) at the center of one of the display screens. The overall screen resolution was set to 1024×860 , and each display area occupied slightly more than a quarter of the display screen area.

2.2. Color calibration

The physical intensity of light issuing from pixels on the display screens was a non-linear function of pixel values. We measured this non-linear relation using a Minolta luminance meter (LS-100) and corrected it in software (' γ correction'). The left and right screens were calibrated separately. Since the display area occupied a substantial part of each screen, we tested for possible spatial inhomogeneities. We performed separate measurements at five square regions (each 2×2 cm) of each monitor at the center and four corners of the display area. The measured differences in gun intensities between the center region and the four corners was less than 5% and, accordingly, we decided to apply the same gamma correction at all points in the region of screen used for the stimulus. The maximum luminance for each screen alone was 98 cd/m^2 . The range of luminance in the images used in all experiments was $15\text{--}90 \text{ cd/m}^2$ per screen and the test patch was always held constant at 20 cd/m^2 .

2.3. Spatial layout of the stimuli

The top stereogram in Fig. 5 shows spatial layout of the stimuli used in all experiments. There were objects and surfaces lit by an illuminant from the upper right corner. The actual matte surface colors of all of the small, specular spheres and the background used in each experiment are described for each experiment separately.

2.4. Stimulus preparation: rendering

We used the physics-based rendering package RADIANCE (Larson & Shakespeare, 1997) to render each of the images in a stereo pair. We used the RADIANCE language to specify the layout of a simple scene in space and a lighting model for the scene. The only difference between rendering computations for the two images in a stereo pair was a change in simulated viewpoint: the viewpoint for the left image corresponded to the position of the left eye of the observer in the simulated scene, that of the right image to that of the right eye. The objects within the scene were rendered as if they were, on average, the same distance in front of the observer as the optical distance from each of the observer's eyes to the corresponding display screen. This choice of location minimizes any conflict between accommodation cues and other depth cues. The rendered scene, viewed binocularly, appeared to be floating approximately 70 cm in front of the observer. In each rendered scene, there were spheres randomly placed on a uniform background plane perpendicular to the observer's Cyclopean line of sight.

RADIANCE can be used to simulate light-surface interactions based on the Shafer (Lambertian-specular)

Model discussed above. The RADIANCE rendering package permits us to change the relative balance of matte and specular components for each rendered surface, by changing the coefficient, β in Eq. (3). The β coefficients were set to 0.1 and 0.05, for the spheres and background, respectively.




Fig. 5. Stimulus layout. The same three-dimensional scene was rendered under four different illuminant conditions. Each left and right image formed a stereo pair for crossed fusion. The scene consisted of a number of specular spheres resting on a uniform plane perpendicular to the observer's Cyclopean line of sight. The small square visible in each stereo image was the test patch whose color the observer adjusted until it appeared achromatic. (a) the scene under illuminant D65, (b) the scene under illuminant D65 with the specular cue perturbed to illuminant A, (c) the scene under illuminant A with the specular cue perturbed to illuminant D65, (d) the scene under illuminant A.

The Shafer model is an accurate description of many surfaces found in everyday environments (Tominaga & Wandell, 1989) but not all (Lee, Breneman, & Schulte, 1990). Since the Lee–D'Zmura–Lennie algorithm presupposes the Shafer model, we adopted it as a description of the angular dependence of light-surface interactions.

It is important to realize that RADIANCE and, so far as we can determine, all available rendering packages *do not render color correctly*: the errors encountered are small but visually evident. These packages cannot be used for surface color perception experiments without modification. We developed an efficient way to re-interpret the input and output of any rendering package that allowed us to correctly simulate the spectral changes that light undergoes when it is absorbed and re-emitted by colored Lambertian-specular surfaces. The problem and its solution (*step-function rendering*) are outlined in the Appendix A. In using step-function rendering, we can, in effect, specify not simply the RGB color of the simulated surface but a very good approximation to its exact surface reflectance function. Similarly, we specify not the illuminant RGB color, but a very good approximation to its spectral power distribution.

The matte component of each rendered surface (background, spheres) was rendered so as to match it to a particular Munsell color reference chip from the Nickerson-Munsell collection (Kelley, Gibson, & Nickerson, 1943). We chose to use surfaces drawn from the Munsell collection in order to permit our results to be compared to earlier results in the literature that used real or simulated Munsell chips (e.g. Arend & Reeves, 1986; Arend et al., 1991). We return to this point in Section 7.

The entire scene was illuminated by a combination of a punctate and a diffuse light. The spectral power distributions of both the punctate and diffuse lights were always the same and set to be that of either standard illuminant D65 or standard illuminant A taken from Wyszecki and Stiles (1982). The punctate illuminant was always positioned 1.5 m behind the observer, above and to the right. The square test patch (0.5° of visual angle on a side) was positioned in depth in the scene, in the plane tangent to the front surface of one of the spheres. In pilot studies we learned that, when the test patch was embedded in the uniform background plane of the image, we found no measurable effect of any cues except the uniform background. When displaced out of the plane (in binocular disparity), we found the pattern of results reported in the experiments below. We return to this point in Section 7 below.

Note that, in Fig. 5a, each sphere and even the background exhibit a wide range of chromaticities in both of the stereo images, even though each is 'made' of a single surface material. The stimulus can be de-

scribed parsimoniously in terms of surfaces and illuminants, but the resulting pair of retinal images is much more complex. Even the color signals corresponding to the ‘uniform’ background differ markedly as a consequence of the relative position of the observer and the punctate light source, and the shadows of the spheres. We return to this point in the discussion as well.

2.5. Perturbing specularity cues

The specular cues are perturbed as follows. We can take any scene description and remove the specular component of all surfaces by adjusting the matte-specular parameters in Eq. (3). We can also do the opposite, creating a scene with the same geometric layout as our scene, but with purely specular surfaces. We render the matte and specular versions of the scene separately and then blend them by a weighted mixture of the two resulting rendered images. We can, of course, render the matte and specular components under the same or under distinct illuminants and, in this way, create perturbed and unperturbed versions of the same scene. Note that we do not compute the effect of light scattering from the specular component to the matte and vice versa. For the isolated spheres in our scene, these inter-reflections have negligible effect but could become significant in scenes with large specularities and closely spaced surfaces.

2.6. Observers

Five naïve paid observers, as well as the first author, participated in the experiments. The color vision of all the observers was tested with the H-R-R Pseudoisochromatic plates (Hardy, Rand, & Rittler, 1957) and all fell within the normal range.

2.7. The task

The observer sat on the open side of the apparatus and viewed the fused image through the mirrors and adjusted the color of the test patch in the image color until it appeared achromatic. Subjects tapped one pair of keys to adjust the test patch, in the $L-M$ direction and a second pair to adjust it in the $S-(L+M)$ direction (MacLeod & Boynton, 1979; Krauskopf, Williams, & Heeley, 1982; Derrington, Krauskopf, & Lennie, 1984). The luminance of the test patch was held constant and the observer was adapted to darkness for one minute before each session started. One session consisted of 20 trials and in each trial the observer made an achromatic adjustment for a binocular image pair. The observer was told to freely move his/her eyes across the image during each trial. The observer first practiced the achromatic setting task for two sessions with stimuli that contained only a uniform background and test patch (no spheres).

On each trial the initial color of the test patch was set at random to one of five possible starting points equally distant from the D65 locus in CIE chromaticity space. The choice of starting point had no measurable effect on observers’ final settings, as has been shown in other studies (Brainard, 1998).

3. Experiment 1

3.1. Purpose

This experiment was designed to test whether the visual system makes use of information about the illuminant available from specular cues in the image. We measured the influence of this cue in scenes illuminated by Illuminant D65 and in scenes illuminated by Illuminant A as described above. Each scene contained 11 highly specular spheres and a background with a small specular component.⁴ Since there are two objects with distinct matte components present in the scene (the background and any of the spheres), it is theoretically possible to gain information about the illuminant by means of either the full surface specularity cue or the specular highlight cue. We are consequently testing, in this experiment, whether the human visual system uses *any* specular cue to the illuminant.

3.2. Stimuli

The stereo-pair stimuli used in this experiment were shown in Fig. 5. The test patch (for achromatic setting) was on one of the objects’ surface. Stereograms (a) and (d) in Fig. 5 were rendered under Illuminants D65 and A; stereogram (b) was rendered under Illuminant D65 with the specularity cue perturbed toward Illuminant A and stereogram (c) under Illuminant A with the specularity cue perturbed toward Illuminant D65. The matte component of the each of the spheres was matched to the Munsell chip with coordinates BG 5/4 and the matte component of the background to the Munsell chip with coordinate N 3/ (Kelley et al., 1943).

3.3. Results and discussion

Fig. 6A shows the achromatic settings for those images for four observers. The horizontal and vertical axes are the u' and v' coordinates of the CIE chromaticity diagram. The two circles represent mean achromatic settings when the scene was rendered under Illuminant A (black) and Illuminant D65 (white). The mean achro-

⁴ The β value in Eq. (3) was set to 0.05 for the background, 0.1 for the spheres. We hesitated to completely eliminate the specular component in the background since the apparent three-dimensionality of the scene was also diminished when we did.

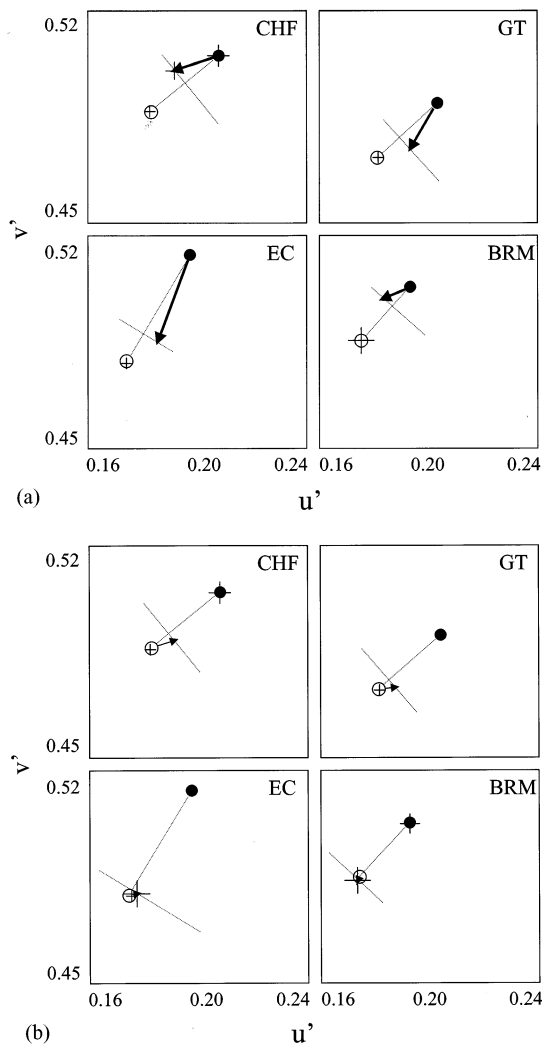


Fig. 6. The specular highlight cue: results of experiment 1. The achromatic settings for four observers are shown, plotted in the $u'v'$ coordinates in CIE chromaticity space. In each small plot, a white circle marks the mean of multiple settings by one observer for the Illuminant D65 consistent-cue condition, a black circle marks the mean for multiple settings by the same observer for the Illuminant A consistent-cue condition, and the head of the arrow marks the mean of multiple settings for the perturbed-cue condition. The base of the vector is connected to the consistent cue setting corresponding to the illuminant signaled by the non-perturbed cues. Horizontal and vertical bars indicate one SE for each setting. The projection of the perturbed setting onto the line joining the unperturbed settings is marked. For all observers, the perturbation from A to D65 led to a strong measured influence on the achromatic settings, while the perturbation from D65 to A led to little or none. (A) The perturbed cue signaled D65, all others, A. (B) The perturbed cue signaled A, all others, D65.

matic setting for the perturbed condition (when the specular highlight illuminant cue alone signaled Illuminant D65, all other cues signaling Illuminant A) is plotted as a vector (precisely, the setting is the center of the triangular head of the vector). S.D.s for each setting are shown as vertical and horizontal bars at the center of each shape. Fig. 6B shows the effect of perturbing

the specular highlight cue toward A, when all of the other illuminant cues signal D65.

Three points can be made about the results. First, the observers' achromatic settings for the two consistent images are clearly different. The observer is responding to changes in the illuminant. The observed changes are similar to those found in previous studies (Arend et al., 1991; Brainard, 1998). A quantitative comparison will be presented later in Section 7.6. Second, the setting points for the perturbed cue fall near the line joining the setting points for the two unperturbed scenes. The perturbed cue has an influence of approximately 0.3 to 0.83, as defined in Eq. (4). Third, the influence is asymmetric, in that the cue perturbation in the direction of Illuminant D65 has a much greater influence than that in the direction of Illuminant A. If all cues in the scene but specularly signal Illuminant A, and specularly signals D65, illumination estimation was much affected, whereas if all cues but specularly signal D65 and specularly signals Illuminant A, illumination estimation was less affected. These results indicate that specularities can be a useful cue to the illuminant, in some scenes, and also that Illuminant D65 and Illuminant A are treated differently by the visual system. This asymmetry is further discussed in Section 7.

We repeated experiment 1 with different choices of Munsell surface for the objects and the background. (10GY 5/6 for the objects and 10P 4/6 for the background). Fig. 7 show, first of all, that when the colors of the objects and background were changed, the achromatic settings changed little, consistent with results reported in previous studies (Brainard, 1998; Kuriki & Uchikawa, 1998). Second, the magnitudes of the influence measures were little affected, and there is still a marked asymmetry in influence between the two perturbation conditions.

4. Experiment 2

4.1. Purpose

We repeated experiment 1, varying the number of specular objects to determine how influence varies with number of specular highlights.

4.2. Stimuli

The stimuli for experiment 2 were identical to those employed in experiment 1 except that the number of objects in the image was varied from 1 to 11.

4.3. Results and discussion

Fig. 8 summarizes the results for the number of objects in the scene. Hurlbert (1989) found little influ-

ence of perturbation when there was only one large ball available in the scene; the results in Fig. 8 show that that is exactly what we observed. This lack of influence was still observed until the number of objects was six. The effect, however, began to show up with nine objects. The overall plot of influence versus number of identical specular objects is evidently non-linear. When the number of objects varied, it changed the scenes. Accordingly, the achromatic settings did change, but not in proportion to the increase in the number of objects. We will return to this result in Section 7.

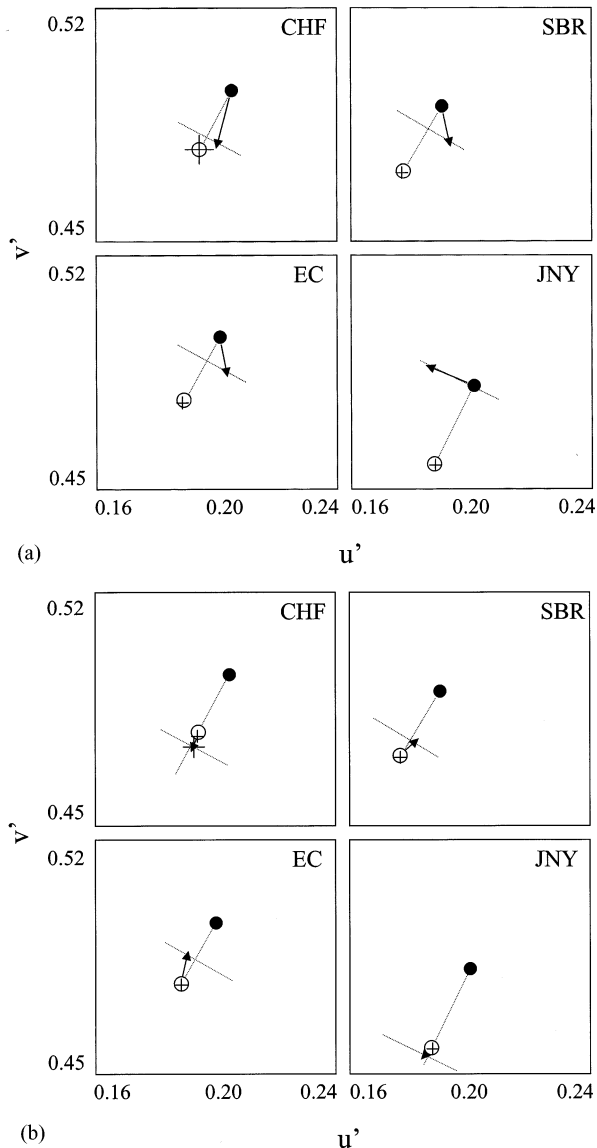


Fig. 7. The specular highlight cue: replication of experiment 1 (with different surface colors). We repeated experiment 1 assigning different Munsell surface reflectance function to surfaces in the scene, changing their apparent colors. The data presentation format is identical to that of Fig. 6. The results here are qualitatively similar to those of experiment 1. (A) The perturbed cue signaled D65, all others, A. (B) The perturbed cue signaled A, all others, D65.




Fig. 8. Influence versus the number of specular objects. We varied the number of objects in the scenes of experiment 1 and measured the influence of the specular cue. Influence is plotted versus number of objects. Measured influence was not measurably different from 0 for one to six objects, but clearly non-zero for nine or 11. Different shapes correspond to different observers. A heavy line joins the means of the observers' influence measures.

5. Experiment 3

5.1. Purpose

This experiment tests whether the full surface specular cue (Lee, 1986; D'Zmura & Lennie, 1986) is used in human vision. Recall that the full surface specular cue is only available in scenes where there are at least two homogeneous surfaces with distinct Lambertian surface spectral reflectance functions and appreciable specularity. The full surface specular cue was available in the simulated scenes of experiments 1 and 2 (the background and any one sphere count as the two objects needed) but, since the specularity of the background was low, the conditions for use of the cue were not optimal. This experiment is identical to experiment 1 except that the 11 specular spheres which all shared a common Lambertian component in experiment 1, now have 11 *distinct* Lambertian surface spectral reflectance functions. There are now multiple, highly specular objects with distinct Lambertian components and, if observers had been using the full surface specular cue in experiments 1 and 2, we might expect a noticeable increase in the influence of specular cues.

5.2. Specific methods: stimuli

The spatial layout of the stereo pairs used resembled that shown in Fig. 5. The matte component of the spheres was matched to the Munsell chips with coordinates BG 2/2, Y 7/10, Y 2/2, P 4/6, RP 2/2, 10R 5/10, PB 2/2, YR 2/2, BG 5/4, R 3/4, and 10GY 5/6 and the

matte component of the background to the Munsell chip with coordinate N 3/ (Kelley et al., 1943).

5.3. Results and discussion

The data are plotted in the CIE chromaticity diagram in $u'v'$ coordinates. Fig. 9A and 9B shows the results of achromatic settings for those images for three observers. The horizontal and vertical axes are the u' and v' coordinates of the CIE chromaticity diagram. Perturbing the specular cues in the direction of either illuminant had little effect on achromatic setting. Statis-

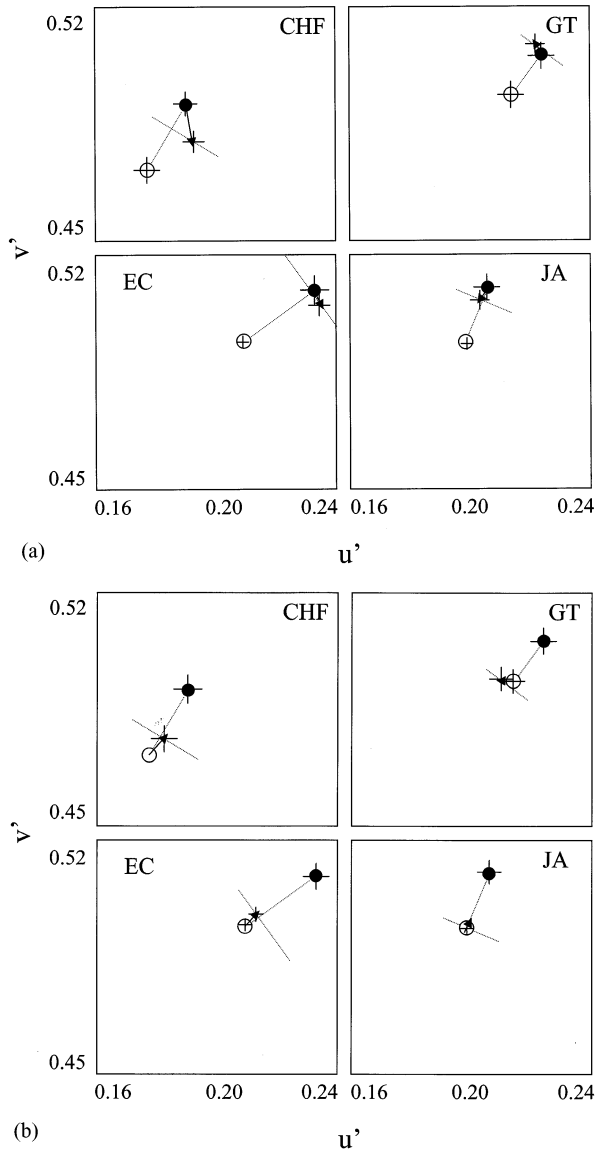


Fig. 9. A test of the full surface specularity cue (Lee-D'Zmura-Lennie). In experiment 3, we tested whether the full surface specularity cue exerts significant influence on achromatic settings. The data presentation format is identical to that of Fig. 6. For all observers, the perturbation from A to D65 or from D65 to A led to little influence on the achromatic settings. (A) The perturbed cue signaled D65, all others, A. (B) The perturbed cue signaled A, all others, D65.

tical tests show no significant influence for any observer.

We conclude that the visual system failed to use the full-surface specularity cue in scenes that would seem to make it maximally available. Moreover, counter to our expectations, the influence of the specular highlight cue is no longer present in a scene with 11 distinct spheres while it proved to be a strong cue in the scene of experiment 1 with 11 identical spheres. We postpone discussion of this result until Section 7.

6. Experiment 4

6.1. Purpose

This experiment was designed to test whether the uniform background cue is used by the visual system in its estimation of the illuminant. As in experiment 1, the scene comprised a uniform background surface, perpendicular to the line of sight and 11 small specular spheres, placed at random, tangent to the surface.

Since the scene was illuminated from upper left, the background cue in the image has evident luminance and chromatic gradients. As discussed above, we define the estimate of the illuminant available from the uniform background cue to be the average of the photoreceptor excitations available from the unoccluded uniform background—not the average photoreceptor excitations of the entire scene. The computation of this cue presupposes that the visual system can identify the parts of the visual field that correspond to the background.

6.2. Stimuli

The spatial layout of the stereo pairs used resembled that shown in Fig. 5. The matte component of the spheres was matched to the Munsell chip with coordinates BG 5/4 and the matte component of the background to the Munsell chip with coordinate N 3/ (Kelley et al., 1943).

6.3. Results and discussion

The data are plotted in the CIE chromaticity diagram in $u'v'$ coordinates. Fig. 10A and B show the results of achromatic settings for those images for three observers. The horizontal and vertical axes are the u' and v' coordinates of the CIE chromaticity diagram. For most observers, the visual system showed little influence following the perturbation of the background in the direction of either illuminant. Perturbation neither in the direction of Illuminant D65 nor in the direction of A has any influence.

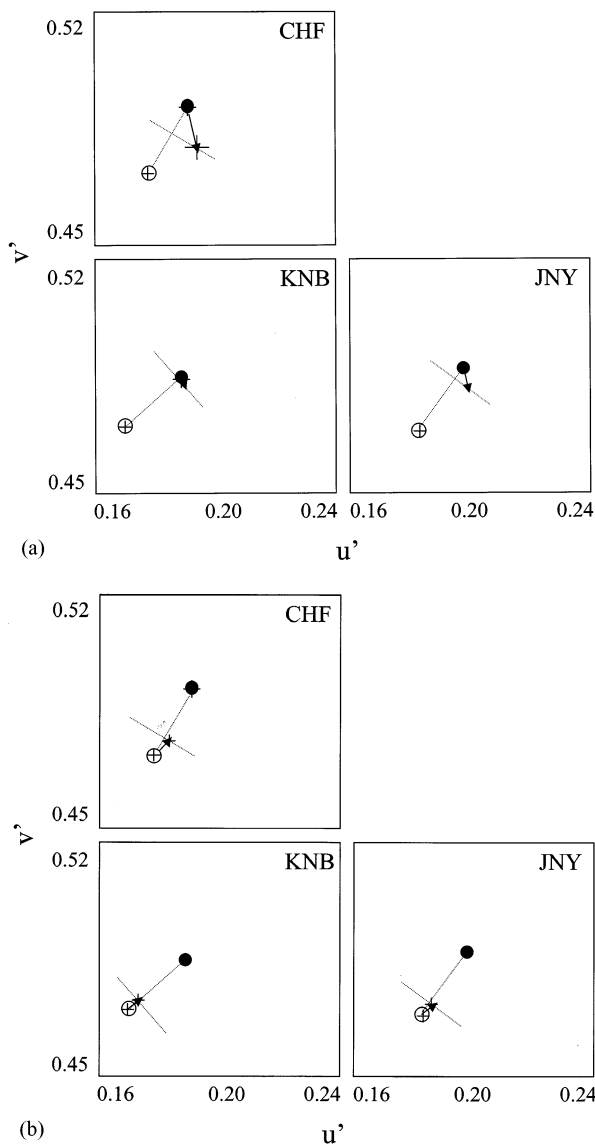


Fig. 10. A test of the uniform background cue. In experiment 4, we tested whether the uniform background cue exerts significant influence on achromatic settings. The data presentation format is identical to that of Fig. 6. For all observers, the perturbation from A to D65 or from D65 to A exerted little influence on the achromatic settings. (A) The perturbed cue signaled D65, all others, A. (B) The perturbed cue signaled A, all others, D65.

Recall that the interpretation of these results needs to be qualified, since the test patch was not located on the same plane as the background, as has usually been the case in previous studies.

7. General discussion

7.1. The equivalent background hypothesis

In perturbing the specularly cue in experiment 1, we are changing the average scene color. It is very natural to ask, could a change of this magnitude in the average

chromaticity of a scene explain the apparent influence of specularly observed here?

We know that the equivalent background hypothesis predicts human vision in simple center-surround scenes to a high degree of accuracy (Whittle, 1992). The issue is whether the equivalent background predicts color appearance in any other sort of scene. The evidence is mixed. Jenness and Shevell (1995) reject the hypothesis in center-surround scenes flecked with white and Brown and MacLeod (1997) and Hahn and Geisler (1995) reject it in simple two-dimensional scenes. There are studies, however, that support some variants of the equivalent background hypothesis (Brenner, Cornelissen, & Nuboer, 1989; Zaidi & Zipser, 1993; Brenner & Cornelissen, 1998). Our results, described next, flatly reject the hypothesis for the kinds of scenes we have used.

In Fig. 11A, the squares (open for D65, filled for A) are the equivalent backgrounds of the two consistent-cue images while the triangles are those for the two inconsistent-cue images, all plotted in the same format as the previous experimental results. The equivalent backgrounds were obtained by averaging all 500×500 pixels for R, G, and B, respectively. The change in equivalent background introduced by perturbation of the specularities is very small. Yet, if we seek to explain the results of experiment 1 through the change of

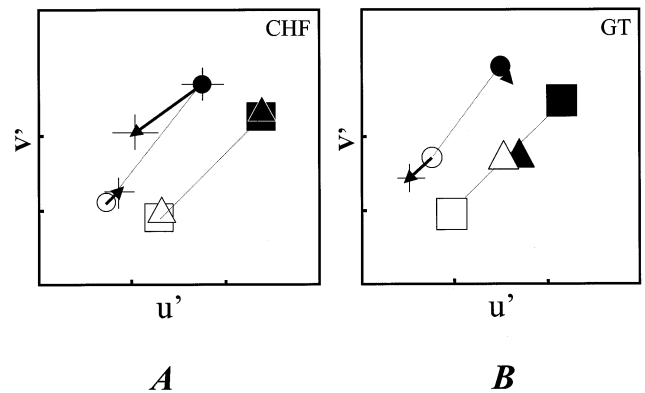


Fig. 11. Equivalent backgrounds. (A) The equivalent background (the arithmetic average of all of the chromaticities in each image of the binocular scene) is plotted for the perturbed (filled triangles) and unperturbed (blank squares) stimuli in experiment 1. The results for one observer in experiment 1 are also shown. Since the specularities were small parts of the image, each pair of perturbed and unperturbed are almost identical. However, in experiment 1, the effect of perturbation was large in some conditions. (B) The equivalent background (the arithmetic average of the scene) is plotted for the perturbed (filled triangles) and unperturbed (blank squares) stimuli in experiment 4. The results for one observer in experiment 4 are also shown. Since the uniform background constituted a large part of the image, the perturbation made a large difference in equivalent backgrounds. But we found little effect of perturbation in experiment 4. Changes in equivalent background do not predict changes in achromatic setting for the sorts of stimuli considered here. See text for further discussion.

equivalent background, then one of these very small changes led to a large change in achromatic setting (when the specular cue was perturbed from A to D65) but the other very small change in equivalent background led to a much smaller change in achromatic setting (when the specular cue was perturbed from D65 to A). Both the magnitude of change in the A-to-D65 condition and the asymmetry are inconsistent with the equivalent background hypothesis.

In contrast, consider the uniform background perturbations of experiment 4, plotted in Fig. 11B. As the plot shows, the perturbation of the uniform background had a large effect on equivalent background, since the uniform background constitutes a relatively large part of the entire scene. Yet in experiment 4, we found little or no effect of the perturbation.

In conclusion, we find that very small changes in equivalent background can have a large effect on achromatic setting (experiment 1), and that very large changes in equivalent background can have little or no effect (experiment 4).

The results of experiment 2 are also difficult to reconcile with the equivalent background hypothesis. Recall that increasing the number of specular objects does not lead to a linear but rather to a non-linear accelerating response. The effect of adding N identical objects on the arithmetic mean of a scene is proportional⁵ to N . If the 3-vector E denotes the equivalent background, then $E = aN + b$. If we normalize the 3-vector of receptor excitations at each point in the scene by the corresponding entry in E , then the effect on achromatic setting should vary as $1/(aN + b)$, a concave function. That is, the more surfaces we add, the less the impact of the most recent surface on achromatic setting. But we found that the change in achromatic setting accelerated with added surfaces. The (arithmetic) equivalent background hypothesis ('the gray world hypothesis') cannot explain our results.

7.2. Dependence on number of highlights

In the introduction, we noted that the classification of regions of a scene as neutral specular highlights (the highlight classification task) is a challenging problem for vision. As we noted before, there are many physical processes that could be mistaken for a neutral specular highlight: distant sources of light (that do not contribute appreciably to the illumination of the scene but are visible to the observer) and non-neutral specular highlights such as highlights on gold or copper. Moreover, even when the chromaticity of a spectral highlight is that of one of the illuminants in a scene, it may not be clear which surfaces are illuminated by that illumi-

nant, and which surfaces are illuminated by other illuminants. If a surface is only modestly specular, then the apparent chromaticity of a highlight may be mixed with that of the matte component. Indeed, the Lee–D'Zmura–Lennie algorithm is effectively a means to correct this matte 'bleed-through'.

It is possible, in principle, to develop heuristics to test whether 'a bright spot' in a scene is a neutral specular highlight by consideration of visual information alone. If, for example, there are many candidate specularities in a scene each of which signals a markedly different illuminant chromaticity, it is not unreasonable to ignore some or all of them. In deciding whether the illuminant chromaticity estimate from a given specular highlight should be used in estimating the surface color of a particular surface patch, it is plausible to take into account the relative three-dimensional location of patch and highlight within the scene. Bloj, Kersten, and Hurlbert (1999), for example, have shown that perceived surface color is affected by perceived three-dimensional scene layout. Last of all, in binocular viewing, specularities can be readily distinguished from 'bright spots' painted onto a surface. Specularities are virtual images of light sources and, viewed binocularly, do not have the disparities corresponding to the surface on which they are formed (Blake & Bühlhoff, 1990). Their disparities correspond to the full optical path from the eye to the light source by way of the surface.

In experiment 2, we found no effect of a single specular highlight (one sphere) on achromatic setting (in agreement with Hurlbert (1989)). We find an effect of specular highlight (in experiments 1 and 2) only when there are many identical specular highlights (eleven spheres). It may be that the visual system requires multiple identical highlights before giving weight to any one of them (the *redundancy check hypothesis*). Alternatively, it may be that only when there are identical highlights at many locations across a region containing the achromatic patch, does the visual system make use of them in estimating the illumination across that region (the *spatial variation hypothesis*). Either hypothesis is consistent with the results for experiment 3. The highlights on the eleven specular spheres with different matte components were of obviously different colors (we used a wide range of matte chromaticities to enhance the applicability of the Lee–D'Zmura–Lennie algorithms). It is possible that these eleven highlights failed a 'redundancy check' or, alternatively, that the differences in highlight signaled a scene with a markedly non-uniform illuminant. For such a scene, there is no plausible reason to extrapolate local estimates of illuminant chromaticity (derived at each specular highlight) to the region containing the achromatic test patch.

⁵ We ignore the effect of mutual illumination among added objects.

7.3. *The observed asymmetry between illuminant conditions*

The observed asymmetry between the two perturbation conditions in experiment 1 also deserves comment. The specular cue had negligible influence when it signaled Illuminant A in a scene otherwise consistent with D65, but it had appreciable influence when it signaled Illuminant D65 in a scene otherwise consistent with A. These observed cue interactions suggest two hypotheses concerning illuminant cue combination, each based in the existing computational literature. The first hypothesis is the *Bayesian*: different illuminants are assigned different prior probabilities and the visual systems takes these priors into account in estimating illuminant chromaticity (D'Zmura, Iverson, & Singer, 1995; Brainard & Freeman, 1997). It is not obvious how to assign these prior probabilities to illuminants or to cues so as to reproduce the cue interactions we observe, but we have not been able to rule out this possible explanation.

The second hypothesis is drawn from the literature on statistical 'robustness' (Landy et al., 1995; Maloney, 1999). A 'robust' estimator is resistant to failures of the underlying assumptions made in collecting data, e.g. the assumption that a sample is drawn from a Gaussian distribution. The illuminant cue model employed here assumes that all of the cues signal estimates of the same illuminant chromaticity, each perturbed by unbiased Gaussian 'noise'. If one or more cues, in fact, signal a markedly different illuminant estimate from the remainder, then this assumption is false and any estimate based on it would be misleading.

As we noted in Section 1, there are physical processes (distant light sources, non-neutral specular highlights, etc.) that could be confused with neutral specular highlights and that could signal potentially discrepant illuminant chromaticities. A 'robust' statistical estimator must decide which cues to use and which to discard when cues disagree. A possible rule appropriate for a 'robust' visual system is: When specular signals a markedly non-neutral illuminant, and other cues do not, then specularity should be given little weight. Of course, this rule merely restates our results. Further experiments are needed to test both the Bayesian hypothesis and the 'robustness' hypothesis and to determine if either plays any significant role in illuminant cue combination.

7.4. *The choice of surfaces and lights*

In these experiments, we rendered the matte component of each surface as if its surface spectral reflectance function was specific that of a specific color sample drawn from the Nickerson Munsell collection (Kelley et al., 1943). We listed these samples by Munsell coordi-

nates in the methods sections and our results can therefore be compared to those of other researchers who've used this reference set or similar reference sets. The Munsell color samples are not representative of the surfaces that we are likely to encounter in everyday life (Bonnardel & Maloney, 2000) and it is natural to ask how our results might be affected had we chosen to simulate a different set of surfaces in rendering. This is a very reasonable question to ask concerning any empirical study of surface color perception, and, for many of the candidate illuminant cues, we would have to specify a specific set of surface spectral reflectances as part of the 'environment' of the algorithm (Maloney, 1999).

The somewhat surprising answer, for our studies, is that we would expect little or no change in our results. The algorithms corresponding to the specular cues we are considering are framed in terms of the color signals⁶ of the matte and specular components of color surfaces. The perturbations only involve the specular component. The matte components are chosen arbitrarily from among the Munsell color samples which, under the illuminants we employ, produce a wide range of color signals. Suppose that we used a non-Munsell surface as the matte component of either of the spheres or background in experiments 1–3. Then, so long as the color signal of this surface matched the color signal of some Munsell surface, we have effectively used that Munsell surface. Of course, the metamer under one illuminant may take on a distinctly different color appearance under the other illuminant. But, in experiments 1–3, we only compare human performance across conditions where the illumination of the matte component is the same (the perturbation only affects the specular component).

In experiment 4, had we found an effect of the uniform background on the achromatic settings, then the precise choice of the matte component of the background would be important. Suppose, for example, that we chose a surface spectral reflectance that was metameric under D65 to our choice of Nickerson Munsell background (this surface cannot be one of the Nickerson Munsell surfaces no two of which are metameric under D65). Then, under the second illuminant A, the apparent color of our choice of background and the background metameric to it could be different. If this apparent color (or the LMS chromaticity of the background) affected achromatic settings, then the choice of surfaces would affect our conclusions. However, we found no effect of the uniform background cue

⁶ The color signal is the LMS-chromaticity of the light re-emitted from a surface under a particular illuminant that arrives at the viewer's eye. Under binocular viewing conditions, each surface gives rise to two color signals but, for simplicity in the following discussion, we refer to only a single color signal.

in experiment 4. Consequently, while the choice of surface reflectances in simulated (or real) scenes could, in principle, affect the outcome of experiments, there is little reason to believe that it affected the outcomes of the experiments reported here.

7.5. The three-dimensional location of the test patch

In pilot work we found that the test patch position was not important so long as it was *not* in the background plane. When the test patch was in the background plane, we found that the specular highlight had no influence under any of the experimental conditions and that the uniform background cue had appreciable influence. There are several reports in the lightness literature suggesting that the spatial interpretation of a scene can profoundly affect perceived lightness (Gilchrist, 1977, 1980; Gilchrist, Delman, & Jacobsen, 1983; Schirillo & Shevell, 1993; Gilchrist, 1994). Gilchrist et al. (1999) proposed that (complex) scenes are segmented into illumination ‘frameworks’ with separate illuminant corrections within each framework. The rules for organizing frameworks and assigning surfaces to them are complex. The three-dimensional layout of a scene certainly influences the segmentation of scenes into frameworks and it is likely that there are analogous effects of three-dimensional layout on color perception (e.g. Bloj et al., 1999). Our results suggest that the rules governing illuminant estimation in color tasks may lead to segmentation of scenes into multiple ‘frameworks’, each with a separate illuminant estimate, and that different frameworks may make use of possibly distinct illuminant cues.

7.6. Degree of color constancy

In the experiments reported here, we quantified the degree of color constancy of each observer using a modified Brunswick ratio:⁷ the values obtained ranged from 0.57 to 0.79, with an average of 0.65. The modified Brunswick ratios obtained by Brainard and colleagues, with observers in real scenes, averaged 0.84 (Brainard, 1998). With simulated images, others report markedly lower modified Brunswick ratios: 0.50 or less (Arend et al., 1991; Kuriki & Uchikawa, 1998). The observers who viewed our simulated scenes are evi-

dently compensating for illuminant changes to a greater extent than did observers in previous experiments using simulated scenes displayed on CRT monitors. Equally evidently, they do not compensate to the same degree as observers in real scenes do.

The enhanced color constancy performance we encountered could be due to any of several factors. (1) Our scenes were three-dimensional and contained a small collection of readily identifiable objects, (2) we presented our scenes binocularly at high-resolution (500 × 500 pixels) and (3) we used more accurate rendering methods than are typical. We do not know which of these factors, if any, contributed to the high observed values of the modified Brunswick index. We also do not know why observers in real scenes exhibit even higher values on the index. The limited field of view (20° × 20° of visual angle), the Lambertian-specular assumption used in rendering, and the lack of ego-motion (observers were confined in a chin rest and instructed not to move) must be numbered among possible factors that might be responsible for the discrepancy in the modified Brunswick index observed.

8. Conclusions

In this study, we examined how achromatic settings in simple, three-dimensional scenes were influenced by three candidate illuminant cues taken from the computational literature. We considered two specular cues, the *specular highlight cue* and the more sophisticated *full surface specular cue*, and concluded that, for the kinds of scenes we employed, the former cue influenced color appearance, the latter did not. We cannot exclude the possibility that, with some other stimulus configuration, we might find that the latter cue does influence achromatic setting. However, it is arguable that the stimulus configuration used (specular balls attached to a vertical plane) is ideal for any specular cue. Further, there is no experimental evidence suggesting that the full surface specular cue is ever active in human vision.

For the specular highlight cue, we found an asymmetry of influence. Under Illuminant A, the cue had considerable influence for all observers, under Illuminant D65 much less. We speculate that a highly colored specular component in a scene where other illuminant cues signal that the illuminant is neutral is, perhaps, suspicious. Such a specular component may be due to a specular surface whose specular component is not spectrally-neutral (such as a gold mirror) or perhaps the specularity is really a self-luminous source such as the red traffic light we discussed earlier. The visual system is possibly organized to somehow avoid confusing spectrally-neutral specularities with other phenomena.

⁷ The degree of achromatic color constancy was calculated using the idea of the equivalent illuminant (Brainard, 1998), which is basically an extension of the Brunswick ratio (Lucassen & Walraven, 1996). The idea is that given two achromatic settings for two different illuminants, one finds the von Kries coefficients that relate the two achromatic settings. Then, these coefficients were applied to the CIE locus of one illuminant to arrive at the estimated second illuminant, or the equivalent illuminant. We symmetrized the results for the two illuminants in computing an overall index. See Brainard (1998) for details.

Illuminant Cue Combination

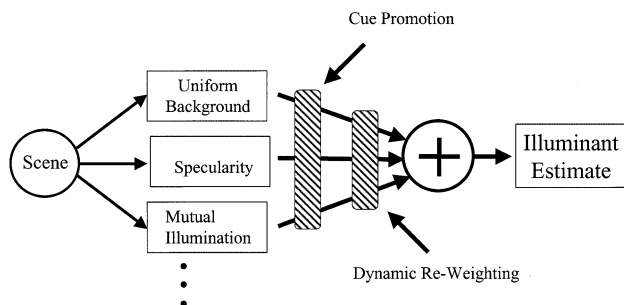


Fig. 12. Illuminant cue combination. The visual system is assumed to compute a separate illuminant estimate based on each illuminant cue. These estimates are then *promoted* (see Maloney & Yang, 2001) to a common format and then combined by a weighted average. The weight assigned to each cue may vary from scene to scene (*dynamic reweighting*) and the influence measures described in the text are estimates of the current weight assigned to a cue. The final combined estimate of the illuminant determines surface color appearance.

The third cue considered, the *uniform background cue*, had little effect under the conditions of experiment 4. As just discussed, our results are inconsistent with an equivalent background hypothesis where the equivalent background has the average chromaticity of the scene. Note however that, in our experiments the test patch that the observer adjusted was not embedded in the background but was placed a short distance in front of it, tangent with the nearer surface of the spheres. We found in pilot testing that, when the test patch was embedded in the background, that *only* the background had influence on color appearance.

8.1. The illuminant estimation hypothesis

Our results, together with the results of others, suggest that at least two illuminant cues are active in human color vision: a specularity cue, and a uniform background cue, where the influence of the latter may be most pronounced in simple center surround scenes (Whittle, 1973) or when the test patch is part of the background. There are other candidate illuminant cues that remain to be tested (see Maloney, 1999), but the two we have so far allow us to formulate a picture of color vision as a cue combination problem where the cues to be combined are illuminant cues (see Fig. 12).

Illuminant estimation can therefore be formulated as a cue combination problem analogous to depth cue combination, as suggested by Maloney (1999).

We can also conclude that the influence assigned to the illuminant cues considered here can vary from scene to scene. This is plausible: any illuminant cue may be absent in some scenes. In a particular scene, there may be no specularity, no visible shadows, etc. If human

color vision made use of only one cue to the illuminant, then, when that cue was present in a scene, we would expect a high degree of color constancy and, when that cue was absent, a catastrophic failure of color constancy. Based on past research, it seems unlikely that there is any single cue whose presence or absence determines whether color vision is color constant, and consequently it is plausible that the human visual system changes the weights assigned to any cue relative to others from scene to scene, a process that Maloney and Landy (Maloney & Landy 1989; Landy et al., 1995) referred to as *dynamic reweighting*.

The results of experiment 2 also suggest that the visual system dynamically reweights the specularity cue: when a small number of specular sources are available they are given little or no weight. When there are more such sources, they are given considerable weight. It is unclear why the visual system would ignore six specularities and pay attention to nine and the precise rules that determine the influence of specularity and how it propagates to other locations in a scene remain a puzzle.

Previous models of surface color perception begin with the assumption that color correction and resultant surface colors are determined by a small number of statistics computed from a single retinal image, ignoring the three-dimensional structure of the scene. The typical statistics are the mean (the 'equivalent background', Helson (1938)), the maximum (Land & McCann, 1971), the geometric mean (Land, 1983, 1986) and the variance (Webster & Mollon, 1995; Webster, 1996; Zaidi, Spehar, & DeBonet, 1998; and Mausfeld & Andres described in Mausfeld, 1997) in the three receptor channels. It is not obvious how to go from such retinal statistics to estimates of surface color that are constant or nearly so (see, for example, Brainard & Wandell, 1986).

In contrast, we seek to examine human surface color perception in three-dimensional scenes containing physical cues to the illuminant that have been shown to be of value in discounting the illuminant. We seek to determine which cues the visual system uses and the rules that determine which cues are used in a particular scene. Our results suggest that there are multiple cues to the illuminant and that the rules for cue usage are complicated and non-linear, and that there is no ready generalization from visual performance in estimating surface colors in two-dimensional scenes to visual performance in three-dimensional scenes.

Acknowledgements

The research described here is part of a doctoral dissertation by the first author (Yang, 1999). It was supported in part by Grant EY08266 from the National

Institute of Health. The second author thanks Dr Shigeru Akamatsu of the ATR Corporation, Kyoto, Japan, for providing support and facilities, during the second author's preparation of the manuscript. We especially thank Michael S. Landy for advice and comments during both research and writing.

Appendix A. Step-function rendering

Computer graphics rendering packages represent the spectral information associated with both light and surface as three numbers ('spectral coordinates') that we denote as $[R, G, B]$. Rendering proceeds as follows: When a light with spectral coordinates $[R, G, B]$ is absorbed and re-emitted by a surface with spectral coordinates $[r, g, b]$, the emitted light is assigned spectral coordinates $\gamma [rR, gG, bB]$ where γ is a scalar determined by the relative position of the light source, the surface, and the observer. A typical interpretation of the spectral coordinates of a light or surface is that they are the photoreceptor excitations experience when the observer directly views the light:

$$\begin{aligned} R &= \int E(\lambda) R_1(\lambda) d\lambda \\ G &= \int E(\lambda) R_2(\lambda) d\lambda \\ B &= \int E(\lambda) R_3(\lambda) d\lambda \end{aligned} \quad (\text{A1})$$

or views the surface under a spectrally-neutral illuminant:

$$\begin{aligned} r &= \int S(\lambda) R_1(\lambda) d\lambda \\ g &= \int S(\lambda) R_2(\lambda) d\lambda \\ b &= \int S(\lambda) R_3(\lambda) d\lambda. \end{aligned} \quad (\text{A2})$$

If we compare Eqs. (A1) and (A2) with Eq. (1), ignoring γ , we see that rendering packages implicitly assume that:

$$\int E(\lambda) S(\lambda) R_k(\lambda) d\lambda \stackrel{?}{=} \int E(\lambda) R_k(\lambda) d\lambda \int S(\lambda) R_k(\lambda) d\lambda. \quad (\text{A3})$$

As the question mark indicates, this is a false assumption. Although three numbers can completely characterize the effect of light on the photoreceptors of a trichromatic observer, the light-surface interaction cannot be captured by recording three numbers for the light, three for the surface, and multiplying them component-wise. An example in Evans (1948) illustrates that the consequences for human perception of this

rendering assumption can be very large. He devised two illuminants that are identical in appearance (same $[R, G, B]$) and illustrated how the colors of objects in a simple scene (composed of bric-a-brac) had dramatically different appearances (bright yellow becomes bright red) under the two lights. Of course, the surfaces did not change their $[r, g, b]$'s and the lights have the same $[R, G, B]$ by construction. If the approximation in Eq. (A3) is employed, objects that should be yellow can end up looking red and vice versa. This is not to say that a wide range of lights and a wide range of surfaces cannot be well approximated by models with only three parameters (Maloney, 1999). However, with such parameterizations the light-surface interactions cannot be modeled by component-wise multiplication.

One way to get around this limitation is to modify the rendering package so that it computes separate images for an arbitrary number of selected wavelengths (Meyer, Rushmeier, Cohen, Greenberg, & Torrance, 1986; Trussel & Kulkarni, 1996). The wavelengths chosen could be simply the nominal wavelengths associated with measurements made with a radiospectrophotometer. For a finite collection of continuous surface reflectance functions and light spectral power densities, we can always choose a high-enough sampling density to carry out rendering that is spectrally-correct to any accuracy we choose. The practical drawback to this approach is that a modern radiospectrophotometer can take very many measurements across the visible spectrum and, using this approach, the rendering time increases linearly with number of wavelengths.

We take a different approach, approximating surface spectral reflectances and spectral power distributions by arbitrary step-functions which do not correspond to particular wavelengths and which, for empirically measured surface spectral reflectances and spectral power distributions, do not depend on the arbitrary choice of sampling density used in measuring the surface spectral reflectance or spectral power distribution so long as it is large enough.

Step-function rendering: for any N , the N -dimensional spectral coordinates $[c_1, \dots, c_N]$ are interpreted as the step heights of a step function (Fig. 13) whose step intervals are fixed. All of the step functions with these fixed step intervals form a *step-function family* and, it can be easily shown that such a family of step functions is closed under addition, scalar multiplication, and point-wise multiplication of functions. Moreover, the correspondence between the spectral coordinates $[c_1, \dots, c_N]$ and the step-functions in a step-function family is an isomorphism of algebra's: addition, scalar multiplication, and component-wise multiplication of spectral coordinates correspond precisely to addition, scalar multiplication, and point-wise multiplication of the step functions associated with them. Typical render-




Fig. 13. Step-function rendering. A surface spectral reflectance function is shown together with step function approximations with 3, 6, and 12 steps. Currently-available rendering packages typically use 3-step approximations, but can be easily modified to step function approximations with any number of steps. Scenes containing Munsell chip surface spectral reflectances and the spectral power distributions of the reference illuminants (as in the experiments reported here) are not rendered accurately when 3-step approximations are used. Rendering is satisfactory when 9-step or 12-step approximations are used. See text.

ing packages can be readily modified so that, in adding, scaling, and multiplying spectral coordinates, they are, in effect, performing the corresponding operations on step functions: if the spectral power distribution of the light emitted from a source is such a step function and all surfaces have surface spectral reflectances in the same step function family, then we can follow the light as it travels from surface to surface by simply multiplying the spectral coordinates of light and surfaces. We must of course include the scalar factors determined by the relative positions of light, surfaces, and the viewer. By interpreting rendering computations as operations on step functions, we can use standard rendering algorithms to compute light-surface interactions as modeled by the Shafer model (Cook & Torrance, 1982; Shafer, 1985) for surface spectral reflectances and spectral power distributions confined to a single step function family.

The problem with standard rendering methods is that step-function families with $N = 3$ do not provide good approximations to the spectral power distributions of illuminants of interest and to the surface spectral reflectance functions of surfaces of interest. By increasing N , we can find step function families that provide acceptable approximations to sets of lights and surfaces of interest, and where component-wise multiplication approximates the correct light-surface interaction to whatever degree we desire. We find that $N = 9$ to $N = 12$ permits accurate simulations of the surfaces (the Munsell color samples of Kelley et al. (1943)) and the standard illuminants we employ. We need only modify the dimensionality of the spectral descriptors used in rendering, select step intervals, approximate the surface

spectral reflectances and spectral power distributions of interest by step functions, and then use RADIANCE to compute the Shafer model rendering for these step functions.

References

- Arend, L. E., & Reeves, A. (1986). Simultaneous color constancy. *Journal of the Optical Society of America A*, 3, 1743–1751.
- Arend, L., Reeves, A., Schirillo, J., & Goldstein, R. (1991). Simultaneous color constancy: papers with diverse Munsell values. *Journal of the Optical Society of America A*, 8(4), 661–672.
- Benzschawel, T. (1992). CIE chromaticity diagrams and application. In H. Widdel, & D. Post, *Color in electronic displays* (pp. 48–69). New York: Plenum.
- Blake, A., & Bülthoff, H. (1990). Does the brain know the physics of specular reflection? *Nature*, 343(6254), 165–168.
- Bloj, M., Kersten, D., & Hurlbert, A. C. (1999). Perception of three-dimensional shape influences color perception through mutual illumination. *Nature*, 402, 877–879.
- Bonnardel, V., & Maloney, L. T. (2000). Daylight, biochrome surfaces, and human chromatic response in the Fourier domain. *Journal of the Optical Society of America A*, 17, 677–686.
- 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., & Freeman, W. T. (1997). Bayesian color constancy. *Journal of the Optical Society of America A*, 14, 1393–1411.
- Brainard, D. H., & Wandell, B. A. (1986). Analysis of the retinex theory of color vision. *Journal of the Optical Society of America A*, 3(10), 1651–1661.
- 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.
- Brenner, E., & Cornelissen, F. W. (1998). When is a background equivalent? Sparse chromatic context revisited. *Vision Research*, 38(12), 1789–1793.

- Brenner, E., Cornelissen, F., & Nuboer, J. F. W. (1989). Some spatial aspects of simultaneous colour contrast. In J. J. Kulikowski, C. M. Dickinson, & I. J. Murrays, *Seeing contour and colour* (pp. 311–316). Oxford: Pergamon.
- Brill, M. H. (1978). A device performing illuminant-invariant assessment of chromatic relations. *Journal of Theoretical Biology*, 71(3), 473–478.
- Brown, R. O., & MacLeod, D. I. (1997). Color appearance depends on the variance of surround colors. *Current Biology*, 7(11), 844–849.
- Buchsbaum, G. (1980). A spatial processor model for object color perception. *Journal of the Franklin Institute*, 310(0), 1–26.
- Cook, R., & Torrance, K. (1982). A reflectance model for computer graphics. *ACM Transactions on Graphics*, 1, 7–24.
- Derrington, A. M., Krauskopf, J., & Lennie, P. (1984). Chromatic mechanisms in lateral geniculate nucleus of macaque. *Journal of Physiology (London)*, 357, 241–265.
- 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. (1993). 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(10), 2148–2165.
- D'Zmura, M., & Lennie, P. (1986). Mechanisms of color constancy. *Journal of the Optical Society of America A*, 3(10), 1662–1672.
- D'Zmura, M., Iverson, G., & Singer, B. (1995). Probabilistic color constancy. In R. D. Luce, M. D'Zmura, D. Hoffman, G. J. Iverson, & A. K. Romney, *Geometric representations of perceptual phenomena; papers in honor of Tarow Indow on his 70th birthday* (pp. 187–202). Mahwah, NJ: Lawrence Erlbaum.
- Evans, R. (1948). *An introduction to Color*. New York: John Wiley and Sons.
- Funt, B., Drew, M., & Ho, J. (1991). Color constancy from mutual reflection. *International Journal of Computer Vision*, 6(1), 5–24.
- 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 and Psychophysics*, 28, 527–538.
- Gilchrist, A. L. (1994). *Lightness, brightness, and transparency*. Hillsdale, NJ: Lawrence Erlbaum.
- Gilchrist, A. L., Delman, S., & Jacobsen, A. (1983). The classification and integration of edges as critical to the perception of reflectance and illumination. *Perception and Psychophysics*, 33, 425–436.
- Gilchrist, A., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., Spehar, B., Annan, V., & Economou, E. (1999). An anchoring theory of lightness perception. *Psychological Review*, 106, 795–834.
- Hahn, L. W., & Geisler, W. S. (1995). Adaptation mechanisms in spatial vision-I. Bleaches and backgrounds. *Vision Research*, 35(11), 1585–1594.
- Hardy, L. G. H., Rand, G., & Rittler, M. C. (1957). *H-R-R pseudoisochromatic plates*. New York: American Optical Company.
- van Hateren, J. H. (1993). Spatial, temporal and spectral pre-processing for color vision. *Proceedings of the Royal Society London Series B*, 251, 61–68.
- von Helmholtz, H. (1962). *Helmholtz's Treatise on Physiological Optics*. New York: Dover (edited by J.P.C. Southall).
- Helson, H. (1938). Some factors and implications of color constancy. *Journal of the Optical Society of America*, 33, 555–567.
- Helson, H. (1943). 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.
- Hurlbert, A. (1989). *The computation of color*. PhD Dissertation, Harvard Medical School/Massachusetts Institute of Technology, MA.
- Hurlbert, A. C. (1998). Computational models of color constancy. In V. Walsh, & J. Kulikowski, *Perceptual constancies; why things look as they do* (pp. 283–322). Cambridge: Cambridge University Press.
- Ives, H. (1912). The relation between the color of the illuminant and the color of the illuminated object. *Transactions of Illuminating Engineering Society*, 7, 62–72.
- Jenness, J. W., & Shevell, S. K. (1995). Color appearance with sparse chromatic context. *Vision Research*, 35(6), 797–805.
- Kelley, K. L., Gibson, K. S., & Nickerson, D. (1943). Tristimulus specification of the Munsell book of color from spectrophotometric measurements. *Journal of the Optical Society of America*, 33, 355–376.
- Krauskopf, J., Williams, D. R., & Heeley, D. W. (1982). Cardinal directions of color space. *Vision Research*, 22(9), 1123–1131.
- Kuriki, I., & Uchikawa, K. (1996). Limitations of surface-color and apparent-color constancy. *Journal of the Optical Society of America A*, 13(8), 1622–1636.
- Kuriki, I., & Uchikawa, K. (1998). Adaptive shift of visual sensitivity balance under ambient illuminant change. *Journal of the Optical Society of America A*, 15(9), 2263–2274.
- 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(3), 389–412.
- Larson, G. W., & Shakespeare, R. (1997). *Rendering with radiance*. San Francisco: Morgan Kaufmann.
- Lee, H. C. (1986). Method for computing the scene-illuminant chromaticity from specular highlights. *Journal of the Optical Society of America A*, 3(10), 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.
- Lucassen, M. P., & Walraven, J. (1996). Color constancy under natural and artificial illumination. *Vision Research*, 36(17), 2699–2711.
- MacLeod, D., & Boynton, R. (1979). Chromaticity diagram showing cone excitation by stimuli of equal luminance. *Journal of the Optical Society of America A*, 69(0), 1183–1186.
- 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, 1673–1681.
- Maloney, L. T. (1999). Physics-based models of surface color perception. In K. R. Gegenfurtner, & L. T. Sharpe, *Color vision: from genes to perception* (pp. 387–416). Cambridge: Cambridge University Press.
- Maloney, L. T., & Landy, M. S. (1989). A statistical framework for robust fusion of depth information. *Visual Communications and Image Processing IV*, 1199, 1154–1163.
- 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, 29–33.
- Maloney, L. T., & Yang, J. N. (2001). The illuminant estimation hypothesis. In R. Mausfeld, & D. Heyer, *Colour vision: from light to object*. Oxford University Press, Oxford (in press).
- Mausfeld, R. (1997). Colour perception: from Grassman codes to a dual code for object and illuminant colours. In W. Backhaus, R. Kliegl, & J. Werner, *Color vision* (pp. 1–44). Berlin: De Gruyter.

- Meyer, G. W., Rushmeier, H. E., Cohen, M. F., Greenberg, D. P., & Torrance, K. E. (1986). An experimental evaluation of computer graphics. *ACM Transactions on Graphics*, 5(1), 30–50.
- 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.
- Romero, J., Garcia-Beltran, A., & Hernandez-Andres, J. (1997). Linear bases for representation of natural and artificial illuminants. *Journal of the Optical Society of America A*, 14, 1007–1014.
- Schirillo, J. A., & Shevell, S. K. (1993). Lightness and brightness judgments of coplanar retinally noncontiguous surfaces. *Journal of the Optical Society of America A*, 10, 2442–2452.
- Shafer, S. (1985). Using color to separate reflection components. *Color Research and Applications*, 4(10), 210–218.
- Tominaga, S., & Wandell, B. (1989). Standard surface-reflectance model and illuminant estimation. *Journal of the Optical Society of America A*, 6(4), 576–584.
- Trussel, H. J., & Kulkarni, M. S. (1996). Sampling and processing of color signals. *IEEE Transactions on Image Processing*, 5(4), 677–681.
- Ullman, S. (1976). On visual detection of light sources. *Biological Cybernetics*, 21, 205–212.
- Vrhel, M. J., Gershon, R., & Iwan, L. S. (1994). Measurement and analysis of object reflectance spectra. *Color Research and Applications*, 19, 4–9.
- 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.
- Whittle, P. (1973). The brightness of coloured flashes on backgrounds of various colours and luminances. *Vision Research*, 13, 621–638.
- Whittle, P. (1992). Brightness, discriminability and the ‘crispening effect’. *Vision Research*, 32(8), 1493–507.
- Wyszecki, G., & Stiles, W. S. (1982). *Color science; concepts and methods, quantitative data and formulas* (2nd edn). New York: Wiley.
- Yang, J.N. (1999). Analysis of illuminant estimation in surface color perception. PhD Dissertation, New York University, New York.
- Zaidi, Q., & Zipser, N. (1993). Induced contrast from radial patterns. *Vision Research*, 33, 1281–1286.
- Zaidi, Q., Spehar, B., & DeBonet, J. S. (1998). Color constancy in variegated scenes: the role of low-level mechanisms in discounting illumination changes. *Journal of the Optical Society of America A*, 14, 2608–2621.