Testing limits on matte surface color perception in three-dimensional scenes with complex light fields

K. Doerschner [a,*], H. Boyaci [a], L.T. Maloney [b,c]

[a] *Department of Psychology, University of Minnesota, 75 Elliot Hall, S218, Minneapolis, MN 55455, USA*
[b] *Center for Neural Science, New York University, 4 Washington Place, New York, NY 10003, USA*
[c] *Department of Psychology, New York University, 6 Washington Place, New York, NY 10003, USA*

Abstract

We investigated limits on the human visual system's ability to discount directional variation in complex lights field when estimating Lambertian surface color. Directional variation in the light field was represented in the frequency domain using spherical harmonics.

The bidirectional reflectance distribution function of a Lambertian surface acts as a low-pass filter on directional variation in the light field. Consequently, the visual system needs to discount only the low-pass component of the incident light corresponding to the first nine terms of a spherical harmonics expansion [Basri, R., Jacobs, D. (2001). Lambertian reflectance and linear subspaces. In: *International Conference on Computer Vision II*, pp. 383–390; Ramamoorthi, R., Hanrahan, P., (2001). An efficient representation for irradiance environment maps. *SIGGRAPH 01*. New York: ACM Press, pp. 497–500] to accurately estimate surface color. We test experimentally whether the visual system discounts directional variation in the light field up to this physical limit. Our results are consistent with the claim that the visual system can compensate for all of the complexity in the light field that affects the appearance of Lambertian surfaces.
© 2007 Elsevier Ltd. All rights reserved.

Keywords: Matte surface color perception; Light field; Directional variability; Color constancy; Spherical harmonics; Spatial frequency

1. Introduction

In everyday scenes, illumination typically varies from location to location. At any fixed point within a scene, the light arriving from different directions may also vary. As a consequence, the light absorbed by a small matte patch in the scene can depend on both its location and its orientation and any attempt to estimate the surface albedo or color of the patch must effectively discount this variation in the light field (Gershun, 1939).

Researchers have shown that the visual system does compensate in part for changes in surface orientation and position when estimating the albedo or color of matte surfaces in three-dimensional scenes (Hochberg & Beck, 1954; Flock & Freedberg, 1970; Gilchrist, 1977, 1980; Bloj, Kersten, & Hurlbert, 1999; Boyaci, Maloney, & Hersh, 2003; Boyaci, Doerschner, & Maloney, 2004; Ripamonti et al., 2004; Doerschner, Boyaci, & Maloney, 2004; Maloney, Boyaci, & Doerschner, 2005; Snyder, Doerschner, & Maloney, 2005; Boyaci, Doerschner, & Maloney, 2006).

The lighting in these experiments typically consisted of a small number of light sources—a single punctate source and occasionally an additional diffuse source—by adding light sources, one could create ever more complex light fields.

This leads to an important question: can the visual system continue to estimate matte surface color as well as it has under simple lighting conditions as we increase the complexity of illumination in the scene? Alternatively put, for what range of light fields can the visual system compensate in estimating the albedo or color of matte surfaces?

In this study we address this question by first describing an upper limit on the possible complexity of light fields. This limit, corresponding to a spherical harmonic

---

subspace, is based on analytic work by Basri and Jacobs (2001) and Ramamoorthi and Hanrahan (2001). We then report an experiment intended to test a conjectured lower limit, framed in the frequency domain. In order to describe the results, analytical and experimental, and their implications for surface color perception, we develop notation for the light field, introduce spherical harmonics and describe a remarkably simple characterization in the frequency domain of how light and matte surfaces interact.

1.1. The light field

The light field of a given space is the illumination as a function of direction and position (similar to the plenoptic function of Adelson & Bergen, 1991). The *local light field* $E_\mathbf{p}(\theta, \varphi, \lambda)$, is the illumination arriving at a point $\mathbf{p}$ in three-dimensional space, specifying the intensity of light as a function of direction, with elevation θ ranging from 0 to π, and azimuth ranging from 0 to 2π, and wavelength λ in the electromagnetic spectrum. The local light field (or *environment map*) can be visualized as a collection of light rays, each with a spectral power distribution, originating from all directions and impinging on point $\mathbf{p}$ (Fig. 1A). We concentrate on the *spatial component* of the local light field $E_\mathbf{p}(\theta, \varphi)$, and how it varies as a function of elevation θ, azimuth φ, at a fixed point $\mathbf{p}$ (See Fig. 1A and B, and also Appendix A). We can unwrap the sphere and also plot the
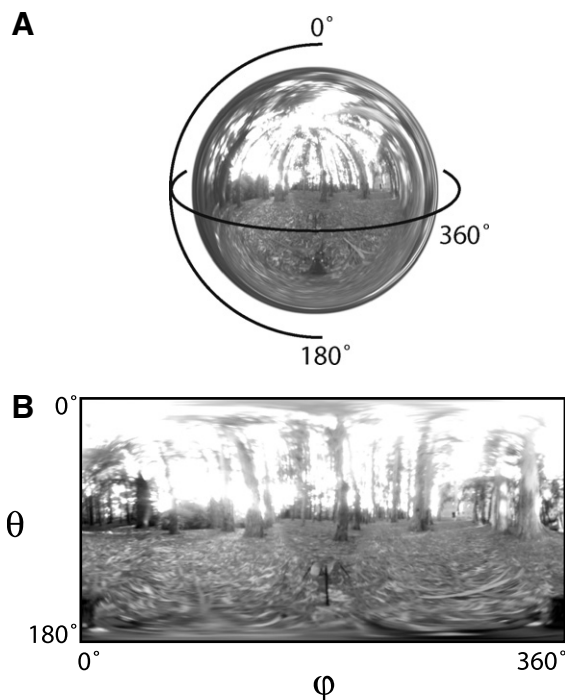

Fig. 1. Illustration of the light field. (A) A light field, φ denotes azimuth, θ denotes elevation. (B) An 'unwrapped' version where azimuth and elevation now correspond to horizontal and vertical axes, respectively. The environment map for this illustration was obtained from http://www.debevec.org/Probes.
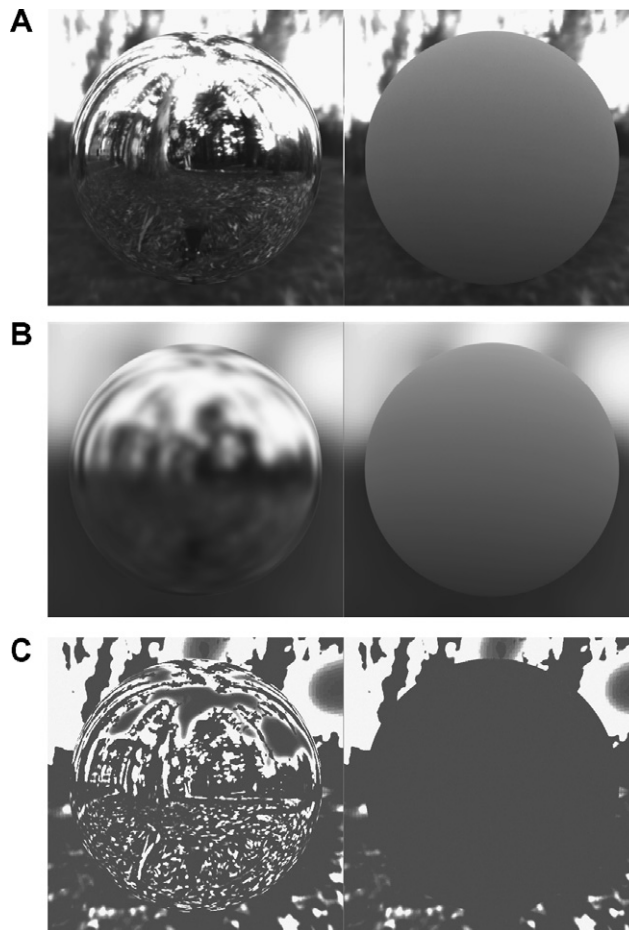

Fig. 2. Light field directional structure and surface appearance. Images were rendered with *Radiance* (Ward, 1994). (A) A mirror sphere (left) and matte sphere (right) were rendered using the same light probe (http://www.debevec.org/Probes). (B) We low-pass filtered the light probe, removing fine spatial detail. Under this illumination the mirror sphere ceases to look chrome-like and changes its appearance to that of a brushed metal. The matte sphere however does not change in appearance. In the text we explain why. (C) A was subtracted from B. The values of the resultant difference image were squared Dark gray indicates zero difference.

light field[1] by elevation and azimuth as illustrated in Fig. 1B.

1.2. Light field structure and surface appearance

As we noted above, in every day scenes the spatial structure of light fields can be complex considering the full range of skylight as well as the complexities of shadows and light absorbed and re-emitted by other surfaces in the scene (Fig. 1). However, not all of the directional variation in the light field may need to be taken into account by the visual system for every task.

As an example consider Fig. 2. Fig. 2A depicts two spheres: while the left one is mirror-like and the right one

---

[1] We refer from this point on to the *local light field* as *light field* (or interchangeably as environment map).

appears to be covered with a matte paint, both spheres are illuminated by the same light field. If we reduced the spatial complexity of the light field by blurring $E(\theta, \varphi)$ at each point of the mirror sphere (Fig. 2B), the sphere would now appear to be made from a different, brushed metal-like material (Fleming, Dror, & Adelson, 2003; and te Pas & Pont, 2005, 2006). However, blurring the light field has little effect on the appearance of the matte sphere in Fig. 2B. Matte surfaces seem to be 'indifferent' to the fine spatial details in the light field that are lost in blurring. (Fig. 2C shows the difference-images of 2A and 2B.)

This 'indifference' is due to the special properties of the Lambertian bidirectional reflection distribution function (BRDF; Nicodemus, Richmond, Hsia, Ginsberg, & Limperis, 1977; Dutre, Bekaert, & Bala, 2003; Cabral, Max, & Springmeyer, 1996). Basri and Jacobs (2001, see also Ramamoorthi and Hanrahan, 2001) using a spherical harmonic expansion of both a kernel that represents Lambert's reflectance (Haralick & Shapiro, 1993) and the spatial component of the light field, thereby taking their analysis into the spatial frequency domain, demonstrated that a Lambertian surface essentially acts as a low-pass filter on the light field. This 'low-pass filtering' property is precisely what accounts for the nearly unchanged appearance of the matte sphere in Fig. 2. Consequently, if the visual system's task is to judge the color or albedo of a Lambertian surface it could disregard the higher spatial frequencies in the light field. In order to quantify this limit we rely on analytical results by Basri and Jacobs (2001) and Ramamoorthi and Hanrahan (2001) described next.

1.3. 4D and 9D subspace models of the light field

Notation: in this article we are concerned with directional variation in the light field $E_{\mathbf{p}}(\theta, \varphi)$ at a single fixed point $\mathbf{p}$. For simplification, we will omit the subscript $\mathbf{p}$ in $E_{\mathbf{p}}(\theta, \varphi)$.

1.4. Spherical harmonics

Spherical harmonics, $Y_{nm}(\theta, \varphi)$, provide an orthonormal basis with which one can express a function defined on the surface of a sphere. They have been employed by numerous researchers in perception and computer graphics to study light and surface reflectance (e.g. D'Zmura, 1991; Cabral et al., 1996; Nimeroff, Simoncelli, & Dorsey, 1994; Teo, Simoncelli, & Heeger, 1997; Basri & Jacobs, 2001; Ramamoorthi & Hanrahan, 2001). Any light field $E(\theta, \varphi)$ can be written as a weighted sum of the harmonic basis functions,

$$E(\theta, \varphi) = \sum_{n=0}^{\infty} \sum_{m=-n}^{n} \varepsilon_{nm} Y_{nm}(\theta, \varphi), \tag{1}$$

where ε_{nm} denote the *lighting coefficients*.[2]

The spherical harmonic expansion (Eq. (1)) defined on the sphere is analogous to Fourier series defined on an interval and, like Fourier series, they provide a convenient way to analyze spatial structure.

1.5. Spherical harmonic subspaces

The projection of any light field $E(\theta, \varphi)$ represented by a spherical harmonics series (Eq. (1)) onto the Nth subspace is formed by simply truncating the spherical harmonics expansion, as we might truncate a Fourier series after a fixed number of terms,

$$\mathbb{P}_N E(\theta, \varphi) = \sum_{n=0}^{N} \sum_{m=-n}^{n} \varepsilon_{nm} Y_{nm}(\theta, \varphi). \tag{2}$$

The projection operation replaces the light field by its 'low-pass' approximation confined to the specified subspace.[3] We will be primarily concerned with $\mathbb{P}_N$ for the cases $N = 1$ (the 4D subspace) and $N = 2$ (the 9D subspace), as explained below.

When $N = 1$, the first four spherical harmonic basis functions[4] $\{Y_{00}, Y_{1,-1}, Y_{10}, Y_{11}\}$ (Fig. 3) span a subspace that can represent one maximum of intensity (one 'bump' of light) anywhere on the sphere. Any light field in this subspace is completely specified by the lighting coefficients $\varepsilon_{00}, \varepsilon_{1,-1}, \varepsilon_{10}, \varepsilon_{11}$. The first nine spherical harmonic components $\{Y_{00}, Y_{1,-1}, \cdots, Y_{20}, Y_{21}, Y_{22}\}$ (Fig. 3) span a 9D subspace (corresponding to $N = 2$) that can represent light fields with up to two maxima of intensity (two 'bumps' of light).

1.6. The 4D subspace conjecture

If a visual system had an accurate estimate of the nine coefficients $\varepsilon_{00}, \varepsilon_{1,-1}, \varepsilon_{10}, \varepsilon_{11}, \varepsilon_{2,-2}, \varepsilon_{2,-1}, \varepsilon_{20}, \varepsilon_{21}, \varepsilon_{22}$ which specify the low-pass component[5] $\mathbb{P}_2 E(\theta, \varphi)$ of the light field in the 'two bump' subspace it could estimate matte surface color with high accuracy (See Appendix B). However, an even tighter conjecture can be formulated.

Previous results (e.g. by Boyaci et al., 2003, 2004, 2006; Ripamonti et al., 2004; Doerschner et al., 2004; Maloney et al., 2005) demonstrated that observers discount directional variation in the light field, in the case where $E(\theta, \varphi)$ consisted of a single punctate source (and, in some cases, an additional diffuse source) when estimating the color of matte surfaces. However, all of these results could

---

[2] Lighting coefficients are computed as $\varepsilon_{nm} = \int_0^\pi \int_0^{2\pi} E(\theta, \varphi) Y_{nm}(\theta, \varphi) \sin\theta \mathrm{d}\theta \mathrm{d}\varphi$.

[3] For any $N \geqslant 0$, the spherical harmonics functions form a basis for a subspace that is closed under rotation. The subspace corresponding to $N = 0$ contains only the rotationally invariant constant function Y_{00}. As $N \to \infty$ subspaces can represent light fields with increasing spatial complexity.

[4] Note that when m is negative in the subscript nm, we insert a comma to improve readability.

[5] We will use the terms 'low frequency', 'high frequency', 'low-pass', etc. by analogy to Fourier series.
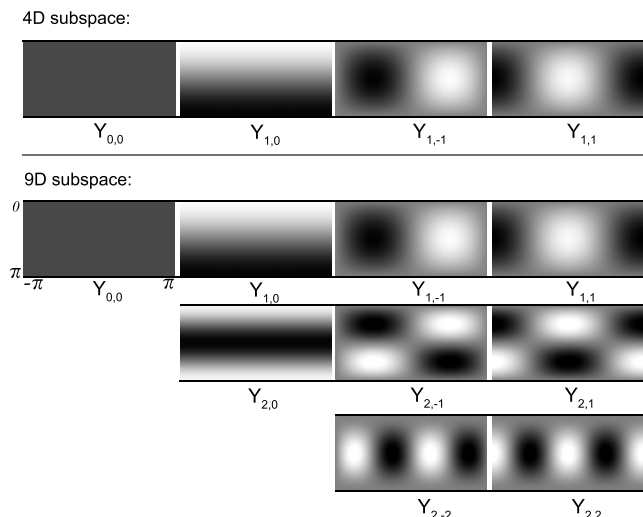
4D subspace:

9D subspace:

Fig. 3. The 4D and 9D spherical harmonics subspaces. Shown are the first nine spherical harmonics basis functions in spherical coordinates. Since they reside on the sphere, the dimensions of these basis functions are $\pi \times 2\pi$ (elevation × azimuth). Middle gray denotes zero, white are positive, and black are negative values. Vertical axes correspond to the elevation in the range $[0\,\pi]$, horizontal axes correspond to azimuth in the range $[-\pi\,\pi]$. When approximating an arbitrary light field with the 1st order harmonics (or the 4D subspace: $Y_{0,0}, Y_{1,-1}, Y_{1,0}, Y_{1,1}$), one can have at most one maximum of intensity (one-bump) in the expansion. If a light field is approximated with 2nd order harmonics (or the 9D subspace: $Y_{0,0}, Y_{1,-1}, Y_{1,0}, Y_{1,1}, Y_{2,-2}, Y_{2,-1}, Y_{2,0}, Y_{2,1}, Y_{2,2}$), two (sufficiently distant from each other) maxima of intensity in the light field can be detected in the expansion.

be accounted for if observers discounted merely the 4D approximation $\mathbb{P}_1 E(\theta, \varphi)$, the 'one bump' subspace. We note that a 4D subspace representation of the light field still accounts for 94% of the appearance variation in Lambertian surfaces (Basri & Jacobs, 2001).

But, if the visual system only resolves the directional variation in the illumination up to a 4D subspace, then observers should not be able to discount the effects of a light field which consists of two widely spaced distant punctate sources (two maxima of intensity). In order to discount more than one punctate source the visual system needs to resolve the directional variation in the light field at least up to a 9D spherical harmonic subspace. We therefore test the 4D subspace conjecture by examining human perception of matte surface color in scenes with two punctate sources as described in the Section 2.2.

2. Experiment

2.1. Introduction

The main goal of the experiment was to assess the visual system's resolution of the directional variation in the illumination in the scenes used. More specifically, we wished to discover whether observers discount the directional variation in the illumination equivalent up to the frequency limit of a 4D or of a 9D subspace. We accomplished this by measuring observers' achromatic settings (Helson & Michels, 1948) of a Lambertian test patch embedded in rendered scenes which were illuminated by a superposition of three light sources: two were angularly separated punctate

lights placed sufficiently far from the scene so as to act as collimated sources. The third light source was diffuse. The two punctate sources differed in chromaticity from the diffuse light (see Section 2.2 for details). As explained below, the two punctate sources were positioned so that predictions of observers' performance for discounting directional variation in the illumination up to a 4D subspace would clearly differ from the predictions for an observer discounting the directional variation up to a 9D subspace.

2.2. Methods

2.2.1. Familiarization procedure

Prior to the main experiment observers participated in a short procedure intended to familiarize them with the task.

2.2.2. Setup

Stimuli were two 5×5 arrays of colored paper chips. Each array contained exactly the same colored chips; however the position of a chip in array 1 and 2 was randomized. Each chip was indexed by a number, referring to its position. The two arrays were placed on the back wall of a black-painted box. The box was open along one side and contained a separation wall at the center, creating two compartments. The left compartment was illuminated by a reddish light, the right by a neutral-yellowish light. The light sources were not directly visible to the observer. Observers were seated in an otherwise dark room, about 40 cm away from the open side of the black box (about 80 cm from the stimuli), and were allowed to view the two arrays of colored chips freely.

2.2.3. Task

First, the experimenter read out loud a number corresponding to the position of a chip located in the left array and the observer's task was to name the number of, or point to the chip in the right array which he or she thought was made from the same piece of colored paper. This was continued for a total of 10 trials. After that, the experimenter placed colored chips in the left compartment, (illuminated by the reddish- light) and the observer was asked to pick up each chip and to move it from the left compartment to the right (in effect changing the illumination arriving upon that chip). Observers were not given any feedback about their performance upon completion of the session. Completion of all trials took about 5–7 min.

Note that, this experiment was not a screening test, and merely served to familiarize the observer with judging surface color under illuminants differing in chromaticity and to clarify the instructions given to the observer in the main experiment. Therefore performance was not analyzed and is not reported here.

2.3. Main experiment

2.3.1. Stimuli

Our stimuli were computer-generated scenes (software: *Radiance,* Ward, 1994) each composed of a matte ground plane and a number of objects of various shapes (such as spheres, cubes, or cylinders), sizes, and reflectance properties. These objects served as possible cues to the spatial distribution, intensities and chromaticities of the light sources (Yang & Maloney, 2001; Boyaci et al., 2006) in our scenes. Their positions varied randomly from trial to trial. The background in the scenes was dark blue. All scenes contained a smooth, Lambertian rectangular test patch at the center and were rendered twice from slightly different viewpoints corresponding to the positions of the observers' eyes. The stimuli were viewed in a computer-controlled stereoscope. Fig. 4 depicts a typical stereo pair of stimuli.

2.3.2. Tri-stimulus rendering

Most rendering packages render a three-channel approximation to light-surface interaction in scenes. This approximation is not always accurate (See the 'RGB Heuristic', Maloney, 1999) particularly when light
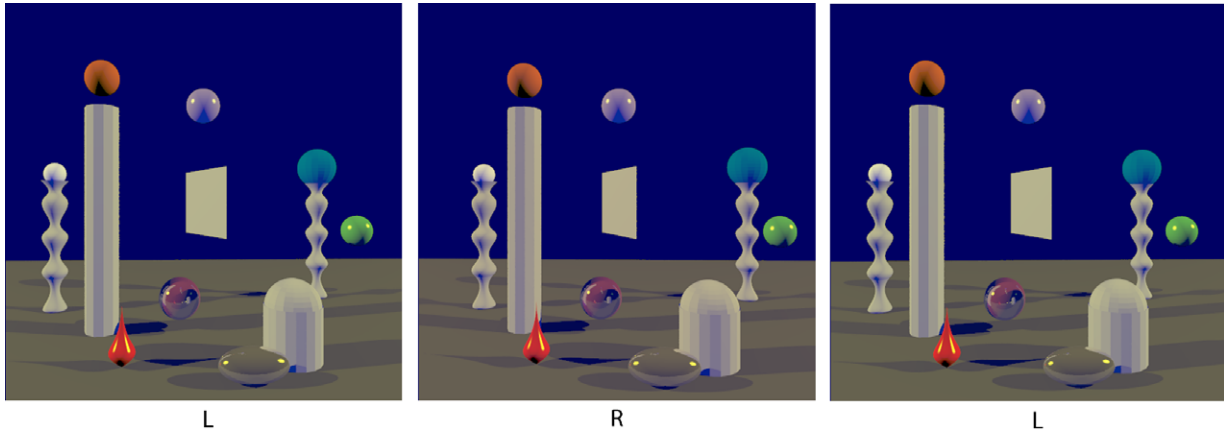
Fig. 4. Example of a stereo pair of stimuli (for crossed and uncrossed fusion). Stimuli were computer-generated scenes each composed of a matte ground plane and a number of objects of various shapes, sizes, and reflectance properties. Scenes contained a Lambertian, rectangular test patch at the center. Observers viewed stimuli in a computer-controlled stereoscope.

from one surface illuminates a second (inter-reflection). To avoid such problems we employed the techniques described in Doerschner et al. (2004, pp. 96), in particular, we designed scenes so that light from other surfaces in the scene could not illuminate the test surface.

2.3.3. Illumination: Spherical coordinates

The azimuths of the two punctate sources are denoted φ_i, $i = 1,2$, and the elevations θ_i, $i = 1,2$. The elevations were always set to $\theta_i = 65°$, $i = 1,2$. The azimuths were varied between conditions. In the 160°-condition, the two punctate lights were placed symmetrically about the line of sight with $\varphi_1 = 80°$ and $\varphi_2 = -80°$. The separation on azimuth of the punctate sources was therefore 160°. In the 90°-condition, the two punctate lights were placed symmetrically about the line of sight with $\varphi_1 = 45°$ and $\varphi_2 = -45°$. The separation in azimuth of the punctate sources was 90° in this condition (Figs. 5 and 6A). The punctate light sources were placed sufficiently far (665 cm) from the test patch and we treat them as collimated sources (Kaufman & Christensen, 1972).

2.3.4. Illumination: Chromaticities

The light radiating from surfaces was described as a weighted mixture of three abstract primary lights (red, green, blue), whose spectra coincide with those of the corresponding guns of the monitors. The three primaries were linearized versions of the guns. The intensities of the three primaries were measured in arbitrary units proportional to their luminance, such that a mixture of the primaries with equal intensities appeared roughly achromatic to most observers. The *tri-stimulus values* (essentially the weights of the three primaries) (Wyszecki & Stiles, 1982, p.120) that describe the light at a particular location on the monitor are referred to by us for convenience as *RGB codes*. When making an achromatic setting the observer selects the RGB code that makes the test patch appear achromatic, as described below.

There were two punctate light sources in every scene and the intensity and chromaticity of the two punctate light sources were always the same. For convenience, in defining the parameters that affect the intensity and chromaticity of the punctate light sources, we will use the subscript P to denote either light source (e.g. L_P). When we need to distinguish one light source form the other, we will add a subscript, e.g. P_1. We use the subscript D for the single diffuse source.

The RGB codes of the two punctate and the diffuse light sources were denoted by L_P^R, L_P^G, L_P^B and L_D^R, L_D^G, L_D^B. We define $L_P = L_P^R + L_P^G + L_P^B$, $L_D = L_D^R + L_D^G + L_D^B$. The chromaticity coordinates (Wyszecki & Stiles, 1982) of the punctate light sources were specified as $\pi^R = L_P^R/L_P$, $\pi^G = L_P^G/L_P$, $\pi^B = L_P^B/L_P$ and the chromaticity coordinates of the diffuse light source as $\delta^R = L_D^R/L_D$, $\delta^G = L_D^G/L_D$, $\delta^B = L_D^B/L_D$. Further we define the diffuse-punctate ratio $\Delta = L_D/(2L_P)$, a measure of the relative intensity of the diffuse light source and the two punctate light sources. The values
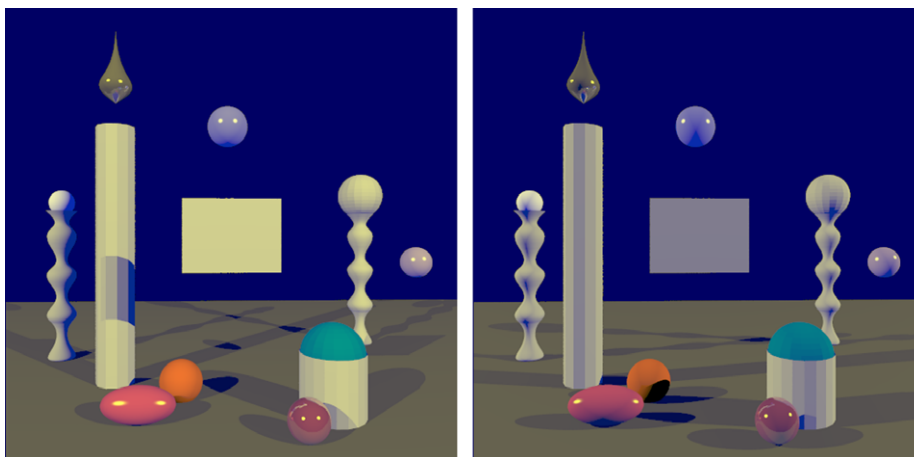

Fig. 5. Example of scenes under the 90° (left) and 160° (right) illumination conditions. Only the left image of each stereo pair is shown. Scenes are illuminated by a composition of a diffuse *blue* source and two *yellow* punctate sources either 90° apart (left), or 160° apart (right).
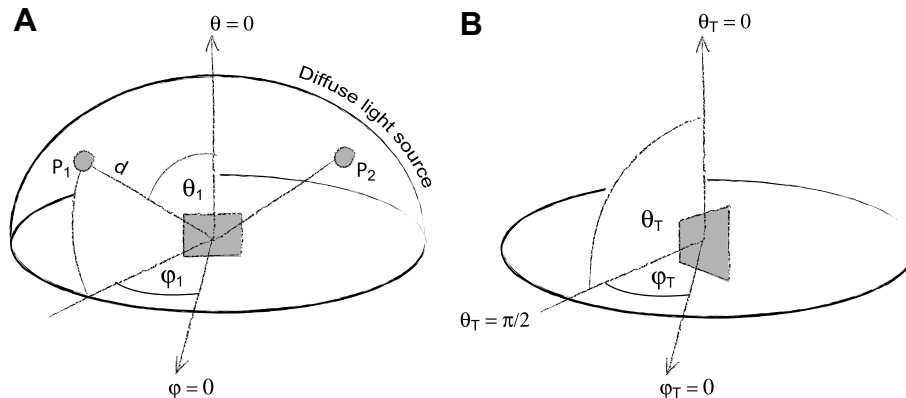
Fig. 6. Lighting and test surface coordinates. (A) Lighting coordinates. Scenes in both experiments were illuminated by a mixture of two *yellow* punctate sources (P_1, P_2) and one diffuse *blue* light. The azimuth and elevation of the *yellow* sources were denoted by θ_i and φ_i ($i = 1,2$), respectively. At $\varphi_i = 0°$ lies the line of sight of the observer. The punctate source was placed sufficiently far ($d = 665$ cm) from the test patch so that it was effectively a collimated source. Its elevation was $\theta_i = 65°$ and never varied throughout the experiment. (B) Test surface coordinates. The azimuth of the test patch was denoted φ_T ($\varphi_T = 0°$ corresponds to the line of sight of the observer, i.e. the test patch surface normal would be pointing directly towards the observer). The patch was simulated to approximately 70 cm away from the observer, and its azimuth (φ_T) could take on one out of nine orientations {$-65°$ $-45°$ $-25°$ $-10°$ $0°$ $25°$ $45°$ $65°$}. The elevation of the test patch $\pi/2$ was never varied throughout the experiment.

used in rendering our scenes were $\pi^R = 0.5$, $\pi^G = 0.5$, $\pi^B = 0$, (*yellow* punctate), $\delta^R = 0$, $\delta^G = 0$, $\delta^B = 1$ (*blue* diffuse). The diffuse-punctate ratio was set to $\Delta = .245$.[6] The terms *yellow* and *blue* will be used as mnemonics for light source chromaticities for the remainder of the paper.

We used chromatic light sources to improve the power of the experiment, i.e. the local light field varies in intensity as well as chromaticity across directions, and to maintain comparability with past work (Boyaci et al., 2004). We do not claim they are representative of illuminants in the everyday environment.

2.3.5. Test surface: Spherical coordinates

Each scene contained a Lambertian test patch at the center whose albedo was 0.55. The patch was simulated to be located about 70 cm away from the observer, and its azimuth (φ_T) could take on one out of nine orientations {$-65°$ $-45°$ $-25°$ $-10°$ $0°$ $10°$ $25°$ $45°$ $65°$} on a given trial. The elevation θ_T of the test patch was equal to $\pi/2$ and was never varied (Fig. 6B).

2.3.6. Angle of incidence

The cosine of the angle between test patch normal and a given punctate source is given by $\mathbf{n} \cdot \mathbf{p}$, where $\mathbf{n} = \mathbf{U}(\theta_T, \varphi_T)$ is the direction of the test patch normal, and $\mathbf{p} = \mathbf{U}(\theta_P, \varphi_P)$ is the unit vector in the direction to the punctate source. For the special case of our experiment, where the elevation of the test surface is always $\theta_T = \pi/2$, this reduces to,

$$\cos \eta = \sin \theta_P \cos(\varphi_T - \varphi_P). \tag{3}$$

Naturally we have to consider two angles of incidence, one associated with each punctate source, hence we will consider two cosines terms $\cos \eta_1$ and $\cos \eta_2$ in the derivations below.

2.3.7. Task

Observers were asked to adjust the color of the test patch until it was perceived to be achromatic ('cut from a gray piece of paper'). Let $[S^R, S^G, S^B]$ denote the RGB setting of the test patch which is under the control of the observer. This setting could be adjusted by using a computer keyboard, and could be varied in the *blue–yellow* (tradeoff between S^B and $S^R + S^G$) and *red–green* directions (tradeoff between S^R and S^G), while holding $S^R + S^G + S^B$ constant.

---

[6] For more information on realism & rendering and calibration (lookup table preparation & luminance measurements) please refer to the method section in Boyaci et al. (2004).

2.3.8. Randomized initial test patch settings

Before presenting a given scene we replaced the test patch's RGB code with a randomly chosen one, keeping the sum of the intensity of three primaries of the test patch constant.

2.3.9. Procedure

The experiment was blocked by illumination condition. The order was randomized for each observer; half the observers started with the 90°-condition, the other half with the 160°-condition. In each illumination block observers made achromatic adjustments for nine test patch orientations {$-65°$ $-45°$ $-25°$ $-10°$ $0°$ $10°$ $25°$ $45°$ $65°$}, each orientation was repeatedly presented 20 times. The order of test-patch orientations was randomized. There was no time constraint for the completion of a trial. Observers were allowed to practice the task for a few trials before starting the experiment. Each block took the observers less than an hour. Typically the two experimental blocks were completed on two different days.

2.3.10. Observers

Four New York University undergraduate or graduate students participated as observers in the experiment. All four observers were unaware of the purpose of the study. Observers had normal or corrected to normal visual acuity and normal color vision.

3. Analysis and results

3.1. Geometric chromaticity functions

In this section we derive what the settings of an ideal observer would be for a spectrally unselective test surface as a function of test surface orientation φ_T for different configurations of the punctate light sources. We emphasize that observers in the experiment never saw the actual gray patch whose RGB code was consistent with the rest of the rendered scene.

We define

$$T^B(\eta_1, \eta_2) = \alpha(L_P^B(\cos \eta_1 + \cos \eta_2) + L_D^B), \tag{4}$$

as the total light in the *B*-component reaching the observer's eye, where α is the albedo of the test patch and L_P^B
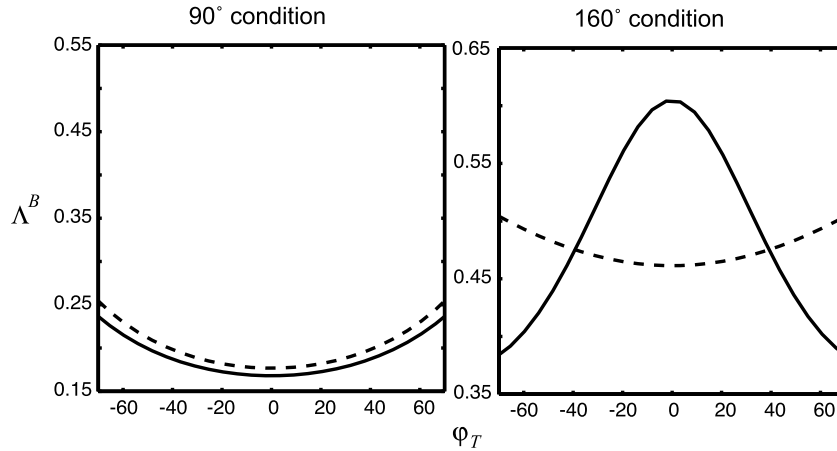
Fig. 7. Predictions of relative blue in the 90° and 160° illumination condition, for a 4D and 9D subspace expansion of the lighting model. Shown are the predictions of relative blue settings (Λ^B) as a function of test patch orientation (φ_T) for both illumination conditions (left 90°, right 160°). Graphs depict Λ^B of an ideal observer discounting the directional variation in the illumination up to a 4D subspace (dashed line) or a 9D subspace (solid line).

is the maximum possible contribution of *either* punctate light source to the *B*-component of the light. We define similar expressions for the *R*- and *G*-components, thereby specifying the tri-stimulus coordinates

$$\mathbf{T}(\eta_1, \eta_2) = [T^R(\eta_1, \eta_2), T^G(\eta_1, \eta_2), T^B(\eta_1, \eta_2)], \tag{5}$$

of the total light that would reach the eye from a neutral test patch.

We can describe the change in chromaticity of the light reflected by the test patch across different orientations by means of *geometric chromaticity functions* $\Lambda(\eta_1, \eta_2) = [\Lambda^R(\eta_1,\eta_2), \Lambda^G(\eta_1,\eta_2), \Lambda^B(\eta_1,\eta_2)]$. The *geometric b-chromaticity function* is defined as

$$\Lambda^B(\eta_1, \eta_2) = \frac{T^B(\eta_1, \eta_2)}{T^R(\eta_1, \eta_2) + T^G(\eta_1, \eta_2) + T^B(\eta_1, \eta_2)}. \tag{6}$$

$\Lambda^R(\eta_1,\eta_2)$ and $\Lambda^G(\eta_1, \eta_2)$ (the geometric r- and g-chromaticity functions, respectively) are defined analogously. We verified that predicted r-, b- and g-chromaticity values were in good agreement with those of a rendered neutral test patch at the correct position and orientation within each scene (mean squared error <0.000004).

3.2. 4D and 9D low-pass chromaticity functions

As shown in Fig. 7, $\Lambda^B(\eta_1, \eta_2)$ differs when 4D or 9D subspace projections of the light field are used. For details on the spherical harmonic expansion of diffuse and punctate sources refer to Appendix C. The chromaticity of the light reflected from the test patch is a mixture of chromaticities from *yellow* punctate and *blue* diffuse sources. As the test patch rotates away from a *yellow* punctate source the intensity of the light reflected from this source decreases with the cosine of the angle between test patch normal and light direction; consequently the relative *blue* content of the test patch increases.

The minima in the Λ^B-curves depend on the directions of the *yellow* punctate light sources ($\varphi_T = \varphi_1$ or $\varphi_T = \varphi_2$)

and the resolution of the subspace model as described next. *160°-condition.* The directional resolution of the 4D subspace of $E(\theta, \varphi)$ is limited to represent one maximum of intensity (e.g. one punctate light). A visual system discounting illumination only up to this subspace cannot resolve the two punctate sources 160° apart. In this case 4D-Λ^B (4D blue-chromaticity function) has a single minimum at 0° (Fig. 7, right plot), which does not correspond to the direction of either punctate source.

The above can be contrasted with the 'successful' predictions of the 9D subspace model. The directional resolution of this subspace is sufficient for 9D-Λ^B (9D-blue-chromaticity function) to have two 'correct' minima, each roughly corresponding to the *yellow* punctate source locations at $\varphi_1 = 80°$, $\varphi_2 = -80°$.

It should be noted that if the visual system discounted an altogether different illumination setup, namely a single *blue* punctate source located near $\varphi = 0°$ and a *yellow* diffuse source, then the resulting 4D-Λ^B would crudely approximate the shape of 9D-Λ^B with an illumination setup consisting of two *yellow* punctate sources and a *blue* diffuse source (as described previously). In the results section below we will test the possibility that observers' discounting may be consistent with this alternative effective illumination setup. *90°-condition.* We included the 90°-condition in order to assess observers' discounting behavior across different light conditions, i.e. to exclude the possibility that they adopt one particular strategy independent of the specific illumination setup. In this condition, predictions of Λ^B for 4D and 9D light field models agree, forming one effective maximum of intensity near 0° (Fig. 7, left plot). Neither light field model can resolve two punctate sources which are 90° apart,[7] however, this is not in contradiction with the response of a Lambertian

---

[7] The minimum angular separation between two punctate sources that results in two maxima of intensity in the 9D light field model is approximately 98°.

surface to the 90° lighting setup. The punctate sources are too close to each other to be resolved by the Lambertian kernel.

3.3. Relative blue

To assess the *B*-component of an observer's achromatic setting we can rewrite Eq. (6) with re-parameterization as

$$\Lambda^B = \frac{\omega\pi^B + \delta^B\Delta}{\omega + \Delta}, \tag{7}$$

We refer to Λ^B as the *observer's geometric b-chromaticity function*. π^B, δ^B, and Δ are defined as described above, and $\omega = \sum_{n=0}^{N}\sum_{n=-m}^{m}\beta_n\varepsilon_{nm}Y_{nm}$, where β_n is a scalar that is tied directly to the *n*th order coefficient of the harmonic expansion of the Lambertian kernel (Basri & Jacobs, 2003); $\beta_0 = 3.14$, $\beta_1 = 2.09$ and $\beta_2 = 0.79$ remain constant. ε_{nm} are the lighting coefficients of the spherical harmonic expansion of the punctate sources (see Appendix C). Since we assume separability of light fields (Appendix A), above derivations can be made analogously for Λ^R and Λ^G.

An observer's visual system can compute what the *b-chromaticity* of an achromatic surface should be if estimates of the parameters in Eq. (7) are available. We will denote them as $\hat{\pi}^B$, $\hat{\delta}^B$, $\hat{\Delta}$, and $\hat{\omega}$. Then Eq. (7) becomes

$$\hat{\Lambda}^B = \frac{\hat{\omega}\hat{\pi}^B + \hat{\delta}^B\hat{\Delta}}{\hat{\omega} + \hat{\Delta}}, \tag{8}$$

When obtaining estimates of these parameters from an observer's data we impose the restriction that the punctate sources are placed symmetrically over the line of sight of the observer, and never change elevation ($\theta_{1,2} = 65°$ remains constant). With those restrictions in place, when determining $\hat{\omega}$ from observers' data we only need to estimate $\hat{\varepsilon}_{11}$ when $N = 1$ ($\varepsilon_{00} = .5642$, and $\varepsilon_{1,0}$, $\varepsilon_{1,-1} = 0$ remain constant), and similarly when $N = 2$ only $\hat{\varepsilon}_{11}$ and $\hat{\varepsilon}_{22}$ need to be estimated ($\varepsilon_{00} = .5642$, $\varepsilon_{20} = -.6308$, and $\varepsilon_{1,0}$, $\varepsilon_{1,-1}$, $\varepsilon_{2,-1}$, $\varepsilon_{2,-2} = 0$ remain constant).

It is possible that observers use erroneous estimates of the parameters in Eq. (8) to arrive at the achromatic point. We refer to incorrectly estimated but consistently used parameters of the lighting model as Equivalent Illumination Models (EIM) following Brainard (1998) who used the term *equivalent illuminant* (also see Boyaci et al., 2003; Bloj et al., 2004).

As noted before, the chromaticities of the lights in this experiment vary only in the *blue–yellow* direction, hence $\hat{\Lambda}^B$ is the dependent variable of primary interest to us.

3.4. Maximum likelihood estimation and hypothesis testing

We use a maximum likelihood fitting procedure to estimate values of the lighting model parameters ($\hat{\pi}^R, \hat{\delta}^R, \hat{\pi}^G, \hat{\delta}^G, \hat{\pi}^B, \hat{\delta}^B, \hat{\Delta}, \hat{\varepsilon}_{11}, \hat{\varepsilon}_{22}$[8] and $\hat{\sigma}$ the observer's esti-

---

[8] $\hat{\varepsilon}_{22}$ is explicitly estimated only in 9D model fits.

mated standard deviation) that best accounted for each observer's data separately. Likelihoods were obtained by fitting relative red, blue and green data simultaneously. Table 1 shows results of estimated parameters for all observers.

For all fitting procedures we imposed the following restriction on the estimation algorithm: $\varphi_2 = -\varphi_1$, assuming that the two punctate sources will always be perceived to be placed symmetric over the line of sight of the observer. This was a reasonable assumption to make, as observer's raw data did not exhibit any asymmetry. As mentioned above, θ_i was assigned its true value, since it was never varied in the experiment and was neither directly nor indirectly assessed.

3.5. Testing the 4D subspace conjecture

Each observer's achromatic setting data (160°-condition) were fit twice, with the light field approximated with a 4D or 9D spherical harmonic subspace. The likelihoods of both fits were compared by means of a nested hypothesis test (Mood, Graybill, & Boes, 1974, p. 440), the null hypothesis being that observer's data would be fit equally well by both models.

We nested the hypothesis that $\hat{\varepsilon}_{22} = 0$ (4D subspace approximation) within a model in which $\hat{\varepsilon}_{22}$ was free to vary (9D subspace approximation). The log likelihood of the constrained model λ_0 (4D) must be less than or equal to that of the unconstrained model λ_1 (9D). Under the null hypothesis, twice the difference in log likelihoods is asymptotically distributed as a χ_1^2 variable, with degrees of freedom equal to the difference of freely varying parameters of unconstrained and constrained model.

The hypothesis that the 4D and 9D model fits are equally likely given the data was rejected for all observers (160°-condition, $\chi_1^2 \geqslant 19.6$, $p < .0001$). Fig. 8 shows observers' data overlaid with their estimated b-chromaticity function for both, 4D and 9D model fits.

The 4D model fit in Fig. 8 is a concave-down curve, which is different from the predicted concave-up curve in Fig. 7 (160°-condition). How can this discrepancy be explained? Above it was noted that it may be possible that if observers discounted a light field consisting of a single *blue* punctate source, located near $\varphi = 0°$ and a *yellow* diffuse source then the 4D subspace approximation would have the maximum in Λ^B at $\varphi = 0°$. While in Fig. 7 the 4D ideal curve was computed with 'true' light source chromaticities ($\pi^B = 0$, $\delta^B = 1$) we observe a reversal in the estimated chromaticities of punctate and diffuse source in the 4D model fits to the data ($\hat{\pi}^B > \hat{\delta}^B$) consistent with a *blue* punctate and a *yellow* diffuse source. Table 2 shows estimates of $\hat{\pi}^B, \hat{\delta}^B$ for the 4D and 9D model fits for all observers. Since the 4D light field model was rejected to account for observers' discounting behavior, the possibility that observers discount an erroneous light field consisting of a *blue* punctate and *yellow* diffuse

Table 1
Lighting model parameter estimates for 9D model fits

Veridical	$\max(L_P(\theta,\varphi))$ from $\hat{\varepsilon}_{11}, \hat{\varepsilon}_{22}$	$\hat{\Lambda}$	$\hat{\pi}^R$	$\hat{\delta}^R$	$\hat{\pi}^G$	$\hat{\delta}^G$	$\hat{\pi}^B$	$\hat{\delta}^B$
	$\pm 80°$	0.245	0.5	0	0.5	0	0	1
160°								
AS	($\pm 86.7°$	0.197) $p = .91$	(0.48	0.14	0.30	0.28	0.2	0.59)*
MS	($\pm 89.5°$	0.204) $p = .02$	(0.39	0.24	0.34	0.36	0.26	0.39)*
SHK	($\pm 87.4°$	0.267) $p = .17$	(0.45	0.4	0.30	0.24	0.36	0.45)*
IB	($\pm 87.7°$	0.160)*	(0.37	0.34	0.3	0.28	0.34	0.38)*
Veridical	$\pm 0°$	0.245	0.5	0	0.5	0	0	1
90°								
AS	($\pm 0°$	0.237) $p = .45$	(0.62	0.00	0.30	0.35	0.06	0.68)*
MS	($\pm 0°$	0.279) $p = .39$	(0.47	0.43	0.35	0.13	0.18	0.42)*
SHK	($\pm 0°$	0.247) $p = .99$	(0.58	0.2	0.24	0.35	0.18	0.42)*
IB	($\pm 0°$	0.220) $p = .99$	(0.4	0.52	0.36	0.20	0.24	0.23)*

Note: Values are obtained from combined maximum likelihood fits (geometric r-, g- and b-chromaticity functions) to the data. We tested the hypotheses that observers' estimates of lighting model chromaticities were equal to the veridical values. With a Bonferroni correction for 36 tests (4 observers, 9 parameters) the significant level corresponding to an overall Type I Error rate of 0.05 is 0.00139. All observers' lighting model chromaticity estimates were significantly different from the true chromaticities ($\chi^2_2 \geqslant 94.7$; $p < .0001$). Note there are two possible ways of interpreting these results, either in terms of misestimating the chromaticity of the illumination or in terms of misestimating the surface color of the test patch. In this experiment we cannot distinguish between these two possibilities. The deviations of the chromaticity parameters from veridical do not affect our main finding.

Estimates of punctate source azimuth φ_i shown in column two were obtained by finding the maxima in the approximation of $L_P(\theta,\varphi) = \sum_{n=0}^{2}\sum_{m=-n}^{n}\varepsilon_{nm}Y_{nm}(\theta,\varphi)$ for $\theta = 65°$ (see Appendix B), using estimates of the lighting coefficients $\hat{\varepsilon}_{11}, \hat{\varepsilon}_{22}$. Except for IB ($\chi^2_3 = 11.3$, $p < .0001$) all observers' estimates of the directional lighting model parameters ($\hat{\varepsilon}_{11}, \hat{\varepsilon}_{22}$, and $\hat{\Delta}$ combined) were not significantly different from the veridical values.
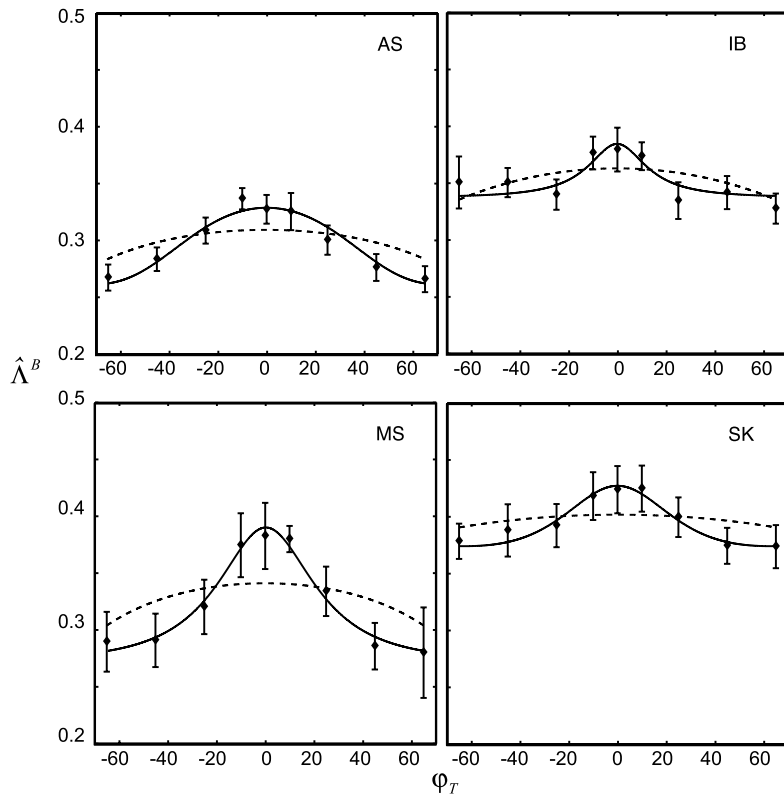

Fig. 8. Data and fits for relative blue, 9D vs 4D lighting model expansion, 160°-condition. We plot Λ^B as a function of test patch orientation φ_T. The figure shows observers' data (diamond symbols, error bars are $\pm 2SE$ of the mean, which corresponds approximately to the 95% confidence interval). The figure illustrates clearly that the data is fit better when an observer's EIM is approximated with a 9D spherical harmonic subspace (solid line) than with a 4D harmonic subspace (dashed line), indicating that the visual system can resolve at least directional variation in the illumination up to a 9D subspace. All fits are obtained by means of maximum likelihood estimation as described in the text.

Table 2
$\hat{\pi}^B$ and $\hat{\delta}^B$ in 9D and 4D model fits

	9D		4D	
	$\hat{\pi}^B$	$\hat{\delta}^B$	$\hat{\pi}^B$	$\hat{\delta}^B$
160°				
AS	.17	.68	.38	.01
MS	.26	.39	.43	.03
SHK	.27	.36	.48	.02
IB	.34	.38	.52	.02

Note: Contrasted are b-chromaticities of punctate $\hat{\pi}^B$ and diffuse $\hat{\delta}^B$ sources as obtained from 9D and 4D model fits. The 4D fit suggests a single *blue* punctate source at $\varphi_i = 0°$. However, the 4D model does not account for observers' data (see Fig. 8).

source has been rejected implicitly. We furthermore verified this finding by comparing likelihoods of the fits where light sources were modeled with their full spatial frequency spectrum ($p < .0001$).

3.6. Comparing discounting strategies across conditions

The 90°-condition was included to check whether observers adopted a single discounting strategy regardless of illumination condition. Employing the 9D light field model we nested the hypothesis that observers' parameters in the 90°-condition were equal to those estimated from the 160°-condition, allowing only $\hat{\sigma}$ and $\hat{\Delta}$ to vary freely, within a model in which all parameters ($\hat{\varepsilon}_{11}$, $\hat{\varepsilon}_{22}$, $\hat{\pi}^R$, $\hat{\delta}^R$, $\hat{\pi}^G$, $\hat{\delta}^G$, $\hat{\pi}^B$, $\hat{\delta}^B$, $\hat{\Delta}$, $\hat{\sigma}$) were free to vary. The hypothesis was rejected for all observers ($\chi_8^2 \geq 83.4$, $p < .0001$). Fig. 9 shows achromatic settings and model fits for all observers in the 90°-condition.

As predicted we found no significant difference between the 4D and 9D model fits in this condition (90°-condition, $\chi_1^2 \leq 1.3$, $p > .11$).

4. Discussion

Given previous results in matte surface color estimation (such as Boyaci et al., 2003, 2004, 2006; Ripamonti et al., 2004; Maloney et al., 2005; also te Pas & Pont, 2006) where the light field had only one maximum of intensity we proposed a conjecture that the visual system discounts directional variation in the illumination only up to a 4D spherical harmonic subspace. Data from the present experiment where the light field had two distinct maxima of intensity allowed us to reject this 4D subspace conjecture showing that the visual system discounts directional variation in the light field outside the 4D spherical harmonic subspace (Table 1).
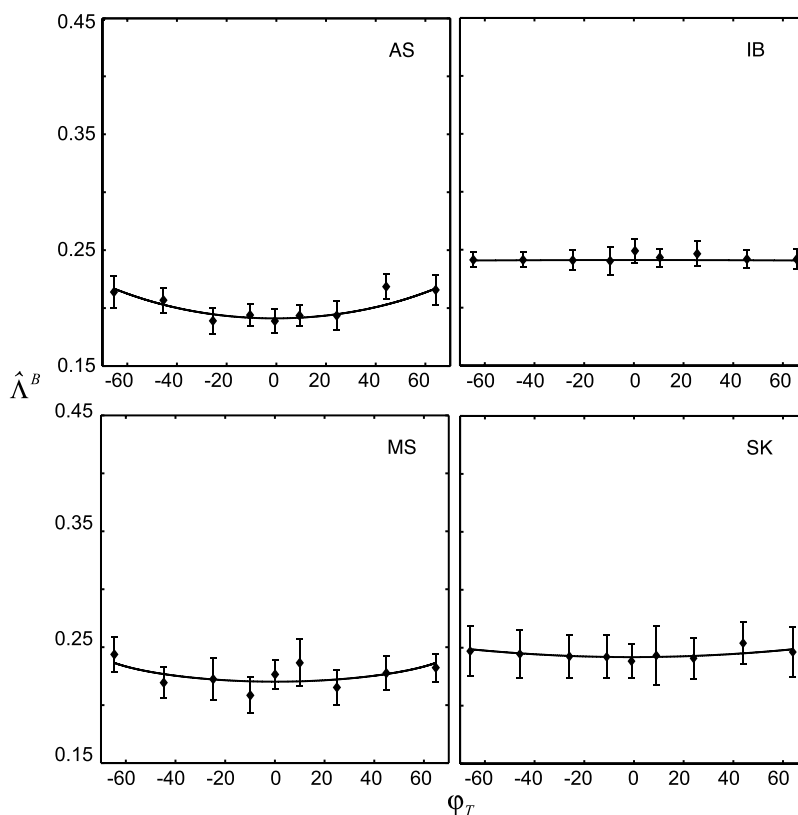

Fig. 9. Data and fits for relative blue 90°-condition. Symbols and SE are defined as in Fig. 9. Λ^B is plotted as a function of test patch orientation φ_T. Note that the 4D and 9D curve fits for this condition are virtually indistinguishable.

We tested two illumination conditions 90° and 160° to probe whether observers discount the given lighting geometry. As detailed above (also see Fig. 7) the minimum of the geometric b-chromaticity function is a rough indicator of the punctate source azimuth (or the location of the effective maximum of light source intensity). If observers discounted the lighting geometry in each case (90° and 160°) approximately correct, we expected to find a shift in the minimum of the observer's geometric b-chromaticity function, from near 0° (in the 90° condition) to approximately ±80° (in the 160° condition). This is what we found (Table 1).

While it is certainly possible to choose an alternative light field model and to parameterize the illumination in our scenes as, for example, a function of numbers of punctate sources, their intensities and directions, we would like to emphasize that doing the modeling and analysis in the frequency domain directly shows the relevance of our results to the theoretical findings by Basri and Jacobs (2001) and Ramamoorthi and Hanrahan (2001) which is a key contribution that this study makes.

4.1. High-frequency directional variation in the illumination

If we want to assess the role of higher spatial frequencies in the light field (beyond the 9D spherical harmonics subspace) we need to study a visual task whose successful completion would directly depend on those light field components: for example the assessment of surface specularity. Fig. 1 illustrates this dependency: the mirror sphere reflecting a blurry environment map appeared to be made from a rather different, rougher, material than a mirror. Theoretically, a rough(er) sphere reflecting the un-blurred environment map and a mirror sphere reflecting the blurred version will result in identical images. In fact, it would be impossible for the visual system to disentangle material property (in this case specular blur) from the spatial structure of the light field (Adelson, 2001; te Pas & Pont, 2005; Pont & te Pas, 2006)—if no disambiguating context (e.g. here the image background which indicates which environment map served as illumination) is provided. Recently, researchers (e.g. Pont & te Pas, 2006; te Pas & Pont, 2005, 2006; Khang, Koenderink, & Kappers, 2006) have begun to systematically investigate how shape, material and spatial properties of the light field interact.

4.2. Non-Lambertian surfaces

In this article we developed an experimental test of a hypothesis concerning the visual system's ability to discount spatially complex illumination in estimating the surface color of Lambertian surfaces. We emphasize that the theory and experiments made use of physical constraints on how light interacts with Lambertian surfaces, therefore we cannot generalize our results to materials with other kinds of BRDF's, for example specular surfaces.

4.3. Naturally occurring light fields

The light fields used in this experiment were artificial and carefully chosen to test the 4D subspace conjecture. However, one advantage of framing the experimental question in the frequency domain as we do is that the 4D and 9D components of any light field are always well defined.

4.4. Estimating the equivalent illumination model (EIM)

In the present study we developed an EIM which predicts observers' performance with ideal knowledge about the physical parameters in the scene. Deviations from ideal performance can be meaningfully explained by observers' mis-estimation of EIM parameters. How is the EIM formed? Boyaci et al. (2006) examined what sources of information the visual system uses to estimate the location of neutral punctate light sources and their relative intensities in a scene. They refer to these information sources as "cues to the lighting model". Examining three potential cues—specular highlights, cast shadows and shading of matte surfaces—each individually and all three combined, they find that observers can use each of the cues in isolation when estimating albedo of a matte test patch embedded in the scene, and that observers combine multiple available cues into an estimate that is more reliable than any of the cues in isolation (effective cue combination, Boyaci et al., 2006; see also Oruç, Maloney, & Landy, 2003). The scenes employed in this experiment provided the observer with a rich lighting model cue environment, and we believe that the data in our experiment can be explained parsimoniously by employing the EIM. In general, we do not claim that the parameterization of the illumination we have chosen is the one used by the visual system, only that whatever parameterization the visual system uses is captured by the EIM parameters which we estimated in this experiment. In that sense we simply describe human performance and possible implementations of computational models of that performance (Marr, 1982).

Alternatively, one may like to reject the notion of the EIM in favor of explanations based on specific heuristic strategies. For example observers could simply make use of a cue based exclusively on shading, then their performance is consistent with the EIM and, if the visual system assigns a high weight to the shading cue in these scenes, their performance as predicted by the EIM would closely resemble that predicted by this heuristic. There is no inconsistency between the claim that the visual system builds an EIM and uses it in estimating surface color and the claim that its operation can be summarized as the coordinated activity of a small number of heuristic algorithms. We do note that the visual

system must make use of heuristics other than the one just described since it can partially discount changes in effective illumination due to changes in surface orientation even when surface shading cues are not present in scenes (Boyaci et al., 2006) or when surfaces with identical orientations are assigned different lightnesses by observers (Snyder et al., 2005).

In interpreting the outcome of our experiment we are assuming that observers are in fact adjusting the test patch so that it appears achromatic on each trial. We can also consider the possibility that observers are not, in fact, making achromatic settings but instead simply copying the chromaticities of matte surfaces at similar orientations in the scene. When observers complete a trial using the copying heuristic, the test surface need not appear achromatic; the observer has deliberately failed to carry out the achromatic setting task as instructed. We reject this possibility for the following reasons. First, the observers were not aware of the hypothesis under test. It is difficult to see how observers would fail to carry out the task correctly or why they would have any motive to deviate from the instructions given them. Second, after the experiment, we explicitly asked observers what 'strategies' they had used to make their settings. None reported to have used the available matte (or other specific) objects in the scenes as references. In fact, some observers reported that they had tried to ignore the scene altogether when making their achromatic settings. Third, we (the experimenters) experience these effects and are in a position to critically evaluate whether our settings are achromatic. For these reasons, we believe that observers are not employing conscious strategies that lead to settings inconsistent with the instructions they have been given.

4.5. The 9D conjecture

Our results demonstrate that, under our experimental conditions, the human visual system is able to discount the directional variation in the local light field beyond what can be captured by a 4D spherical harmonic subspace. We can say a bit more. Let H denote the set of light fields for which the visual system can compensate and $\mathbb{P}_2H$ the components of these light fields projected into the 9D subspace. Then we can state our conclusions as $\mathbb{P}_2H$ is not contained in the 4D subspace. But what can we say about the stronger claim that $\mathbb{P}_2H$ is precisely the 9D subspace? We refer to this claim as the "9D conjecture".

We have not established this stronger claim. However, we end by stating necessary and sufficient conditions that it be true. If $\mathbb{P}_2H$ is (1) rotationally invariant and (2) closed under superposition of light fields, then $\mathbb{P}_2H$ is precisely the 9D subspace. This claim is easily shown: subspaces are closed under superposition and the only rotationally-invariant subspace within the 9D subspace that contains

more than the 4D subspace is the 9D subspace itself (MacRobert & Sneddon, 1967).

The requirement of rotational invariance is effectively the requirement that rotations of the entire scene or, equivalently, changes in the viewpoint of the observer, do not affect performance. We emphasize that this is an empirical question.

Closure under superposition implies that if the visual system can compensate for each of two light fields, it can compensate for their superposition. Some of the alternative models considered above would not be closed under superposition—e.g. the superposition of light fields based on lighting models that can represent up to two punctate sources and a diffuse could have as many as four punctate sources. Of course, the components of these scenes in the 9D subspace, however, sum to another light field in the 9D subspace.

We are left with following conclusions. $\mathbb{P}_2H$ is not included in the 4D subspace. If $\mathbb{P}_2H$ is not identical to the 9D subspace then it must fail either rotation invariance or superposition. If it is coextensive with the 9D subspace, then the visual system has evolved to compensate for all of the complexity in the light field that affects the appearance of Lambertian surfaces.

Appendix A. Separable light fields and the special tristimulus case

We consider the special case where the light field is separable in direction (θ, φ) and wavelength λ: $E_{\mathbf{p}}(\theta, \varphi, \lambda) = E_{\mathbf{p}}(\theta, \varphi)\,G(\lambda)$. This would be the case in scenes where all light sources have the same relative spectral power distribution (that is, a monochrome world) and all surfaces and other physical media that could affect the spectral power distribution of light are spectrally unselective (hence, the spectrum of light never changes when being absorbed and re-emitted by the surfaces). We refer to $E_{\mathbf{p}}(\theta, \varphi)$ as the *spatial component* of the light field and $G(\lambda)$ as the *spectral component*.

Next, we expand the space of possible light fields to those that can be represented by the weighted sum of three light primaries ('guns', such as those of a CRT monitor),
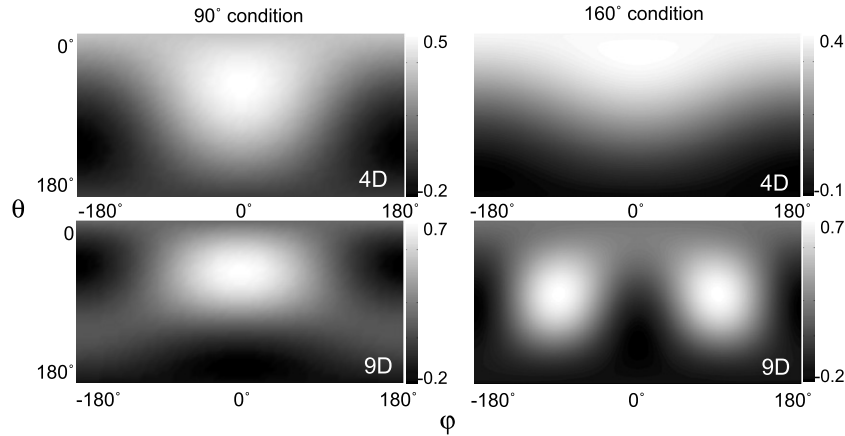
Fig. 10. 4D and 9D harmonic expansion of L_P. Illustrated is the expansion of L_P for one channel coding channel activity on a black-white scale shown next to each expansion. If P_1 and P_2 are placed sufficiently far from each other, as in the *160°-condition* in the 9D expansion of L_P will have two maxima of intensity. Conversely, if P_1 and P_2 are located closer together, as in the *90°-condition*, 4D and 9D spherical harmonic expansion of L_P will both have only one maximum of intensity.

$$E(\theta, \varphi, \lambda) = \sum_{k=1}^{3} E^k(\theta, \varphi) G^k(\lambda). \tag{9}$$

For any tri-stimulus light field we can expand each of the light fields $E^k(\theta, \varphi), k = 1, 2, 3$ separately as a spherical harmonic series and examine its low-pass components with respect to either the 4D or 9D subspace,

$$\mathbb{P}_N E(\theta, \varphi, \lambda) = \sum_{k=1}^{3} \mathbb{P}_N E^k(\theta, \varphi) G^k(\lambda), \quad N = 1, 2. \tag{10}$$

The light field components $E^k(\theta, \varphi), k = 1, 2, 3$ corresponding to the different 'guns'. $\mathbb{P}_N E^k(\theta, \varphi)$, $k = 1, 2, 3$ need not be the same. We will use the symbols, R, G, B as synonyms for the values $k = 1, 2, 3$, respectively. Thus $E^B(\theta, \varphi)$ is a synonym for $E^3(\theta, \varphi)$, $G^B(\lambda)$, a synonym for $G^3(\lambda)$, etc.

Appendix B. The low-pass Lambertian kernel

The light received and re-emitted by a Lambertian surface can be written as $R(\theta', \varphi') \alpha(\lambda) G(\lambda)$ where $\alpha(\lambda)$ is the surface reflectance function and

$$R(\theta', \varphi') = \int \int E(\theta, \varphi) \langle \mathbf{U}(\theta, \varphi), \mathbf{U}(\theta', \varphi') \rangle \mathrm{d}\varphi \mathrm{d}\theta, \tag{11}$$

where $\langle \mathbf{U}(\theta, \varphi), \mathbf{U}(\theta', \varphi') \rangle$ is the inner product of the unit surface normal vector $\mathbf{U}(\theta', \varphi')$ and the unit vector $\mathbf{U}(\theta, \varphi)$ indicating the direction of a 'light ray'. We refer to $R(\theta', \varphi')$ as the directional component of the *Lambertian response function*. Given the linearity of the integral in Eq. (4) and the low-pass filtering property of the Lambertian surfaces as shown by Basri and Jacobs, if $E(\theta, \varphi)$ is replaced by $\mathbb{P}_2 E(\theta, \varphi)$ then the response function of the Lambertian surface

$$R(\theta', \varphi') \approx \mathbb{P}_2 R(\theta', \varphi'), \tag{12}$$

is (almost) unaffected.

Appendix C. Spherical harmonic expansion of the lighting model

The light field in our experiment consisted of two *yellow* punctate light sources (90° or 160° apart) and a *blue* diffuse (directionally uniform) light source.

The diffuse light (we denote it $L_D(\theta, \varphi) = c$, where c is a constant) can be expanded using Eq. (1). With uniform lighting the lighting coefficients (see Eq. (2)) simply become

$$\varepsilon_{Dnm}^B = c \int_{\theta=0}^{\pi} \int_{\varphi=0}^{2\pi} Y_{nm}(\theta, \varphi) \sin \theta \mathrm{d}\theta \mathrm{d}\varphi$$
$$= c\sqrt{4\pi} \delta_{n0} \delta_{m0}, \tag{13}$$

using the orthonormality property of spherical harmonics (for an overview of the properties of spherical harmonics see MacRobert & Sneddon, 1967; Arfken, 1985a, 1985b; or Byerly, 1959). Consequently the expansion of L_D is a constant, namely c.

The punctate sources are the sole source of directional heterogeneity in our particular illumination setup. To simplify the harmonic expansion, we will treat each punctate source as a delta function that has a nonzero value only at location (θ_1, φ_1) and (θ_2, φ_2), respectively. We write the expansion of the punctate sources

$$L_P(\theta, \varphi) = \sum_{n=0}^{\infty} \sum_{m=-n}^{n} \varepsilon_{Pnm} Y_{nm}(\theta, \varphi) \tag{14}$$

where

$$\varepsilon_{Pnm} = \int_{\theta=0}^{\pi} \int_{\varphi=0}^{2\pi} L_P(\theta, \varphi) Y_{nm}(\theta, \varphi) \sin \theta \mathrm{d}\theta \mathrm{d}\varphi. \tag{15}$$

As mentioned above, L_P consists of two delta functions at $\delta(\theta_1, \varphi_2)$ and $\delta(\theta_2, \varphi_2)$. Because light fields combine linearly we can write $\varepsilon_{Pnm} = \varepsilon_{P_1nm} + \varepsilon_{P_2nm}$. Using

$$\int_{\theta=0}^{\pi} \int_{\varphi=0}^{2\pi} \delta(\theta', \varphi') Y_{nm}(\theta, \varphi) \sin \theta \, d\varphi \, d\theta$$
$$= Y_{nm}(\theta', \varphi') \sin \theta' \qquad (16)$$

(see Basri & Jacobs, 2003) we get

$$\varepsilon_{P_1 nm} = Y_{nm}(\theta_1, \varphi_1) \sin \theta_1, \qquad (17)$$
$$\varepsilon_{P_2 nm} = Y_{nm}(\theta_2, \varphi_2) \sin \theta_2.$$

Consequently, $\varepsilon_{Pnm} = Y_{nm}(\theta_1, \varphi_1)\sin \theta_1 + Y_{nm}(\theta_2, \varphi_2)\sin\theta_2$. Fig. 10 shows the 4D and 9D expansion of L_P for each illumination condition. For the remainder we refer to the lighting coefficients of the punctate source expansion ε_{Pnm} simply as ε_{nm}.

References

Adelson, E., & Bergen, J. (1991). The plenoptic function and elements of early vision. In M. S. Landy & A. J. Movshon (Eds.), *Computational models of visual processing* (pp. 3–20). Cambridge: MIT Press.

Adelson, E. H., (2001). On seeing stuff: The perception of materials by humans and machines. In: Rogowitz, B. E., Pappas, T. N. (Eds.), *Proceedings of the SPIE, Vol 4299. Human Vision and Electronic Imaging* VI. (pp. 1–12).

Arfken, G. (1985a). Spherical harmonics. In *Mathematical methods for physicists* (3rd ed., pp. 680–685). Orlando: Academic Press.

Arfken, G. (1985b). Integrals of the products of three spherical harmonics. In *Mathematical methods for physicists* (3rd ed., pp. 698–700). Orlando: Academic Press.

Basri, R., Jacobs, D., (2001). Lambertian reflectance and linear subspaces. In: *International Conference on Computer Vision II*, pp. 383–390.

Basri, R., & Jacobs, D. (2003). Lambertian reflectance and linear subspaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 25*, 218–223.

Bloj, M. G., Kersten, D., & Hurlbert, A. C. (1999). Perception of three-dimensional shape influences color perception through mutual illumination. *Nature, 402*, 877–879.

Bloj, M., Ripamonti, C., Mitha, K., Hauck, R., Greenwald, S., & Brainard, D. H. (2004). An equivalent illuminant model for the effect of surface slant on perceived lightness. *Journal of Vision, 4*, 735–746, http://journalofvision.org/4/9/6/.

Boyaci, H., Maloney, L. T., & Hersh, S. (2003). The effect of perceived surface orientation on perceived surface albedo in binocularly-viewed scenes. *Journal of Vision, 3*, 541–553, http://journalofvision.org/3/8/2/.

Boyaci, H., Doerschner, K., & Maloney, L. T. (2004). Perceived surface color in binocularly-viewed scenes with two light sources differing in chromaticity. *Journal of Vision, 4*, 664–679, http://journalofvision.org/4/9/1/.

Boyaci, H., Doerschner, K., & Maloney, L. T. (2006). Cues to an equivalent lighting model. *Journal of Vision, 6*, 106–118, http://journalofvision.org/6/2/2/.

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

Byerly, W. E. (1959). *An elementary treatise on Fourier's series, and spherical, cylindrical, and ellipsoidal harmonics, with applications to problems in mathematical physics*. New York: Dover, pp. 195–218.

Cabral, B., Max, N., & Springmeyer, R. (1996). Bidirectional reflection functions from surface bump maps. In: *Proceeding of SIGGRAPH 87*, pp. 273–281.

Doerschner, K., Boyaci, H., & Maloney, L. T. (2004). Human observers compensate for secondary illumination originating in nearby chromatic surfaces. *Journal of Vision, 4*, 92–105, http://www.journalofvision.org/4/2/3/.

Dutre, P., Bekaert, P., & Bala, K. (2003). *Advanced global illumination*. Natick: A K Peters, pp.31–41.

D'Zmura, M. (1991). Shading ambiguity: Reflectance and illumination. In M. S. Landy & A. J. Movshon (Eds.), *Computational models of visual processing* (pp. 187–207). Cambridge: MIT press.

Fleming, R. W., Dror, R. O., & Adelson, E. H. (2003). Real-world illumination and the perception of surface reflectance properties. *Journal of Vision, 3*, 347–368, http://journalofvision.org/3/5/3/.

Flock, H. R., & Freedberg, E. (1970). Perceived angle of incidence and achromatic surface color. *Perception and Psychophysics, 8*, 251–256.

Gershun, A. (1939). The light field (Translation by Moon, P., & Timoshenko G.). *Journal of Mathematics and Physics, 18*, 51–151.

Gilchrist, A. L. (1977). Perceived lightness depends on spatial arrangement. *Science, 195*, 185–187.

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

Haralick, R. M., & Shapiro, L. G. (1993). *Computer and robot vision* (Vol. 2). Reading, MA: Addison-Wesley.

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

Hochberg, J. E., & Beck, J. (1954). Apparent spatial arrangement and perceived brightness. *Journal of Experimental Psychology, 47*, 263–266.

Kaufman, J. E., & Christensen, J. F. (1972). *IES lighting handbook: The standard lighting guide* (5th ed.). New York: Illuminating Engineering Society.

Khang, B. G., Koenderink, J. J., & Kappers, A. M. L. (2006). Perception of illumination direction in images of 3-D convex objects: Influence of surface materials and light fields. *Perception, 35*, 625–645.

MacRobert, T. M., & Sneddon, I. N. (1967). *Spherical harmonics: An elementary treatise on harmonic functions, with applications* (3rd ed.). Oxford, UK: Pergamon Press.

Maloney, L. T., Boyaci, H., & Doerschner, K., 2005. Surface color perception as an inverse problem in biological vision. In: Bouman, C. A., Miller, E. L. (Eds.), *Computational Imaging III Proceedings of SPIE-IS& T Electronic Imaging, SPIE, Vol. 5674*.

Maloney, L. T. (1999). Physics-based approaches to modeling surface color perception. In K. R. Gegenfurtner & L. T. Sharpe (Eds.), *Color vision: From genes to perception* (pp. 387–422). Cambridge, UK: Cambridge University Press.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco, CA: W.H. Freeman.

Mood, A., Graybill, F. A., & Boes, D. C. (1974). *Introduction to the theory of statistics* (3rd ed.). New York: McGraw-Hill.

Nicodemus, F. E., Richmond, J. C., Hsia, J. J., Ginsberg, I. W., & Limperis, T. (1977). Geometric considerations and nomenclature for reflectance. In: *Monograph 161*. National Bureau of Standards (US).

Nimeroff, J. S., Simoncelli, E., & Dorsey, J. (1994). Efficient Re-rendering of naturally illuminated environments. *Fifth Eurographics Workshop on Rendering*, 337–359.

Oruç, I., Maloney, L. T., & Landy, M. S. (2003). Weighted linear cue combination with possibly correlated error. *Vision Research, 43*, 2451–2468.

Pont, S. C., & te Pas, S. F. (2006). Material-illumination ambiguities and the perception of solid objects. *Perception, 5*(10), 1331–1350.

Ramamoorthi, R., Hanrahan, P., (2001). An efficient representation for irradiance environment maps. *SIGGRAPH 01*. New York: ACM Press, pp. 497–500.

Ripamonti, C., Bloj, M., Hauck, R., Kiran, K., Greenwald, S., Maloney, S. I., et al. (2004). Measurements of the effect of surface slant on perceived lightness. *Journal of Vision, 4*, 747–763, http://journalofvision.org/4/9/7/.

Snyder, J. L., Doerschner, K., & Maloney, L. T. (2005). Illumination estimation in three-dimensional scenes with and without

specular cues. *Journal of Vision, 5*, 863–877, http://journalofvision.org/5/10/8/.

Teo, P. C., Simoncelli, E. P., & Heeger, D. J., (1997). Efficient linear re-rendering for interactive lighting design. Technical Report. UMI Order Number: CS-TN-97-60, Stanford University.

te Pas, S. F., & Pont, S. C., (2005). Comparison of material and illumination discrimination performance for real rough, real smooth and computer generated smooth spheres. In: Bülthoff, H., Troscianko, T. (Eds.), *Proceedings APGV 2005*, ACM SIGGRAPH, USA, pp. 75–81.

te Pas, S. F., & Pont, S. C. (2006). Illumination discrimination under varying complexity of shape and light sources [Abstract]. *Journal of Vision, 6*, 469a, http://journalofvision.org/6/6/469/.

Ward, G., "The RADIANCE Lighting Simulation and Rendering System", *Computer Graphics*, July 1994.

Wyszecki, G., & Stiles, W. S. (1982). *Color science; concepts and methods, quantitative data and formulae* (2nd ed.). New York: Wiley.

Yang, J. N., & Maloney, L. T. (2001). Illuminant cues in surface color perception: Tests of three candidate cues. *Vision Research, 41*, 2581–2600.