Perceived surface color in binocularly viewed scenes with two light sources differing in chromaticity

Huseyin Boyaci
Department of Psychology and Center for Neural Science, New York University, New York, NY, USA

Katja Doerschner
Department of Psychology, New York University, New York, NY, USA

Laurence T. Maloney
Department of Psychology and Center for Neural Science, New York University, New York, NY, USA

We examined the effect of perceived orientation on the perceived color of matte surfaces in rendered three-dimensional scenes illuminated by a blue diffuse light and a yellow punctate light. On each trial, observers first adjusted the color of a matte test patch, placed near the center of the scene, until it appeared achromatic, and then estimated its orientation by adjusting a monocular gradient probe. The orientation of the test patch was varied from trial to trial by the experimental program, effectively varying the chromaticity of the light mixture from the two light sources that would be absorbed and reemitted by a neutral test patch. We found that observers' achromatic settings varied with perceived orientation but that observers only partially discounted orientation in making achromatic settings. We developed an equivalent illuminant model for our task in which we assumed that observers discount orientation using possibly erroneous estimates of the chromaticities of the light sources and/or their spatial distribution. We found that the observers' failures could be explained by two factors: errors in estimating the direction to the punctate light source and errors in estimating the chromaticities of the two light sources. We discuss the pattern of errors in estimating these factors across observers.

Keywords: color perception, color constancy, rendering, binocular disparity

Introduction

In a field, on an ordinary sunny day, the light impinging on each object is a mixture of direct sunlight, light from sky and clouds, and light absorbed and reemitted by other objects in the scene. Even when the contribution from other objects is neglected, the light absorbed and reemitted by any surface is composite, a mixture that changes with each passing cloud.

In Figure 1, we illustrate how light is absorbed and reemitted by a Lambertian (matte) surface illuminated by a punctate source and a diffuse source. The spectral power distribution of the punctate source at each wavelength λ is denoted by $E_P(\lambda)$ and the spectral power distribution of the diffuse light by $E_D(\lambda)$. The angle between the punctate light direction and the surface normal $\mathbf{n}$ is denoted by θ, and the angle between the surface normal and the direction to the viewer is denoted by ν. In the Lambertian model, the intensity of emitted light does not depend on the direction to the viewer, so long as the viewer and the light source are on the same side of the surface. When this condition is satisfied, the intensity of the light reflected from an achromatic Lambertian surface at any wavelength λ is given by

$$E(\lambda) = \alpha\left(E_P(\lambda)\cos\theta + E_D(\lambda)\right), \quad (1)$$

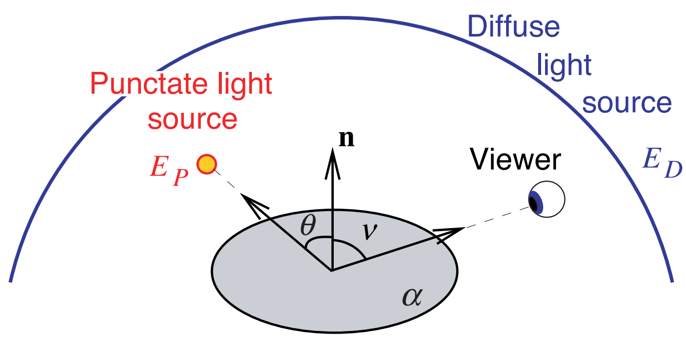

Figure 1. An achromatic Lambertian surface. The intensity of the light that is absorbed by a surface is proportional to the cosine of the angle between the rays of light from the punctate light source and the surface normal, $\mathbf{n}$. A uniform diffuse light source contributes with a constant amount to the total intensity on the surface. Light absorbed by a Lambertian surface is reemitted uniformly in all directions; the intensity of light reaching the viewer does not depend on viewing angle ν, as long as the surface is visible. An achromatic Lambertian surface reflects all incoming light equally independent of its wavelength. For such an ideal surface, the ratio of the reflected light to the incoming light is determined by a wavelength-independent reflectance.

doi:10.1167/4.9.1 Received September 15, 2003; published August 23, 2004 ISSN 1534-7362 © 2004 ARVO

where α is the wavelength independent albedo (reflectance) of the achromatic surface. The spectral power distribution of the effective illuminant $E(\lambda)$ is a weighted mixture of the spectral power distributions of the diffuse and punctate sources and the mixture changes as the orientation of the test patch changes. If the diffuse and punctate sources differ in chromaticity, then the chromaticity of the mixture will also change as a function of angle.

Our goal is to examine how human observers perceive surface colors in scenes with composite lighting. In particular, we will test whether they correctly discount surface orientation in estimating the surface color of matte surface patches in simulated scenes illuminated by a yellow punctate source and a blue diffuse source.

Previous research

A number of researchers have investigated how the spatial arrangement of lights and surfaces in a three-dimensional (3D) scene affects the perception of lightness or color of a particular matte surface in the scene.

Gilchrist (1977, 1980; Gilchrist et al., 1999) examined scenes where the intensity of illumination varied with depth and found that perceived depth affected the perceived lightness of achromatic surfaces. The experimental setup was composed of two rooms that differed in illumination, and were connected by a doorway. Observers viewed this construction through a pinhole. In the two conditions of this experiment, the target patch, whose lightness was to be matched, would be perceived to be located as coplanar with either a brightly illuminated far wall or a dimly illuminated near wall. The actual position and brightness of the test patch were not altered. The test patch was perceived to be white in the near wall condition and almost black in the far wall condition. Gilchrist argued that perceived lightness depends on the relationship between the target and regions with which it is seen as coplanar (the *coplanar ratio hypothesis*).

Boyaci, Maloney, and Hersh (2003) set out to determine whether perceived orientation affects perceived lightness. Their stimuli were binocularly viewed, computer-rendered scenes illuminated by a neutral punctate light source and a neutral diffuse light source. Observers first estimated the orientation of an achromatic test patch with Lambertian reflectance properties and then matched its perceived albedo ('lightness') to a reference scale. Their results clearly showed that observers systematically discounted the perceived orientation of a surface when estimating its albedo. In fact, observers' performances closely matched the predictions of the Lambertian model.

How dramatically the perception of 3D shape can influence perceived surface reflectance was demonstrated by Bloj et al. (1999). They used a chromatic version of the Mach card; one side of the card was painted magenta and the other white, the magenta side casting a pinkish gradient on the white areas. With binocular disparity as the only cue to shape, observers viewed the card (1) in its actual concave shape and (2) through a pseudoscope, which reversed the disparities in left and right eye so that the card appeared to be convex. Bloj and colleagues found significant evidence that observers incorporate information about the shape of the object into their estimates of surface color: The color of the white side of the card was judged by observers to be more pinkish in the apparently convex condition than in the actual concave condition.

Bloj and colleagues studied the effect of 3D shape on color perception for only two viewing conditions. What happens for different angles between the mutually illuminated surfaces? According to the physics of light, the strength of the mutual illumination will depend on the angle between the emitting and receiving surface. Doerschner, Boyaci, and Maloney (2004) showed that observers systematically discount the angle between a brightly colored surface and an adjacent light gray surface, when setting the color of the latter to be achromatic.

Yang and Shevell (2003) investigated surface color appearance in simulated scenes illuminated by two punctate light sources differing in chromaticity. The scenes were rendered and presented binocularly, and each consisted of two side-by-side rooms separated by an opaque partition oriented along the line of sight. The back wall of each room was covered with lozenge-shaped specular objects. The light sources were placed so that each primarily illuminated one of the rooms and only secondarily the other (thus the primary light source for each room was the secondary for the other). Yang and Shevell (2003) varied the contribution of the secondary light source to a room by varying the height of the partition dividing the two rooms. When the partition was at its maximum height, each room was illuminated only by its primary light source.

Color constancy was greatest when each room was illuminated exclusively by its primary light source and decreased with increasing admixture of the secondary. In these scenes, each light source created a distinct specular highlight on each of the specular lozenges that it illuminated. Yang and Shevell (2003) perturbed the chromaticity of these highlights to show that they were effective as cues to the chromaticity of the light reaching each of the room, confirming earlier results (Yang & Maloney, 2001; Yang & Shevell, 2002; Maloney & Yang, 2003).

The results of Yang and Shevell (2003) suggest that the human visual system can partially discount the relative contributions of two light sources in a scene, at least when there is sufficient information in the scene to permit estimation of the chromaticities and spatial distribution of the light sources. In the experiment we describe next, we will provide considerable visual information about the chromaticity of the light sources and the direction to the punctate light source. Before describing the experiment, we consider its possible outcomes and how we might interpret each.

Possible outcomes

One possibility is that the observer may simply make settings that do not vary with the perceived orientation of the test patch. We would conclude that the observer's visual system does not take orientation into account in estimating color, at least not in the scenes we employ. This outcome would be consistent with the previous literature just described, although we might wonder why the visual system discounts the effect of orientation in estimating lightness (Boyaci et al., 2003), but not in estimating color in general.

Alternatively, the observer could discount the effect of perceived orientation perfectly or nearly so. In the next section, we develop a model that allows us to predict the pattern of settings that correspond to this outcome. We will compare the observer's settings to these predictions. To achieve such perfect constancy despite changing orientation, the observer must effectively take into account the spatial distribution of the light sources, the direction to the punctate light source, the relative intensities of the two sources, and the chromaticities of the two sources. We refer to these quantities together as a lighting model. It is evidently implausible that the observer estimates the lighting model accurately in scenes such as ours with one of the light sources not even visible.

The results of our experiment will allow us to reject both the hypothesis of no constancy and the hypothesis of perfect constancy just described. A third and more realistic possibility lies between these two extremes. It is that the observer makes achromatic settings that are incorrect, but that would be correct for a different lighting model than the one illuminating the scene. A pattern of discounting consistent with incorrect estimates of the lighting model is referred to by Brainard (1998) as an equivalent illuminant model. Brainard's equivalent illuminant model consisted of an estimate of the chromaticity of a single light source. He finds that observers' deviations from color constancy can be parsimoniously explained by the assumption that they have misestimated the chromaticity of the illuminant. We will consider a more complex equivalent illumination model appropriate for a combination of punctate and diffuse sources. We fit this model to observers' data and determine whether we can reject it. We will discover that it provides a parsimonious summary of observers' performance in the experiment described next, and we will discuss the pattern of estimates of the lighting model parameters across observers.

Experiment

Introduction

We asked observers to carry out two tasks on each trial. They first set a Lambertian test patch to be achromatic (Helson & Michels, 1948), and then set a gradient probe to indicate its orientation. The test patch was embedded in a scene illuminated by a mixture of a blue diffuse light source and a yellow punctate light source. We varied the orientation of the test patch with respect to the punctate light source from trial to trial, thereby varying the chromaticity of the composite illuminant striking the test patch. We sought to determine whether observers compensated for changes in perceived orientation (and illuminant chromaticity) in their achromatic settings. Note that the orientation task is only included as a way of estimating the observer's perceived orientation. The focus of the experiment is on the achromatic setting task.

Methods

Stimuli

The stimuli were computer-rendered 3D-complex scenes composed of simple objects with different shapes (such as spheres and boxes), and with various reflectance properties (such as shiny, matte, and transparent). Each scene contained a matte test patch at the center. The scenes were rendered with the Radiance software package (Larson & Shakespeare, 1996). Each scene was rendered twice, from slightly different viewpoints corresponding to the positions of the observer's eyes. A stereo image pair for a typical scene is shown in Figure 2. The other objects in the scene were varied randomly from trial to trial and were included as possible cues to the spatial distribution and chromaticities of the light sources (see Yang & Maloney, 2001).

Apparatus

The experimental apparatus was a Wheatstone stereoscope (Figure 3). The left and right images of each stereo pair were presented to the corresponding eye of the observer on two 21" Sony Trinitron Multiscan GDM-F500 monitors placed to the left and right of the observer. The screens on these monitors are close to physically flat, with less than 1 mm of deviation across the surface of each monitor.

The stereoscope was contained in a box 124 cm on a side. The front face of the box was removed, and the observer sat there in a chin/head rest. Two small mirrors were placed directly in front of the observer's eyes. These mirrors reflected the images displayed on the left and right monitor to the corresponding eye of the observer. The interior of the box was coated with black-flocked paper (Edmund Scientific) to absorb stray light. Only the stimuli on the screens of the monitors were visible to the observer. The casings of the monitors and any other features of the room were hidden behind the non-reflective walls of the enclosing box. Additional light baffles were placed near the observer's face to prevent light from the screens reaching the observer's eyes directly. The optical distance from each of the observer's eyes to the corresponding computer screen was 70 cm. To minimize any conflict between binocular disparity and accommodation depth cues, the center of the test patch was rendered to be exactly 70 cm in front of the

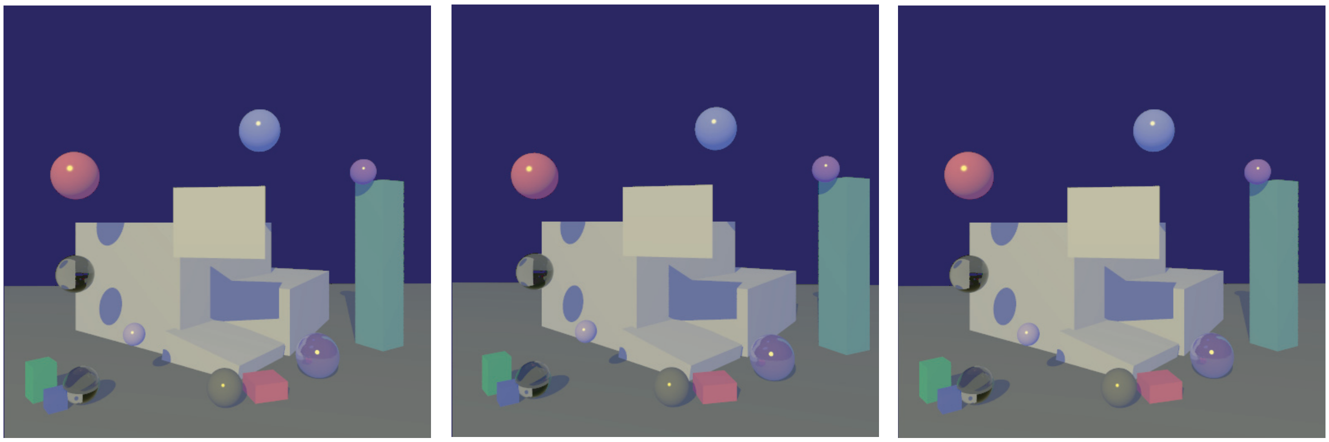

Figure 2. Sample stimulus. The stimuli were computer-rendered images of complex scenes. Each scene was rendered twice, from slightly different viewing points corresponding to the two eyes of the observers. The reader can fuse the left and center images (uncrossed fusion) or the center and right images (crossed-fusion). The virtual scene was illuminated by a yellow punctate light source positioned behind the observer either to his or her left or right and a blue diffuse light source. A test patch was located at the center of the scene. The test patch was the closest object to the observer in the scene (except the floor). By doing so, we eliminated any possible secondary illumination of the test patch by light emitted from other surfaces in the scene. Various additional objects, with a variety of surface reflectance properties (matte, shiny, or transparent) were included in the scene. The locations and properties of these objects were varied at random from trial to trial in the experiment.

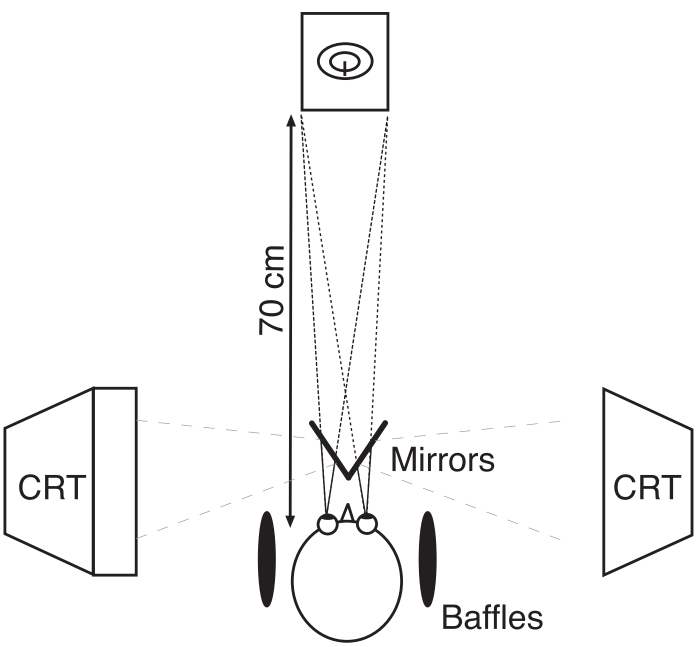

Figure 3. Apparatus. Stimuli were displayed in a computer-controlled Wheatstone stereoscope. The left and right images of a stereo pair were displayed on the left and right monitors of the stereoscope. These images were viewed by means of small mirrors placed in front of the observer's eyes. Baffles placed to either side of the observer's head blocked stray light from the monitors that might otherwise reach the eyes. In the fused image, the test surface appeared approximately 70 cm in front of the observer. This distance was also the optical distance to the screens of the two computer monitors, minimizing any mismatch between accommodation cues and other depth cues.

observer. The monocular fields of view were 55 × 55 deg. The observer's eyes were approximately at the same height as the center of the scene being viewed, which was also the position of the center of the test patch.

Spatial coordinate system and spatial arrangement

We used a spherical coordinate system (ψ,φ,r) to specify a simulated scene (Figure 4). This coordinate system has the origin at the center of the test patch. The spherical coordinate system (ψ,φ,r) and the Cartesian coordinate system underlying it are explained in the figure legend.

Color coordinate system

We will describe all light sources and the light radiating from surfaces as weighted mixtures of three abstract primary lights referred to as red, green, and blue (RGB). For convenience, the spectra of these lights coincide with those of the corresponding guns of the monitors, and the three primaries can be thought of as linearized versions of the guns, for that is what they are. We measure the intensities of these three primaries in arbitrary units proportional to their luminance, chosen so that a mixture of the three lights with equal intensities appears roughly achromatic to most observers. We denote the intensities by E^R, E^G, and E^B, respectively, and refer to the tristimulus coordinates (E^R, E^G, E^B) that describe the light at a particular location on the monitor as an RGB code.[1] In making an achromatic setting, the observer in effect selects the RGB code for the test patch that makes it appear to be an achromatic surface, as described in more detail below. We report the u'v' chromaticities (Wyszecki & Stiles, 1982, p. 165) of the guns (and, therefore, of the primaries) in the "Calibration" section below.

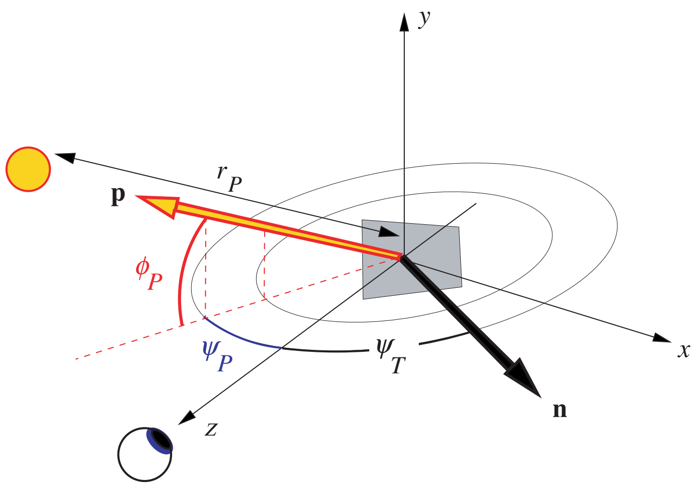

Figure 4. Spatial coordinate system. We used a spherical coordinate system based on a Cartesian coordinate system to describe the geometry of the test patch and the punctate source. Both coordinate systems had their origins in the center of the test patch. Note that the center of the test patch was always in the same location throughout the experiment. In the Cartesian system (*x*,*y*,*z*), the *z*-axis fell along the observer's line of sight, the *y*-axis was vertical, and the *x*-axis horizontal as shown. In the spherical coordinate system, a point in the 3D space is denoted by three numbers (ψ,φ,r): r is the distance of the point from the origin. ψ is the angle between the observer's line of sight (*z*-axis) and the projection of the point on the horizontal plane (*xz*-plane), φ is the angle between the horizontal plane and the line connecting the origin and the point. The position of the punctate source is denoted by (ψ_P,φ_P,r_P). The unit vector in the direction of the punctate source is denoted by $\mathbf{p}$, the unit vector normal of the test patch is denoted by $\mathbf{n}$. The angle between the punctate direction and surface normal is calculated with the law of cosines: $\cos\theta = \mathbf{n}\cdot\mathbf{p}$.

Realism and rendering

Typical rendering packages used in computer graphics represent spectral information about surfaces and light sources by 3D vectors, often referred to as RGB codes. When light with RGB code (E^R,E^G,E^B) strikes a Lambertian surface with RGB code (S^R,S^G,S^B), the light emitted from the surface is assigned the RGB code ($E^R\times S^R,E^G\times S^G,E^B\times S^B$), scaled by a factor that takes into account the orientation of the surface with respect to the light (see the discussion leading up to Equation 1). Yang and Maloney (2001; Maloney, 1999) point out that this rendering interpretation ("the RGB heuristic" in Maloney, 1999) does not always lead to accurate simulation of light-surface interactions.

However, the scenes that we use are designed to avoid the limitations of typical rendering packages. First, we define an achromatic Lambertian surface to be one that multiplies the RGB code of the light that it absorbs and reemits by a constant factor that depends on the albedo of the surface and the direction from which the light arrives. If we assign this neutral surface the RGB code (α,α,α), then typical rendering packages will simulate light-surface interaction correctly. So long as our chromatic lights interact with only neutral surfaces, the resulting RGB codes assigned to the light reemitted will be accurate. There are other surfaces in our scenes that are rendered, but the RGB codes of these surfaces are assigned at random and change from trial to trial. Consequently, errors in rendering, due to using the RGB heuristic, are of no consequence. The intended random color assigned to a surface is just replaced by a different random color.

We can derive three equations that break Equation 1 into three RGB-code components. For B, we have,

$$E^B(\theta) = \alpha\left(E_P^B\cos\theta + E_D^B\right) \tag{2}$$

and there are two analogous equations for R and G, respectively. When the test patch is achromatic, each of the components of the RGB-code of the light emitted by the test patch is the same weighted mixture of the corresponding components of the two light sources.

Calibration

Look-up tables were used to correct the nonlinearities in the gun responses and to equalize the display values on the two monitors. The tables were prepared after direct measurements of the luminance values of each gun on each monitor with a Pritchard PR-650 spectrometer. The maximum total luminance achievable on either screen was $114\ cd/m^2$. To test the linear additivity for a monitor, first we measured the isolated spectrum of each gun alone, set to about half of its maximum intensity. Then we measured the spectra of each pair of guns simultaneously set to half of their maximum intensities and compared it to the sum of the isolated spectra for each gun in the pair. Last, we measured the spectrum with all three guns set to half of their maximum intensity and compared it to the sum of the isolated spectra for all three guns. We plot the results of this last test in Figure 5 for both monitors. The red, green, and blue solid lines are the isolated spectra, the gray solid line is the sum of the three isolated spectra, and the black dashed line is the measured spectra when all three guns were simultaneously set to half of their maximum intensities. The curves agree to within 7% or better at each point in the spectrum, for both monitors. The test of additivity for pairs of guns also agreed within 7% or less. The u'v' chromaticity coordinates (Wyszecki & Stiles, 1982, p. 165) for the three primaries are: red (.409, .519), green (.117, .565) and blue (.157, .196) for the left monitor, and red (.430, .528), green (.115, .564), and blue (.160, .189) for the right monitor. The u'v' chromaticity coordinate for the mixture of all three guns at half intensity was (.176, .460) for the left monitor and (.172, .455) for the right monitor.

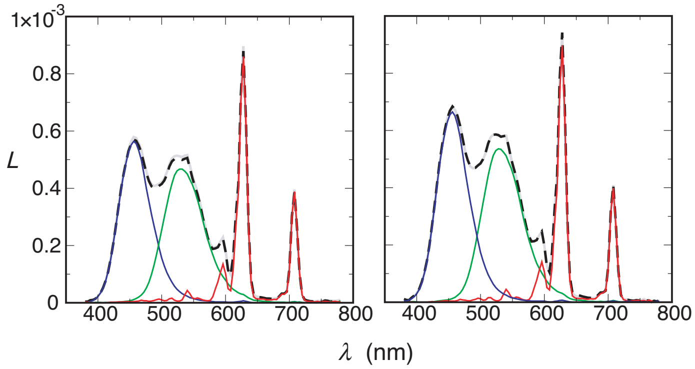

Figure 5. Monitor guns: tests of additivity. We measured the red, green, and blue gun luminances of our 21" Sony Trinitron Multiscan GDM-F500 monitors with a Pritchard PR-650 spectrometer. We first set each gun to approximately half of its maximum possible intensity (pixel value 200) with the other two guns set to zero intensity. Luminance at that intensity is plotted separately across wavelength for each gun. The red, green, and blue solid lines in (A) (left monitor) and (B) (the right monitor) correspond to the red, green, and blue guns. We then computed the sum of the three measured primaries (plotted as a gray solid line) and measured the luminance with all three guns simultaneously set to the same intensities (black dashed lines).

Light sources

A yellow punctate light source and a mostly blue diffuse light source illuminated the scenes. The punctate light source was not directly visible to the observer. The RGB code of the punctate light source is denoted (E_P^R, E_P^G, E_P^B) and that of the diffuse light source is denoted (E_D^R, E_D^G, E_D^B). For convenience, let $E_P = E_P^R + E_P^G + E_P^B$ and $E_D = E_D^R + E_D^G + E_D^B$. To specify the chromaticities of the two light sources and their relative strengths, we define the following set of parameters: $\pi^B = E_P^B / E_P$, the b-chromaticity[2] of the punctate source, $\delta^B = E_D^B / E_D$, the b-chromaticity of the diffuse source, and $\Delta = E_D / E_P$, the diffuse-punctate ratio. The r- and g-chromaticities of the yellow punctate source were always equal, as were those of the diffuse source. The values used in rendering were $\pi^B = 0$, $\delta^B = 0.66$, and $\Delta = 0.37$. In other words, the punctate source had no blue component ($\pi^B = 0$), the diffuse source was mostly blue ($\delta^B = 0.66$), and the ratio of the intensity of the diffuse source to the intensity of the punctate source was 0.37 ($\Delta = 0.37$). The punctate source was always behind and above the observer, and either to his RIGHT or to his LEFT at $(\psi_P, \varphi_P, r_P) = (\pm 15^\circ, 30^\circ, 670\,cm)$ ($\psi_P = +15^\circ$ for RIGHT, $\psi_P = -15^\circ$ for LEFT; Figure 4). The position of the punctate source was varied only from session to session, but in a single session, its position was kept constant. The punctate source was sufficiently far from the test patch so as to treat its light rays collimated. The vector $\mathbf{p} = (\cos\varphi_P \sin\psi_P, \sin\varphi_P, \cos\varphi_P \cos\psi_P)$ is a unit vector pointing from the test patch toward the punctate light source (Figure 4).

Test patch

Each scene contained a test patch at the center, which was rendered as an achromatic Lambertian surface with an albedo of 0.55 (the observer never saw this patch and we used it only to verify that the output of the Radiance program agreed with the predictions of Equation 2 and Equation 3; the parts of the left and the right images corresponding to this patch were replaced by a uniform test patch before the images were shown to the observer). The initial RGB code of the substituted test patch was chosen at random before each trial. The test patch could be displayed with either a rotation in only the ψ direction ($\psi-rotation$) or in only the φ direction ($\varphi-rotation$). The test patch measured 4.8 cm by 3.6 cm; its center was always 70 cm away from the observer along the observer's line of sight. The orientation of the test patch was specified by (ψ_T, φ_T), and its surface normal was $\mathbf{n} = (\cos\varphi_T \sin\psi_T, \sin\varphi_T, \cos\varphi_T \cos\psi_T)$. After a $\psi-rotation$ ($\varphi_T = 0$), ψ_T could take any of the values $\{-60^\circ, -45^\circ, -15^\circ, 15^\circ, 45^\circ\}$ when the punctate source was on the LEFT, and any of the values $\{-45^\circ, -15^\circ, 15^\circ, 45^\circ, 60^\circ\}$ when the punctate source was on the RIGHT. After a $\varphi-rotation$ ($\psi_T = 0$), φ_T could take any of the values $\{0^\circ, 15^\circ, 30^\circ, 45^\circ, 60^\circ\}$. Figure 6 shows a schematic drawing of the two kinds of rotations.

The test patch floated in space in the middle of the scene. It was closer to the observer than all other objects

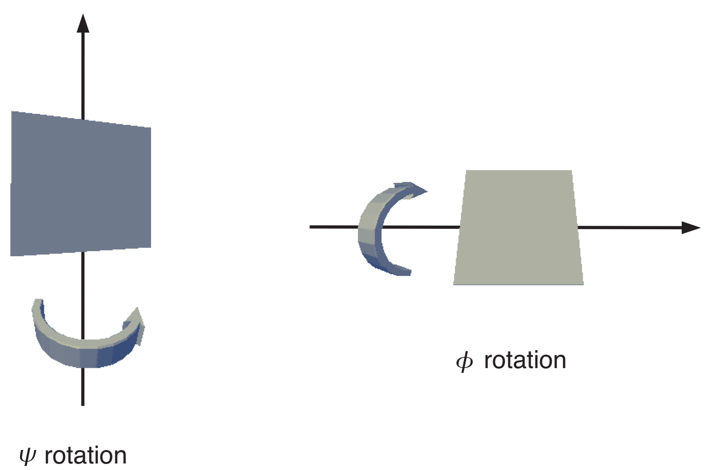

Figure 6. Orientations. In each trial the test patch could appear with one of 10 orientations. Five of them were rotations of the fronto-parallel test patch in the ψ direction; the other five were rotations in the φ direction. After a ψ rotation, the orientation of the surface could be one of $\psi_T = \{-60^\circ, -45^\circ, -15^\circ, 15^\circ, 45^\circ\}$ when the punctate source was positioned on the left or $\psi_T = \{-45^\circ, -15^\circ, 15^\circ, 45^\circ, 60^\circ\}$ when the punctate source was positioned on the right, with $\varphi_T = 0$. After a φ rotation, the orientation of the test patch could be one of $\varphi_T = \{0^\circ, 15^\circ, 30^\circ, 45^\circ, 60^\circ\}$ with $\psi_T = 0$.

and sufficiently high above the floor so that we could eliminate possible mutual illumination effects.

Angle of incidence

The cosine of the angle between the surface normal of the test patch and the direction to the punctate source is found by employing the law of cosines

$$\cos\theta = \mathbf{n}\cdot\mathbf{p} = \cos\varphi_T \cos\varphi_P \sin\psi_T \sin\psi_P + \sin\varphi_T \sin\varphi_P + \cos\varphi_T \cos\varphi_P \cos\psi_T \cos\psi_P \ , \tag{3}$$

or

$$\cos\theta = \begin{cases} \cos\varphi_P \cos(\psi_P - \psi_T) & \text{for } \psi - rotation \\ \sin\varphi_T \sin\varphi_P + & \text{for } \varphi - rotation \\ \cos\varphi_T \cos\varphi_P \cos\psi_P & \end{cases} \ , \tag{4}$$

because $\varphi_T = 0$ for a $\psi - rotation$ and $\psi_T = 0$ for a $\varphi - rotation$.

Tasks

The observer carried out two tasks on each trial.

Achromatic setting task

The observer first adjusted the chromaticity of the test patch such that it looked achromatic. He or she did this by varying the b-, r- and g-chromaticities without altering the total intensity of the test patch (Figure 7a). The observer used the arrow keys on the keyboard to perform this task. Pressing the up arrow key increased the b-chromaticity while decreasing the r- and g-chromaticities by the same amount; the down arrow had the opposite effect. Pressing the right arrow key increased the r-chromaticity while decreasing the g-chromaticity. The left arrow had the opposite effect.

Orientation task

The second task was to estimate the orientation of the test patch (the independent variable) by adjusting a stick-and-circle gradient probe superimposed at the center of the test patch (Figure 7b). The orientation of the probe was controlled by moving a computer mouse. The probe was presented monocularly to the right eye, and the observer had only one degree of freedom on each trial: if it was a $\psi - rotation$ trial, the probe could rotate only in the ψ direction; if it was a $\varphi - rotation$, it could rotate only in the φ direction. Observers reported no difficulty with setting the probe and were unaware that it was visible in only the right eye. Once the observer was satisfied with the setting, he or she clicked the left button on the computer mouse to finalize the task. The purpose of this task was to control for the possibility that observers' perceptions of orientation of the test patch were so different from its actual orientation that it would affect the interpretation of the results. We assume that the observer is using the same cues to test patch orientation during the achromatic setting task as in the orientation task. In these scenes, these cues include binocular disparity and linear perspective. See Landy, Maloney, Johnston, and Young (1995) for a review of cue-combination models for depth and slant.

Software

The experimental software was written by us in the C language. We used the X Window System, Version 11R6 (Scheifler & Gettys, 1996) running under Red Hat Linux 6.1 for graphical display. The computer was a Dell 410 Workstation with a Matrox G450 dual head graphics card and a special purpose graphics driver from Xi Graphics that permitted a single computer to control both monitors. We use the open source physics-based rendering package Radiance (Larson & Shakespeare, 1996) to render the left and right images that comprised the stereo pair for a given virtual scene. The output of the rendering described above was a stereo image pair with floating point RGB codes for each pixel. We translated the output relative intensity values to 24-bit RGB codes, correcting for nonlinearities in the monitors' responses as described above.

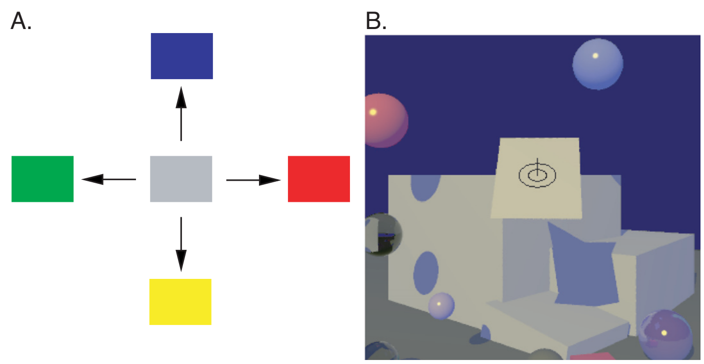

Figure 7. Tasks. On each trial, observers completed two tasks. A. The first one was the achromatic setting task: Observers adjusted the chromaticity of the test patch until it appeared achromatic ("it looked as if it were cut out of an achromatic [gray] piece of paper.") They adjusted the color of the test patch by pressing the arrow keys on a computer keyboard. Pressing the "up" arrow key increased the blue content, while decreasing the yellow (red + green) content by the same amount. Pressing the "down" arrow had the opposite effect. Pressing the "right" arrow increased the red content and decreased the green content by the same amount; pressing the "left" arrow increased the green content and decreased the red content by the same amount. B. The second task was to estimate the orientation of the test patch. Observers indicated the orientation of the test patch by adjusting a monocular gradient probe (presented to the right eye only). The probe consisted of two concentric circles and a stick placed at the center of the circles perpendicular to them. The observer's task was to set the probe such that the stick was perpendicular to, and the circles were tangent to, the test patch.

Procedure

The observers repeated each of the 20 conditions (10 test patch orientations, 2 punctate source positions) of the experiment 20 times for a total of 400 trials. The experiment was split into four sessions, each with 100 trials. In a single session the position of the punctate source (LEFT or RIGHT) was kept constant. The order of presentation was randomized. The observers were allowed to perform a few trials before the actual experiment started, until they were completely comfortable with both tasks. The experiment was paced by the observer. Usually observers completed different sessions on different days and each session took less than an hour.

Observers

Four observers completed the experiment. All were experienced psychophysical observers who were unaware of the purpose of the experiment. One other observer was excused after the first session. She had difficulty doing the task and spent more than 3 hours to finish a single session (which usually took other observers under an hour).

Instructions to the observer

For the color task, we asked the observer to adjust the color of the test patch such that it looked as if it were cut out of an achromatic or gray piece of paper. For the orientation task, the observers were simply instructed to move the mouse until the probe's circles were in the plane of the test patch and the stick was perpendicular to it.

Geometric chromaticity functions

To quantify the observers' perception and compare it with the model predictions, we define

$$\Lambda^B(\theta) = \frac{E^B(\theta)}{E(\theta)} = \frac{E_P^B \cos\theta + E_D^B}{E_P \cos\theta + E_D} \tag{5}$$

as the **geometric b-chromaticity function**. In Equation 5, $E(\theta)$ is the total intensity of the light emitted from the test patch. $E^B(\theta)$ is the blue component of the RGB code of the light emitted from the test patch, as defined earlier. The last term in Equation 5 is gotten by substituting Equation 1 and Equation 2 into the middle term. Equation 5 is the relative intensity of the blue primary in the light emitted from the test patch. We define a geometric g-chromaticity function $\Lambda^G(\theta)$ and a geometric r-chromaticity function $\Lambda^R(\theta)$, analogously, and refer to them collectively as geometric chromaticity functions. Note that when the light sources have the same b-, r- and g-chromaticities, the geometric chromaticity functions are all constant, independent of θ.

Achromatic setting

Suppose that the observer views a matte test patch in a scene illuminated by a combination of blue diffuse and yellow punctate sources. The angle between the normal to the test patch and the punctate light direction is θ. The observer is asked to adjust the chromaticity of the test patch without changing the total intensity until it looks achromatic. We denote this achromatic setting as a function of θ, by $\left(\hat{E}^R(\theta), \hat{E}^G(\theta), \hat{E}^B(\theta)\right)$. Note that his setting is always constrained so that $\hat{E}^R(\theta) + \hat{E}^G(\theta) + \hat{E}^B(\theta) = \text{constant}$, and it is convenient to express the achromatic setting in terms of chromaticities.

We define the observer's geometric blue b-chromaticity function by

$$\hat{\Lambda}^B(\theta) = \frac{\hat{E}^B(\theta)}{\hat{E}(\theta)} . \tag{6}$$

The observer's geometric r- and g-chromaticity functions are defined similarly. If the observer were perfectly color constant, then $\left(\hat{E}^R(\theta), \hat{E}^G(\theta), \hat{E}^B(\theta)\right)$ would coincide with $\left(E^R(\theta), E^G(\theta), E^B(\theta)\right)$, $\hat{\Lambda}^B(\theta)$ would be identical to $\Lambda^B(\theta)$, $\hat{\Lambda}^G(\theta)$ identical to $\Lambda^G(\theta)$, and $\hat{\Lambda}^R(\theta)$ identical to $\Lambda^R(\theta)$.

By means of the color adjustment task just described, we can measure the observer's geometric chromaticity functions and compare them to the theoretical ones for an achromatic Lambertian surface in the scenes we employ.

Analysis and results

Orientation settings

We first tested whether changes in the orientation of the test patch in ψ and φ direction had an effect on observers' orientation settings by separate ANOVAs for each observer. We rejected the hypothesis that the mean orientation setting did not vary with orientation for all observers, for both directions ($p < .0001$ in both ψ and φ directions). With the exception of subject MM in the ψ direction, we found no significant interaction between perceived test patch orientation and punctate light source position (LEFT or RIGHT) for both directions (ψ direction: $p = .206$, .637, and .304 for BH, MD, and RG, respectively, $p = .01$ for MM; φ direction: $p = .852$, .39, .07, and .928 for BH, MD, MM, and RG, respectively). This implies that for all but one observer the position of the light source (LEFT or RIGHT) had no significant effect on how observers made their orientation settings.

Figure 8 shows one observer's (BH) mean settings when the punctate source was positioned on the left. For each observer, we regressed the observer's mean orientation settings on the true orientation settings separately in both the ψ and φ ($\hat{\psi}_T$ vs. ψ_T and $\hat{\varphi}_T$ vs. φ_T) directions. We have plotted the best-fitting regression lines to BH's settings in Figure 8. We report the regression coefficients for similar fits for all observers in Table 1. The estimated regression coefficient a (intercept) is in units of degrees; the regression coefficient b (slope) is unitless. We report p values for

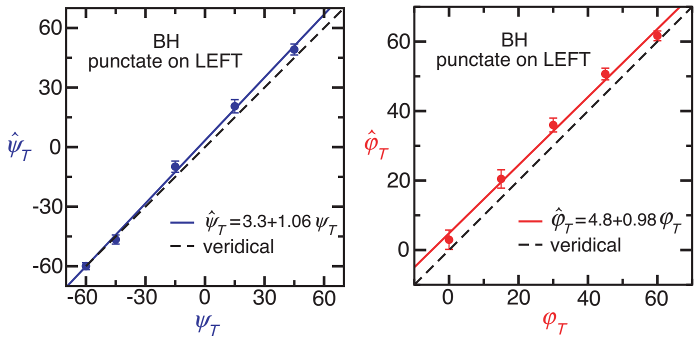

Figure 8. Results: orientation settings. This figure shows the orientation settings of observer BH when the punctate source was on the left. The graph on the left is for rotations of the test patch in the ψ direction, the graph on the right is for rotations in the φ direction. The observer's mean settings $\hat{\psi}_T$ and $\hat{\varphi}_T$ for each angle ψ_T and φ_T are represented by circles (blue for ψ_T, red for φ_T). The error bars represent 95% confidence intervals. The solid lines are the best linear fits, and the regression coefficients are given in the legends. All other observers' results are similar to BH's. Although there were deviations from veridical, those deviations were not large. The regression coefficients for all observers are given in Table 1.

hypothesis tests against the corresponding veridical value (0 for a, 1 for b). In the φ direction, slopes were significantly different from 1 for subject MD (punctate on RIGHT: $p < .001$) and subject MM (RIGHT: $p < .001$). All other subjects' slopes in this direction were not significantly different from 1 for both punctate source positions (LEFT: $p = .934$, .037, and .148 for BH, MM, and RG, respectively; RIGHT: $p = .834$, .01, and .125 for BH, MD, and RG, respectively). The intercepts in the φ direction were significantly different from 0 for all observers ($p < .001$) except for observer RG (LEFT: $p = .782$, RIGHT: $p = .042$).

In the ψ direction slopes were significantly different from 1 for all observers ($p < .001$), except observer BH (LEFT: $p = .009$, RIGHT: $p = .005$). The intercepts in the ψ direction were significantly different from 0 for all observers ($p < .002$) except MM (LEFT: $p = .398$, RIGHT: $p = .038$).

We conclude that observers' perceived orientation changes with orientation but that observers' perception of orientation is not veridical. We will use the observers' own estimates of orientation in analyzing the effect of orientation on achromatic settings. However, we note that it would not alter our conclusions in any important respect if we used the true orientations instead.

Achromatic settings

Observers' achromatic settings are the key dependent variable of our experiment. Consider a model observer who is effectively using Equation 5 to arrive at estimates of surface color appearance but whose estimates of some or all of the parameters in the equation were in error. The achromatic setting of the model observer would not match the geometric chromaticity function of Equation 5 and would instead exhibit characteristic failures due to errors in the parameter estimates. In this section, we examine the effect of errors in each parameter on the model observer's predicted performance.

An equivalent illuminant model

The chromaticities of the light sources in the rendered scenes differ only in the blue-yellow balance; therefore, we are primarily interested in the blue and yellow components of the observers' achromatic settings, and we first consider the geometric b-chromaticity function.

Equation 5 can be rewritten as

$$\Lambda^B(\theta) = \frac{\pi^B \cos\theta + \delta^B \Delta}{\cos\theta + \Delta}, \tag{7}$$

where the variables π^B, δ^B and Δ were defined above. An observer's visual system can compute what the b-

	Punctate source position: LEFT				Punctate source position: RIGHT			
	$\hat{\psi}_T = a + b\,\psi_T$		$\hat{\varphi}_T = a + b\,\varphi_T$		$\hat{\psi}_T = a + b\,\psi_T$		$\hat{\varphi}_T = a + b\,\varphi_T$	
	A	b	*A*	b	a	b	a	b
Veridical	*0*	*1*	*0*	*1*	*0*	*1*	*0*	*1*
BH	**3.33*** *p < .001*	**1.06** *p = .009*	**4.83*** *p < .001*	**0.98** *p = .934*	**3.18*** *p < .001*	**1.06** *p = .005*	**7.42*** *p < .001*	**0.96** *p = .834*
MD	**-2.33*** *p = .002*	**1.14*** *p < .001*	**4.72*** *p < .001*	**1.14*** *p < .001*	**-2.54*** *p = .002*	**1.1*** *p < .001*	**3.55*** *p = .003*	**1.08** *p = .01*
MM	**-0.55** *p = .398*	**0.83*** *p < .001*	**-6.59*** *p < .001*	**0.95** *p = .037*	**-1.76** *p = .038*	**0.76*** *p < .001*	**-4.6*** *p < .001*	**0.87*** *p < .001*
RG	**-4.79*** *p < .001*	**0.71*** *p < .001*	**0.38** *p = .782*	**0.95** *p = .148*	**-5.7*** *p < .001*	**0.66*** *p < .001*	**-1.68** *p = .042*	**0.97** *p = .125*

Table 1. Results: orientation setting regression coefficients. We fit a linear model to the orientation settings ($\hat{\psi}_T$ vs. ψ_T and $\hat{\varphi}_T$ vs. φ_T) separately for each observer. The estimated regression coefficient a (intercept) is in units of degrees, the regression coefficient b (slope) is unitless. We report *p* values for hypothesis tests against the corresponding veridical value (0 for *a*, 1 for *b*). We report exact *p* values when the values are larger than .001. With a Bonferroni correction for 16 tests per observer, the significant level corresponding to an overall Type I Error rate of 0.05 for each subject is .0031. Values whose corresponding *p* values fall below this cutoff are marked with an asterisk.

chromaticity of a gray surface should be if estimates of the parameters in Equation 7 are available. However, if the observer's estimates of the parameters are in error, the achromatic settings would differ from the predicted ones. Let $\hat{\pi}^B$, $\hat{\delta}^B$, $\hat{\Delta}$ and $\hat{\theta}$ denote the observer's estimates of the parameters in Equation 5, then

$$\hat{\Lambda}^B(\theta) = \frac{E^B_{set}\left(\hat{\theta}\right)}{E\left(\hat{\theta}\right)} = \frac{\hat{\pi}^B \cos\hat{\theta} + \hat{\delta}^B \hat{\Delta}}{\cos\hat{\theta} + \hat{\Delta}} \quad . \tag{8}$$

The observer's estimate of the angle of incidence $\hat{\theta}$ depends on his or her estimates of the orientation of the test patch $(\hat{\psi}_T, \hat{\varphi}_T)$ and the direction to the punctate source $(\hat{\psi}_P, \hat{\varphi}_P)$ through Equation 4. Note that the observers explicitly estimated the orientation of the test patch $(\hat{\psi}_T, \hat{\varphi}_T)$ by performing the orientation task.

Suppose that we hold the lighting conditions constant, in particular the parameters

$$\Theta = \left\{\psi_P, \varphi_P, \pi^B, \delta^B, \Delta\right\}, \tag{9}$$

and vary the orientation of the surface by varying ψ_T and φ_T as we do in the experiment. Figure 9a shows the geometric b-chromaticity function Λ^B plotted with respect to the angles ψ_T and φ_T assuming the veridical values of the lighting parameters, Θ. Now suppose that the estimates of the lighting parameters $\hat{\Theta} = \left\{\hat{\psi}_P, \hat{\varphi}_P, \hat{\pi}^B, \hat{\delta}^B, \hat{\Delta}\right\}$ are in error, what kind of distortions would those errors introduce? Misestimating the direction to the punctate source shifts both curves without much effect on their curvatures (Figure 9b). When the test patch is oriented such that it faces the punctate source as directly as possible, that is $\psi_T = \psi_P$ (after a ψ rotation) or $\varphi_T \simeq \varphi_P$ (after a φ rotation), it receives the maximum possible amount of light from the yellow punctate source. However the blue content of the mixture of light falling on it remains fixed, hence the b-chromaticity, Λ^B, assumes its minimum.

More rigorously, the extrema of the function Λ^B are found by taking its derivative with respect to ψ_T and φ_T, and then equating it to zero, which yields

$$\psi_T^* = \psi_P,$$
$$\varphi_T^* = \arctan\left(\frac{\tan\varphi_P}{\cos\psi_P}\right) \approx \varphi_P \quad (\text{for small } \psi_P). \tag{10}$$

(For $\psi_P = \pm 15$, $\varphi_T^* = 30.87°$, only slightly different from $\varphi_P = 30°$.) Note that ψ_T^* and φ_T^* correspond to minima for $\pi^B < \delta^B$, and to maxima for $\pi^B > \delta^B$ (see Equation 13 below). If the model observer misestimated the punctate source direction, his or her achromatic settings would reveal this because the corresponding geometric b-chromaticity function $\hat{\Lambda}^B$ would shift and have its minimum at roughly the estimated direction to the punctate source $(\hat{\psi}_P, \hat{\varphi}_P)$.

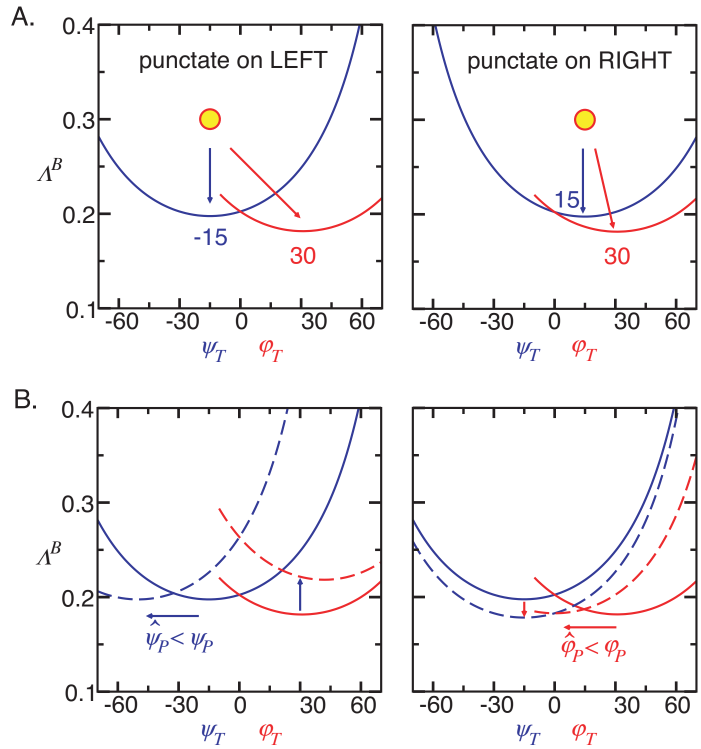

Figure 9. b-chromaticity settings. A. The "right answer." For an ideal Lambertian observer who uses the correct values of the lighting parameters in the right hand side of Equation 7, the geometric b-chromaticity function, Λ^B, calculated from the achromatic settings, would fall on the curves plotted in this figure. We plot Λ^B with respect to both ψ_T and φ_T on the same graphs. The blue solid line is the plot of Λ^B with respect to ψ_T, the red one is with respect to φ_T. The orientation of the test patch affects the geometric b-chromaticity as follows: as the achromatic test patch rotates away from the direction of the yellow punctate, it receives less and less yellow contribution (angle of incidence, θ, increases, $\cos\theta$ decreases (see Equation 1). However, the blue contribution from the diffuse source does not change with this rotation. Therefore, as the test patch rotates away from the punctate source, its b-chromaticity increases. Conversely, as the test patch rotates closer to the direction of the punctate source, its b-chromaticity decreases and reaches a minimum when it faces the punctate source directly. In the experiment, however, the orientation of the test patch could vary either only in the ψ direction or only in the φ direction. Hence Λ^B has minima at $\psi_T = \psi_P$ ($\varphi_T = 0$) and $\varphi_T \simeq \varphi_P$ ($\psi_T = 0$). B. Errors in estimating punctate light direction. What happens if the observers' estimates of the parameters in Equation 7 are in error? Suppose that the observer's estimate $\hat{\psi}_P$ of the direction to the punctate source is in error. If the observer made settings based on erroneous estimate, then the minimum of the blue curve would be at $\hat{\psi}_P$ instead of the correct value, ψ_P, as shown in the upper plots. An error in the estimate of ψ_P also affects the Λ^B versus φ_T curve. The pattern of shifts when $\hat{\psi}_P < \psi_P$ and for $\hat{\varphi}_P < \varphi_P$ are shown in (B). The patterns when $\hat{\psi}_P > \psi_P$ and $\hat{\varphi}_P > \varphi_P$ are just the reverse.

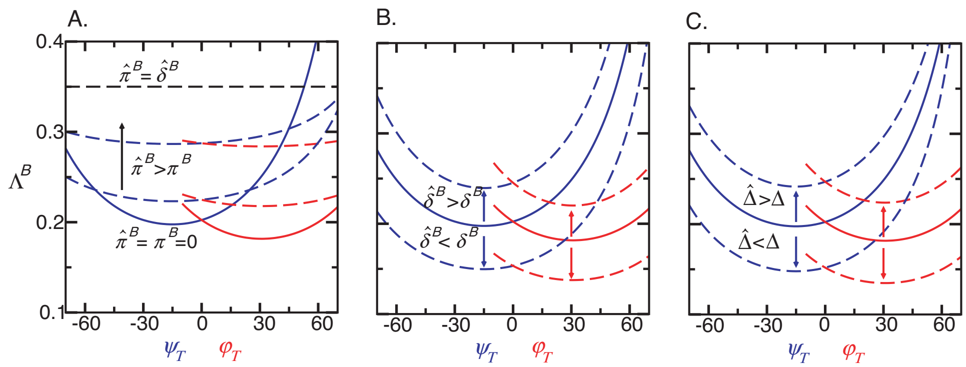

Figure 10. Errors in estimating other lighting parameters. Under- or overestimating the other three lighting parameters π^B, δ^B, and Δ lead to systematic changes in the Λ^B versus ψ_T and the Λ^B versus φ_T curves. They are discussed in the text.

Errors in estimating the parameters π^B, δ^B, and Δ would shift the curves up or down and increase or decrease their curvatures (Figure 10). If the observer's estimates of the b-chromaticity of the punctate source and diffuse source were the same ($\hat{\pi}^B = \hat{\delta}^B$), then the geometric b-chromaticity function would be a constant, because changing the orientation of the test patch would not affect the overall chromatic balance of the light reaching the patch. $\hat{\Lambda}^B$ would be constant also when $\hat{\Delta} = 0$ or $\hat{\Delta} \to \infty$, that is, if the observer estimates that the scene is illuminated either by only a punctate source or by only a diffuse source. However, because veridical values are such that $\pi^B \neq \delta^B$ and Δ is not 0 or infinity, should we find that the observer's geometric b-chromaticity function is constant, then the implication is that the observer does not discount the perceived orientation of the test patch for its color.

Figure 11 shows the empirical geometric b-chromaticity functions for all four observers. As mentioned above, if an

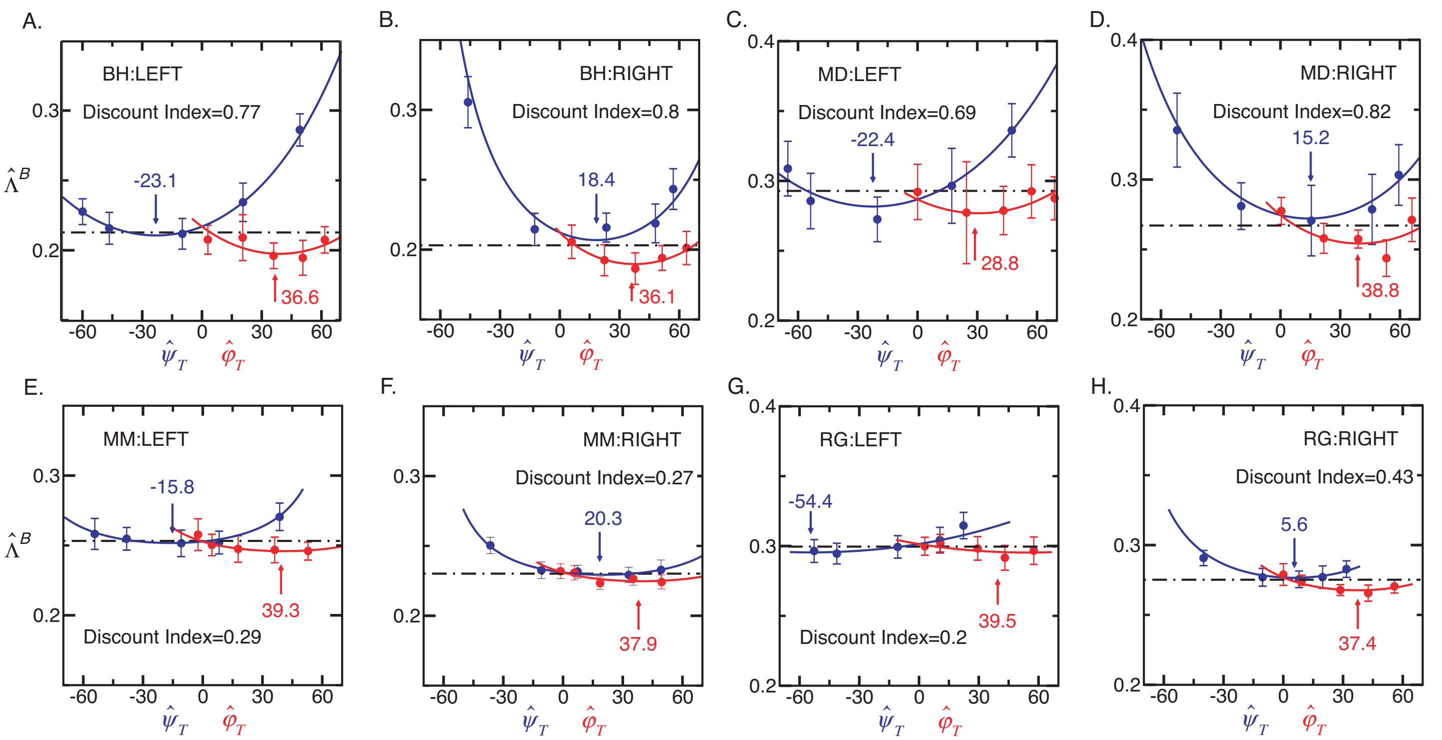

Figure 11. Results. b-chromaticity versus perceived orientation. We plot all four observers' geometric b-chromaticity functions with respect to perceived orientation. The blue filled circles are the mean values of the b-chromaticity of their achromatic settings at the mean perceived angle $\hat{\psi}_T$. The red ones are for $\hat{\varphi}_T$. The solid lines are the best-fitting curves according to the model described in the text. The small blue and red arrows point to the observers' estimates of the punctate source direction. Notice that observers' estimates of the direction to the punctate source are close to the correct values. The direction estimates were not significantly different from correct. Discounting indices are reported in legends (see "Geometric discounting index"). The flat black dashed lines correspond to observers' settings if they did not compensate at all for orientation. Clearly observers are discounting the perceived orientation for perceived color, but the degree of discounting is not as large as the model predicts. Error bars are ± 2 SE of the mean (approximately a 95% confidence interval).

observer were perfectly color constant, then his or her data would fall on the theoretical curve of the geometric b-chromaticity function Λ^B. On the other hand, if the observer were completely ignoring the orientation of the test patch in his or her achromatic judgment, then the ratio would be constant.

It is clear that observers take the orientation of the test patch into account and have some degree of color constancy, although the constancy is not perfect. A comparison of the patterns of data to the family of Λ^B curves in Figure 9 and Figure 10 suggests that observers make settings that are indistinguishable from those of a Lambertian color constant observer who discounts the perceived orientation for estimating color, but who does so using incorrect estimates of the lighting parameters Θ.

We use a maximum likelihood fitting procedure to estimate values of the lighting parameters $\Theta = \left\{\psi_P, \varphi_P, \pi^B, \delta^B, \Delta\right\}$ that best accounted for each observer's data separately (recall that we explicitly measured observers estimates of ψ_T and φ_T.) These estimates are reported in Table 2.

We tested the hypothesis that the observer's estimate of the punctate source direction $(\hat{\psi}_P, \hat{\varphi}_P)$ was equal to the veridical values by means of a nested hypothesis test (Mood, Graybill, & Boes, 1974, p. 440). We nested the hypothesis that $\hat{\psi}_P = \pm 15$ and $\hat{\varphi}_P = 30$ (their true values) within a model in which they were free to vary. We fit both models to the data by the method of maximum likelihood with other parameters allowed to vary freely. The log likelihood of the constraint model (denoted by λ_0) must be less than or equal to that of the unconstraint model (denoted by λ_1). Under the null hypothesis, twice the difference in log likelihoods is asymptotically distributed as a χ^2_2-variable

$$2(\lambda_0 - \lambda_1) = \chi^2_2 \,. \tag{11}$$

We use this result to test whether observers' estimates $(\hat{\psi}_P, \hat{\varphi}_P)$ were significantly different from the true values. We summarize the values of the observers' estimates $(\hat{\psi}_P, \hat{\varphi}_P)$ in Table 2 along with the corresponding p values. None of the observers' punctate source direction estimates were significantly different from their true values under all conditions. Observers' azimuth and elevation estimates of the punctate source direction were within 11 deg of the true direction with one exception (the azimuth estimate for RG with the light on the LEFT; See Table 2). We replot the azimuth and elevation estimates in Table 2 as Figure 12. It is then readily seen that the estimates $\hat{\psi}_P$ are clustered around the true values (with one exception, RG, LEFT) and that the estimates $\hat{\varphi}_P$ are typically too large. We have, in effect, derived a crude estimate of the position of the light from the observer's performance.

We next examine the three outcomes discussed in the "Introduction." We first considered the hypothesis that observers' achromatic settings were not affected by changes in test patch orientation. If this were so, then we would find that the geometric b-chromaticity function was constant. We rejected this hypothesis for all observers in all conditions (p < .0001 in all cases). We conclude that the observers' achromatic settings are affected by changes in test patch orientation.

We next tested the hypothesis that observers correctly discounted the effect of orientation in making achromatic settings (i.e., were observers accurately estimating the b-chromaticities of the punctate and diffuse sources, $\hat{\pi}^B$ and $\hat{\delta}^B$, and the diffuse-punctate balance $\hat{\Delta}$ and using these estimates to discount the effect of rotations?). We tested whether observers' estimates were equal to true values $\pi^B = 0$, $\delta^B = 0.66$ and $\Delta = 0.37$. We nested the hypothesis that the parameters were equal to the true values within a model in which they were free to vary. We rejected the hypothesis that $\left(\hat{\pi}^B, \hat{\delta}^B, \hat{\Delta}\right) = (0, 0.66, 0.37)$ for all observers

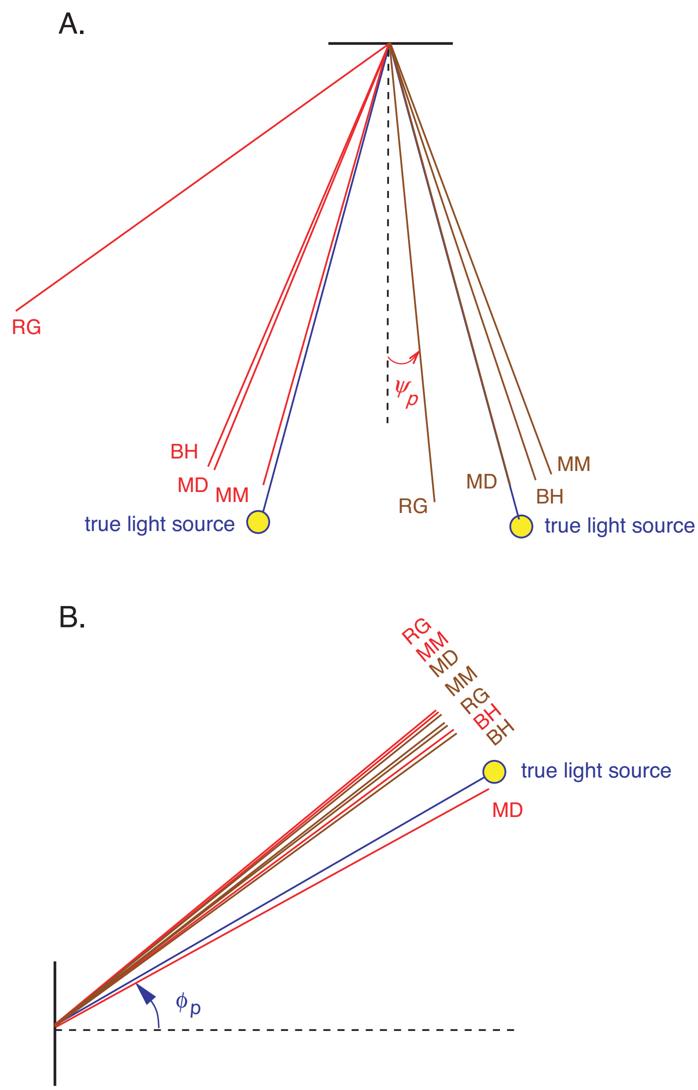

Figure 12. Observer's estimates of punctate light direction. A. ψ_P component of the estimates. B. φ_P component of the estimates. Both sets of estimates are taken from Table 2. The ψ_P component of observers' estimates is always in the correct quadrant and (with one exception) clustered around the true values, whereas the φ_P component is typically overestimated.

	Punctate source position: LEFT		Punctate source position: RIGHT	
	$(\hat{\psi}_P, \hat{\phi}_P)$	$(\hat{\pi}^B, \hat{\delta}^B, \Delta)$	$(\hat{\psi}_P, \hat{\phi}_P)$	$(\hat{\pi}^B, \hat{\delta}^B, \Delta)$
Veridical	**(-15, 30)**	**(0, 0.66,0.37)**	**(15, 30)**	**(0, 0.66,0.37)**
BH	**(-23.1, 36.6)** p = .011	**(0, 0.33, 1.39)*** p < .001	**(18.4, 36.1)** p = .007	**(0.1, 1, 0.12)** p = .036
MD	**(-22.4, 28.8)** p = .63	**(0.1, 0.38, 1.87)*** p < .001	**(15.2, 38.8)** p = .04	**(0, 0.4, 1.6)*** p < .001
MM	**(-15.8, 39.3)** p = .734	**(0.2, 1, 0.03)*** p < .001	**(20.3, 37.9)** p = .426	**(0.2, 1, 0.03)*** p < .001
RG	**(-54.4, 39.5)** p = .029	**(0.2, 0.31, 2.15)*** p < .001	**(5.6, 37.4)** p = .016	**(0.2, 1, 0.05)*** p < .001

Table 2. Achromatic setting: maximum likelihood estimates of lighting parameters. We report the maximum likelihood estimations of the punctate source direction $(\hat{\psi}_P, \hat{\varphi}_P)$ and the lighting parameters $(\hat{\pi}^B, \hat{\delta}^B, \hat{\Delta})$. The parameter π^B is the b-chromaticity of the punctate source, δ^B is b-chromaticity of the diffuse source, and Δ is the ratio of the intensity of the diffuse source to the intensity of the punctate. For each observer, we tested the hypotheses that $(\hat{\psi}_P, \hat{\varphi}_P)$ and $(\hat{\pi}^B, \hat{\delta}^B, \hat{\Delta})$ are equal to the veridical values and report exact p values for the tests when the values are larger than .001. With a Bonferroni correction for 40 tests (four observers, five parameters, two punctate source positions), the significant level corresponding to an overall Type I Error rate of .05 is .00125. Values whose corresponding p values fall below this cutoff are marked with an asterisk. All observers' punctate source direction estimates were not significantly different than the veridical values. However, in contrast with their success in estimating the punctate source direction, observers misestimated the other lighting parameters $(\hat{\pi}^B, \hat{\delta}^B, \hat{\Delta})$. The deviations from the veridical values were significant (except for observer BH when the punctate source was on the right). Those deviations in the lighting parameter estimates result in failure to discount perceived orientation of the test patch for its perceived color exactly as the model predicts.

except observer BH for the punctate-on-the-right condition (p = .036; all other p values are reported in Table 2).

The analysis of the results suggest that one possible reason for the failure of the perfect discounting is that the observers' estimates of the chromaticity balances of the punctate and diffuse sources, and of the diffuse-punctate balance, are in error. All observers slightly overestimated the b-chromaticity of the punctate source $\hat{\pi}^B > \pi^B$ (veridical value is $\pi^B = 0$), and misestimated the b-chromaticity of the diffuse source δ^B and the diffuse-punctate balance Δ. The values are reported in Table 2. It is as if observers are discounting the orientation of the test patch for the achromatic task consistent with the physically correct model but using incorrect estimates of the lighting parameters.

Red-green balance

We also examined the red-green balances of the achromatic settings. We define $\hat{\Lambda}^{R/G} = E^R_{set} / E^G_{set}$ as the red-green balance. We present only one observer's (MM) settings in Figure 13. As expected, the red-green balance of achromatic settings did not change systematically with changes in the orientation of the test patch. All other observers' results were similar.

A neutral light source control

As the test patch rotates away from the yellow punctate light source not only the relative blue content of the patch increases but also the overall intensity of the light from the test patch decreases (as described for Lambertian surfaces). What if these two events are confounded, that is, what if subjects simply assume that darker objects appear more blue? We verified that this is not the case by letting one observer (MD) run an extra session in the experiment where the punctate and diffuse light sources were neutral ($E^R_P = E^G_P = E^B_P$ and $E^R_D = E^G_D = E^B_D$; $\delta^B = \pi^B$). We found no effect of orientation on his achromatic settings, neither

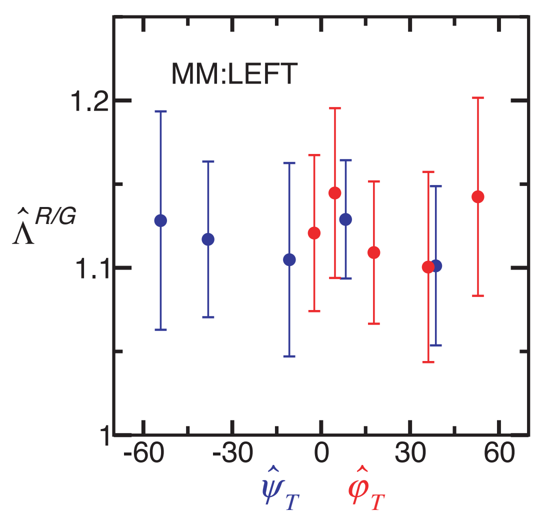

Figure 13. Red-green balance. We checked observers' achromatic settings for the ratio of the red content to green content. Because the red-green balance of the punctate and diffuse source is constant, independent of the orientation of the test surface, we did not anticipate any variation in the red-green balance of observers' settings. This is what we found. We plot one observer's (MM) results here. All other observers' results were similar. Error bars are $\pm 2\ SE$ of the mean (approximately a 95% confidence interval).

for blue-total balance (Figure 14) nor for red-green balance. This result is consistent with Equation 7, and indicates, for example, that observers do not simply increase the blue content of a gray surface as its total intensity decreases.

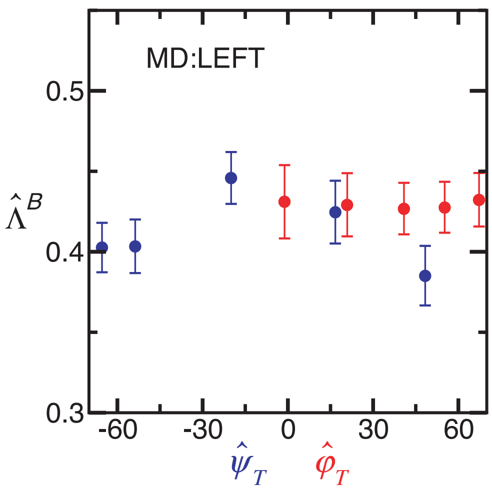

Figure 14. Achromatic setting under neutral light. As a control for the results, we repeated the experiment with neutral light sources ($E_P^R = E_P^G = E_P^B$ and ; $\delta^B = \pi^B$). Only one observer (MD) completed the control experiment. His results show no patterned change with changing orientation of the test patch. This lack of pattern indicates that he does not simply add more blue content as the test patch rotates and gets darker. Error bars are ± 2 *SE* of the mean (approximately a 95% confidence interval).

Geometric discounting index

We can quantify the observer's performance in discounting the perceived orientation for perceived color by, first of all, noting how close his or her estimates of light source direction are to the true direction. This measurement, though, does not reflect errors in the observer's estimates of the remaining lighting parameters. We define a geometric discounting index that effectively compares the curvature of the observer's b-geometric chromaticity function at its minima,

$$DI = 1 - \frac{|\kappa - \hat{\kappa}|}{\kappa + \hat{\kappa}}, \tag{12}$$

where

$$\kappa = \frac{\kappa_\psi + \kappa_\varphi}{2}, \quad \hat{\kappa} = \frac{\hat{\kappa}_\psi + \hat{\kappa}_\varphi}{2},$$

and

$$\kappa_\psi = \left. \frac{\partial^2 \Lambda^B}{\partial \psi_T^2} \right|_{\psi_T = \psi_T^*}$$
$$\kappa_\varphi = \left. \frac{\partial^2 \Lambda^B}{\partial \varphi_T^2} \right|_{\varphi_T = \varphi_T^*}, \tag{13}$$

and $\hat{\kappa}_\psi$ and $\hat{\kappa}_\varphi$ are corresponding curvatures of the observers' b-geometric chromaticity function calculated in the same way. *DI* is a measure of how much the observer's geometric b-chromaticity function is curved compared to the theoretical one. A complete lack of curvature ($\hat{\kappa} = 0$) corresponds to the case where the observer's achromatic settings are unaffected by angle. Then $DI = 0$, and we would conclude that the observer's color estimates are not affected by perceived surface orientation. If, in contrast, $\kappa = \hat{\kappa}$, then $DI = 1$. Note that *DI* is a composite measure of all the parameters in the parameter space $\hat{\Theta} = \{\hat{\psi}_P, \hat{\varphi}_P, \hat{\pi}^B, \hat{\delta}^B, \hat{\Delta}\}$. However, because the estimated parameters $\hat{\psi}_P, \hat{\varphi}_P$ were close to veridical, the errors are effectively due only to the misestimated chromaticity balances of the punctate source and diffuse source and the overall diffuse-punctate balance. The values of *DI* varied between 0.29 and 0.82 (*DI* = 0: no discount; *DI* = 1: perfect discount). The discounting indices are reported in Table 3 and in the legends of Figure 11.

	Discounting index, *DI*	
	Punctate source position: LEFT	Punctate source position: RIGHT
BH	0.78	0.8
MD	0.69	0.82
MM	0.29	0.27
RG	0.2	0.43

Table 3. Discounting Indices. We define a discounting index *DI* to quantify how well observers discounted the effective illumination in the experimental scenes. *DI* is a comparison of how the observers' achromatic settings vary with perceived test patch orientation, and how they should vary if the observer is correctly discounting the effective illumination on the test patch. The exact form of *DI* is given in the text. A zero value of *DI* means no discounting (the achromatic setting is unaffected by perceived orientation). A value of 1 corresponds to perfect discounting. All observers partially compensated for changes in effective illumination due to changes in test patch orientation.

Conclusion

When a scene is illuminated by punctate and diffuse light sources differing in chromaticity, the chromaticity of the light absorbed by any patch or surface depends on its orientation. We report an experiment in which observers were asked to view rendered 3D scenes binocularly. The lighting in these scenes was composite, a mixture of a yellow punctate light source and a blue diffuse. On each trial, observers adjusted a test patch to be achromatic (achromatic setting task) and then adjusted a gradient probe to match the orientation of the patch (orientation task). We varied the orientation of the test patch randomly across trials. The location of the punctate source could be either LEFT or RIGHT, and was fixed for a given session.

We used the results of the gradient probe task to estimate the observer's perceived orientation. We showed that observers systematically take the perceived orientation of the test patch into account in making achromatic settings. However, their settings do not match the settings of an ideal Lambertian observer who has knowledge of the parameters of the composite lighting model (the location of the punctate light source and the chromaticities and intensities of the diffuse and punctate light sources).

We refit the observers' data on the assumption that the observer correctly discounted the illumination arriving at the test patch in making achromatic settings, but that, in doing so, he or she made use of estimates of parameters of the composite lighting model that were in error. We found very good agreement between this equivalent illuminant model of the observer's performance and our data. All of the observers' estimates of the direction to the punctate source were close to veridical. The deviations from the model were for the most part consistent with the hypothesis that observers failed to use correct estimates of the chromaticities of the punctate source and of the diffuse source, and the ratio of intensity of the diffuse source to the intensity of the punctate source.

It has been shown that the perceived geometry of a scene influences the lightness (perceived albedo) of surfaces (Gilchrist, 1977, 1980; Gilchrist et al., 1999; Boyaci et al., 2003). There are also studies that show that the perceived color of surfaces is influenced by the spatial arrangements permitting mutual illumination (Bloj et al., 1999; Doerschner et al., 2004). Because we used several different orientations of the Lambertian test patch, we were able to parametrically determine how observers' surface color estimates are influenced by surface orientation in scenes with composite light models. Observers do take into account the 3D structure of scenes and the lighting model of the scene in arriving at estimates of surface color. Their shortcomings are predominantly consistent with failures to estimate the parameters of the lighting model correctly.

Our results suggest that the observer, in effect, develops a model or estimate of the spatial distribution and chromaticities of light sources in a scene as part of color visual processing. Maloney and Yang (2003; Maloney, 1999, 2002) review previous work related to the hypothesis that the visual system develops an estimate of illuminant chromaticity, and it is natural to extend this "illuminant estimation hypothesis" to include explicit estimation of the spatial distribution of light sources as well.

In retrospect, it is perhaps not surprising that the human observer, living in a world that often is illuminated by sky and sun differing in chromaticity, would be able to compensate for the effect of orientation in arriving at estimates of perceived color. However, this study is, to our knowledge, the first to show that the human visual system does so, even partially. These results for color together with previous results for lightness (Boyaci et al., 2003) and for mutual illumination (Bloj et al., 1999; Doerschner et al., 2004) support the claim that the visual system effectively estimates the spatial and chromatic properties of the illuminant.

How it does so raises a new and exciting set of questions about surface color perception (Maloney, 1999, 2002). How, for example, does the visual system arrive at the admittedly imperfect estimates of punctate source direction that we derive from the observers' own performance? Recall that the punctate light source in our scenes is not even directly visible to the observer. There are several candidate cues that might provide this information as well as information about the chromaticities and relative intensities of the light sources: specular highlights, cast shadows, and attached shadows (shading). It will be of interest to determine which of these cues are actually used by the visual system. Our results also indicate that the observers' estimates of light source chromaticity and spatial distribution of light sources can be markedly in error.

We do not claim that observers are aware of their lighting models or of the cues that signal it (cf. Rutherford & Brainard, 2002). Kafka (1911/1988, p. 63) described the lighting in his room as, "The lights and shadows thrown on the walls and the ceiling by the electric lights in the street and the bridge down below are distorted, partly spoiled, overlapping, and hard to follow. When they installed the electric arc-lamps down below and when they furnished this room, there was simply no housewifely consideration given to how my room would look from the sofa at this hour without any lights of its own." Unlike his room, our scenes were designed carefully and with "housewifely consideration." Yet, looking at the scenes of this experiment, the observers did not have a better understanding of the purposes of the "creator": When asked after the experiment, observers were not even aware that the punctate source changed its position from session to session, or that there was a blue background. Their visual system simply took care of the "overlapping, and hard to follow details" for them.

Acknowledgments

This research was funded in part by National Institute of Health Grant EY08266. HB and LTM were also supported by grant RG0109/1999-B from the Human Frontiers Science Program. We thank Michael Landy for comments on earlier drafts and David Brainard for comments on this work in poster form.

Commercial relationships: none.
Corresponding author: Huseyin Boyaci.
Email: boyaci@cns.nyu.edu.
Address: New York University, Department of Psychology & Center for Neural Science, 6 Washington Place, 877 New York, NY 10003.

Footnotes

[1]We use the term RGB code as a convenient synonym for the tristimulus coordinates based on the three linearized sources ('guns') of the monitor (Wyszecki & Stiles, 1982).

[2]The RGB code is the tristimulus coordinates with respect to the three linearized sources of the monitor and we define the r-, g- and b-chromaticities e.g. b = B/(R+G+B) in the usual manner (Wyszecki & Stiles, 1982).

References

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

Boyaci, H., Maloney, L. T. & Hersh, S. (2003). The effect of perceived surface orientation on perceived surface albedo in binocular-viewed scenes. *Journal of Vision*, *3*(8), 541-553, http://journalofvision.org/3/8/2, doi:10.1167/3.8.2 [PubMed][Article]

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. [PubMed]

Doerschner, K., Boyaci, H., & Maloney, L. T. (2004). Human observers compensate for secondary illumination originating in nearby chromatic surfaces. *Journal of Vision*, *4*(2), 92-105, http://journalofvision.org/4/2/3/, doi:10.1167/4.2.3. [PubMed][Article]

Gilchrist, A. L. (1977). Perceived lightness depends on spatial arrangement. *Science*, *195*, 185-187. [PubMed]

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

Gilchrist, A., Kossyfidis, C., Bonato F., Agostini T., Cataliotti J., Li X. J., et al. (1999). An anchoring theory of lightness perception. *Psychological Review*, *106*, 795-834. [PubMed]

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

Kafka, F. (1988). *The diaries, 1910-1923* (M. Brod, Ed., & J. Kresh, Trans.) New York: Schocken Books. (Original work published in 1911)

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

Larson, G. W., & Shakespeare, R. (1996). *Rendering with radiance; The art and science of lighting and visualization*. San Francisco: Morgan Kaufmann Publishers, Inc.

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.

Maloney, L. T. (2002). Illuminant estimation as cue combination. *Journal of Vision*, *2*(6), 493-504, http://journalofvision.org/2/6/6/, doi:10.1167/2.6.6. [PubMed][Article]

Maloney, L. T., & Yang, J. N. (2003). The illumination estimation hypothesis and surface color perception. R. Mausfeld & D. Heyer (Eds.), *Colour: Connecting the mind to the physical* world (pp. 335-358). Oxford: Oxford University Press.

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

Rutherford, M. D., & Brainard, D. H. (2002). Lightness constancy: A direct test of the illumination estimation hypothesis. *Psychological Science*, *13*, 142-149. [PubMed]

Scheifler, R. W., & Gettys, J. (1996). *X Window System; Core library and standards*. Boston: Digital Press.

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

Yang, J. N., & Shevell, S. K. (2002). Stereo disparity improves color constancy. *Vision Research*, *42*, 1979-1989. [PubMed]

Yang, J. N., & Shevell, S. K. (2003). Surface color perception under two illuminants: The second illuminant reduces color constancy. *Journal of Vision*, *3*(5), 369-379, http://journalofvision.org/3/5/4/, doi:10.1167/3.5.4. [PubMed][Article]

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