Human observers compensate for secondary illumination originating in nearby chromatic surfaces

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

Huseyin Boyaci Department of Psychology and Center for Neural Science, 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

In complex scenes, the light absorbed and re-emitted by one surface can serve as a source of illumination for a second. We examine whether observers systematically discount this secondary illumination when estimating surface color. We asked six naïve observers to make achromatic settings of a small test patch adjacent to a brightly colored orange cube in rendered scenes. The orientation of the test patch with respect to the cube was varied from trial to trial, altering the amount of secondary illumination reaching the test patch. Observers systematically took orientation into account in making their settings, discounting the added secondary illumination more at orientations where it was more intense. Overall, they tended to under-compensate for the added secondary illumination.

Keywords: vision, surface color perception, inter-reflection, mutual illumination

Introduction

In complex scenes, the light emitted by one surface can fall on a second, becoming, in effect, a component of the illumination incident on the second. In Figure 1A, for example, the light gray matte *test surface* marked T absorbs light that reaches it directly from the single light source in the scene. It also absorbs light that arrives from the same light source but only after being absorbed and re-emitted from the nearby orange surface marked C (Figure 1B). Part of the light absorbed and re-emitted by the test patch will in turn be absorbed and re-radiated by the orange surface, initiating an infinite series of inter-reflections between the surfaces. If we denote the spectral power distribution of the original illuminant by $E_{(0)}(\lambda)$ and the surface reflectance functions of the two surfaces by $S_C(\lambda)$ (cube) and $S_T(\lambda)$ (test patch), then the light emitted from any specified small region of the surface toward the observer can be written in the form $E(\lambda)S_T(\lambda)$ where

$$E(\lambda) = \sum_{i=0}^{\infty} \gamma_i E_{(i)}(\lambda) \quad (1)$$

is the *effective illuminant*. It is the weighted sum of the direct illumination, $E_{(0)}(\lambda)$, and the inter-reflected illuminants,

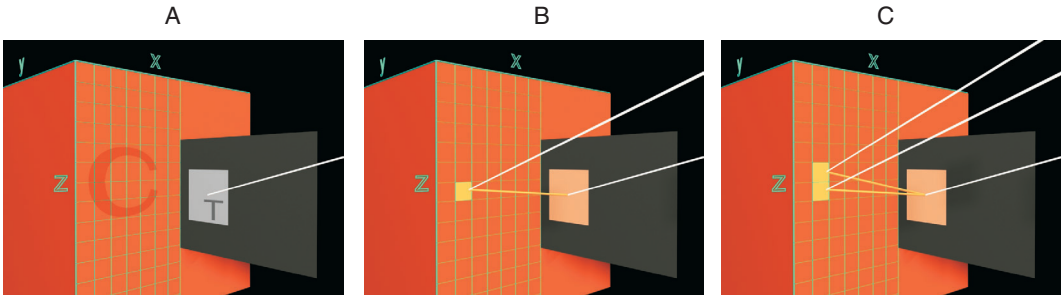

Figure 1. Modeling inter-reflection between Lambertian surfaces. A simple scene consisting of an orange cube with a small light gray test patch embedded in a larger dark gray patch. A. Zero-bounce light. Light from a light source in the scene reaches the test patch directly. B. One-bounce light (single ray). Light from the same light source reaches the test patch after being absorbed and remitted by a small area element on the surface of the cube. C. One-bounce light (multiple rays). The total one-bounce light is the sum of contributions from all such area elements on the face of the cube. The letters C and T on cube and test patch, respectively (Figure 1A), are omitted in Figure 1B and 1C.

doi:10.1167/4.2.3 Received August 29, 2003; published February 27, 2004 ISSN 1534-7362

$$E_{(1)}(\lambda) = E_{(0)}(\lambda)S_C(\lambda);$$
$$E_{(i+2)}(\lambda) = E_{(i)}(\lambda)S_T(\lambda)S_C(\lambda) \quad i = 0,1,\cdots \tag{2}$$

The *geometric factors* γ_i are determined by the sizes and shapes of the two surfaces, their separation, and their orientations with respect to one another and with respect to the primary light source. We will assume that they do not depend on wavelength λ in the electromagnetic spectrum.

To make stable estimates of the surface color and albedo of a surface patch in a scene, independent of scene layout or illumination, a visual system must discount the effective illuminant $E(\lambda)$ at each point in the scene. There is some evidence that human observers do so, if only partly and imperfectly.

Although there is a large body of literature on discounting of the illuminant and color appearance models [for a review on color constancy, see Hurlbert (1998) and Maloney (1999)], there is relatively little research about color perception in complex scenes, or scenes viewed binocularly, and very little concerning discounting of inter-reflection.

Bloj, Kersten, and Hurlbert (1999) demonstrated that perceived surface color changes when perceived spatial organization permits or precludes inter-reflection. Their stimulus was a chromatic version of the Mach card; one side of the card was painted magenta and the other white. The card was folded in such a way that the angle between the magenta and white side measured 70° (concave). The secondary illumination that arose from the magenta caused a pinkish gradient on the white side of the card. The card was viewed directly and also through a pseudoscope that reversed the disparities in left and right eye. As a consequence of this reversal, the card appeared to be convex (290°). If the card were actually convex, then light emitted by the magenta side could not travel directly to the white side. Observers judged the color of the white side to be more "pinkish" in the apparently convex condition than in the actual concave condition, indicating that they incorporated information about the shape of the object into their estimates of surface color and suggesting that they were discounting the effect of inter-reflection when it was perceived to be possible.

Bloj et al. considered only one spatial configuration, and based on their result, we cannot conclude that human observers are capable of discounting the effect of inter-reflection systematically. By manipulating the spatial configuration (see below) under which inter-reflection might occur, we can derive a parametric model that allows us to assess how human observers discount this effect.

In this work, we investigate a wider range of spatial configurations, varying the angle between two flat Lambertian (matte) surfaces. One is large and bright orange (Figure 1). The other (the *test patch*) is small, and its chromaticity is under the control of the observer. The observer is asked to set the second surface to be neutral in appearance. As we explain below, the results of this setting task will permit us to assess how accurately human observers discount inter-reflection. First, we describe how light travels through the scenes that we will use as stimuli.

Modeling secondary illumination

Consider the simple scene in Figure 1A. A light gray small square (the test patch *T*) that is located on a larger dark gray rectangle is rotated toward the side of a brightly colored orange cube at a certain angle τ. All surfaces are Lambertian (we will define precisely what this means in a moment). The scene is illuminated by a neutral punctate light source $E_{(0)}(\lambda)$ placed behind the observer. This light source is sufficiently far away from the small test patch to allow us to assume that the distance to the light source and the angle of incidence θ_{TP} of light from the primary light source on the test patch is constant across the extent of the patch. See Figure 2 for a definition of the angles that are relevant to our discussion.

The angle of incidence θ_{TP} determines the flux of light from the punctate source that falls on the test surface. For Lambertian surfaces, luminance decreases as the angle between surface normal and the direction to the light source increases. If, for example, the test patch is rotated toward the light source so that the angle between its surface normal and the incident light ray is 0° it will receive the maximum amount of light, and its luminance will be at its maximum. Away from this position less light will be received by the test patch, and its luminance will decrease.

The test patch is assumed to be a Lambertian surface and, consequently, the spectral power distribution of the "zero bounce" light that is emitted from the surface is

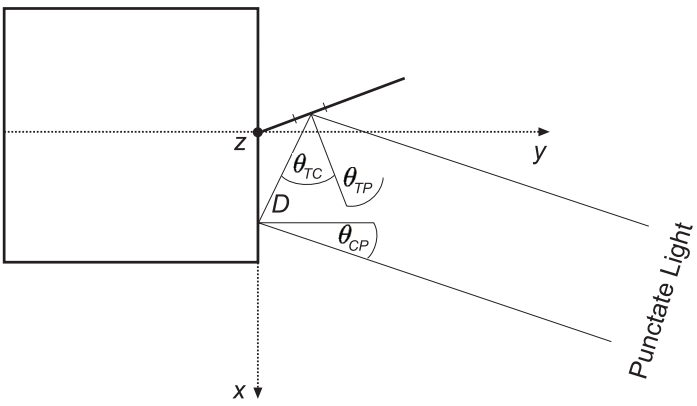

Figure 2. Relevant angles. θ_{TP} is the angle between the punctate light source and the test patch. θ_{CP} is the angle between the punctate light source and the surface of the cube. The punctate light source is effectively collimated and θ_{CP} is constant across the surface of the cube. θ_{TP} does vary as we vary the angle τ (not shown) between the test patch and the cube surface. D is the distance between the position of the area element on the cube surface and the center of the test patch. It varies with both position and with τ.

$\cos\theta_{TP}\, S_T(\lambda) E_{(0)}(\lambda)$, where $S_T(\lambda)$ is the surface reflectance of the test patch. The expression $\cos\theta_{TP} E_{(0)}(\lambda)$ is the first term ("zero bounce" term) in the summation in Equation 1. Comparing terms, we see that the first geometric factor is

$$\gamma_0 = \cos\theta_{TP}. \tag{3}$$

When the test patch is neutral (achromatic), we can replace the surface reflectance function $S_T(\lambda)$ of the neutral test patch by its albedo,

$$S_T(\lambda) = \alpha_T. \tag{4}$$

That is, we define a neutral surface as one that absorbs and re-radiates light without altering its spectral power distribution.

Secondary light

Next we consider the "one bounce" term $E_{(1)}(\lambda)$, the light reflected from the adjacent surface of the cube (the *cube surface* C). The spectral power distribution of the "zero bounce" light that is emitted from the surface of the cube is $\cos\theta_{CP}\, S_C(\lambda) E_{(0)}(\lambda)$, where $S_C(\lambda)$ is the surface reflectance of the cube. If the light arriving at the cube's surface ($E_{(0)}(\lambda)$) is neutral, then the light reflected from the cube would take on the chromaticity of $S_C(\lambda)$. $E_{(1)}(\lambda)$ is the sum of contributions from each area element on C (as shown in Figure 1C). We set up a Cartesian coordinate system (x,z) for the face of the cube (Figure 1) and integrate the contribution from each such element to obtain the total illumination upon the test patch from the surface of the cube. This constitutes the second term in Equation 1

$$\gamma_1 E_{(1)}(\lambda) = \frac{1}{\pi} E_{(0)}(\lambda) S_C(\lambda) \times \iint_C \cos\theta_{CP}\, \cos\theta_{TC}(x,z)\, D^{-2}(x,z)\, dx\, dz\ , \tag{5}$$

where θ_{CP} is the angle of incidence of the light arriving from the collimated punctate light source on the surface of the area element at (x,z). We assume that it is independent of the location (x,z) on the cube, because the punctate light source is far away in our scenes; $\theta_{TC}(x,z)$ is the angle of incidence of the light arriving from the area element at (x,z) on the test patch; and $D(x,z)$ is the distance between the area element at (x,z) on the cube and the center of the test patch.[1] Because $E_{(0)}(\lambda) S_C(\lambda) = E_{(1)}(\lambda)$, the second geometric factor is seen to be

$$\gamma_1 = \frac{1}{\pi}\cos\theta_{CP} \times \left[\iint_C \cos\theta_{TC}(x,z)\, D^{-2}(x,z)\, dx\, dz \right], \tag{6}$$

where $\cos\theta_{CP}$ is moved out of the integral since it is constant across the side of the cube.

This coefficient captures the effect of the spatial layout of the cube and the test patch on the intensity of the one-bounce illuminant. We can compute the coefficients γ_2 and beyond similarly. Expressions for these terms grow rapidly in complexity and, in many scenes, the first two terms of Equation 1 dominate. This, however, need not always be the case. Under a forest canopy, when neither sun nor sky is directly visible, the higher order terms that result from multiple reflections among leaves likely dominate (Endler, 1993).

When the first two terms dominate, the light reaching the small test patch is predominantly a weighted mixture of zero- and one-bounce light,

$$E(\lambda) \approx \gamma_0 E_{(0)}(\lambda) + \gamma_1 E_{(1)}(\lambda), \tag{7}$$

and the weights controlling the mixture vary systematically with the orientations of both the test surface and the cube and the direction to the punctate light source. As the test patch rotates away from the cube, the geometric factor γ_1 of the one-bounce light $E_{(1)}(\lambda)$ decreases (because the $\cos\theta_{TC}$ term decreases, while the distance *D* increases), and the geometric factor γ_0 of the zero-bounce light $E_{(0)}(\lambda)$ also changes (because the angle θ_{TP} changes as the test patch rotates). If we let τ denote the angle between the surface patch and the face of the cube, then, as we change this angle, we can write the effective illumination of the test patch as

$$E(\lambda) \approx \gamma_0(\tau) E_{(0)}(\lambda) + \gamma_1(\tau) E_{(1)}(\lambda), \tag{8}$$

where we have made it explicit that the coefficients γ_0 and γ_1 both depend on τ. In Figure 3, we plot $\gamma_0(\tau)$ and $\gamma_1(\tau)$ versus τ for the scene shown in Figure 1 (See "Appendix" for derivations). The function $\gamma_0(\tau)$ reaches a maximum when the test patch *T* is facing the punctate light source. As τ increases, the contribution of secondary illumination $\gamma_1(\tau)$ decreases.

Boyaci, Maloney, and Hersh (2003) showed that human observers incorporate knowledge of scene geometry, in particular the consequences of the incident angle of the primary light source, into their judgments about the albedo of a matte surface. Boyaci, Doerschner, and Maloney (in press) showed that human observers also take scene geometry into account in scenes with two illuminants differing in chromaticity. In terms of Equation 8, these studies indicate that the visual system can in effect compensate for the first geometric factor $\gamma_0(\tau)$ in estimating surface color and albedo. When a secondary ("one-bounce") illuminant is present, however, the visual system must somehow compensate for changes in both geometric factors $\gamma_0(\tau)$ and $\gamma_1(\tau)$ with changes in τ, or, more generally, scene layout. The second geometric factor depends on many factors and, as it is written in Equation 6, involves a double integration.[2] In this study we investigate whether human vision can compensate for secondary illumination that results from inter-reflection between two surfaces as we vary the angle τ between the surfaces.

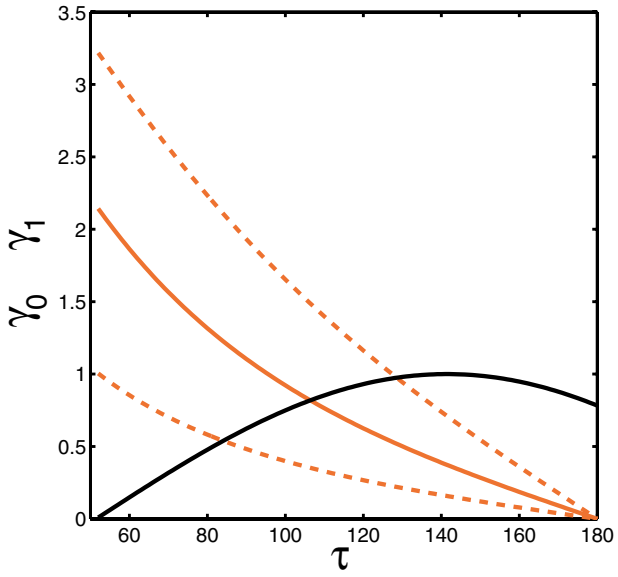

Figure 3. Geometric factors. The two geometric factors $\gamma_0(\tau)$ (solid black) and $\gamma_1(\tau)$ (solid orange) are plotted versus τ for the scene in Figure 1. The first geometric factor $\gamma_0(\tau)$ achieves a maximum when the angle between the surface normal to the test patch and the direction to the punctate light source is smallest. The second geometric factor $\gamma_1(\tau)$ depends upon the size and shape of the orange surface C. The dotted orange lines are plots of $\gamma_1(\tau)$ versus τ with the area of the cube scaled by .55 (lower line) and by 1.83 (upper line). The solid orange curve is a plot of $\gamma_1(\tau)$ with area of the cube scaled by 1 (this was the value used for the experimental stimuli described below).

Experiment

Introduction

In this experiment we tested whether human observers can correctly discount secondary ("one bounce") illumination. We used an achromatic setting task (Helson & Michels, 1948).

Methods

Stimuli

The stimuli were computer-rendered, three-dimensional complex scenes composed of simple objects with various shapes (such as spheres and boxes), and various reflectance properties (such as shiny, matte, and transparent). All scenes were rendered with the Radiance software package (Larson & Shakespeare, 1996). In rendering each scene, we used a four-bounce approximation (rays were permitted to strike up to five surfaces). To verify that the effects of second and higher bounces were negligible for our stimuli, we also rendered each scene with a one-bounce approximation and compared the effective illuminant on the neutral test patch of one- and four-bounce renderings. (Comparison of the RGB pixel values at the center of the scene with one- and four-bounce renderings yielded a mean square error [MSE] of 0.0021, in normalized RGB values). Each scene was rendered twice from slightly different viewpoints corresponding to the positions of the observer's eyes. A stereo pair for a typical scene is shown in Figure 4.

All scenes contained a large orange cube near the center of the scene whose surface properties were never varied and whose location remained unchanged. A rectangular plane containing a smaller square *test patch* was attached to the side of the cube.

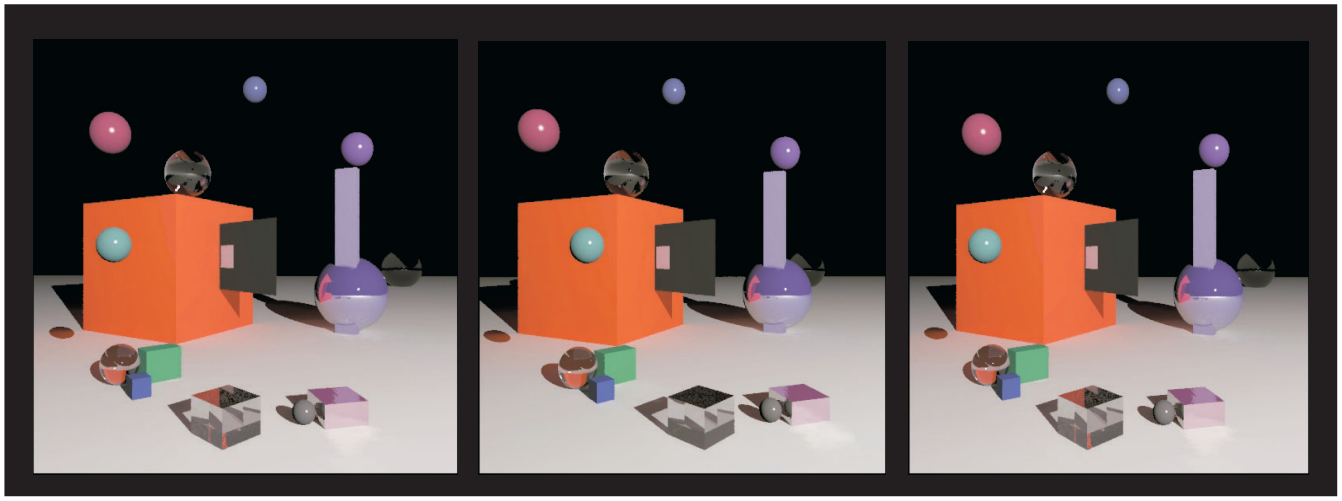

Figure 4. A stereo pair for a typical scene. The left pair can be used for crossed-fusion, the right pair for uncrossed.

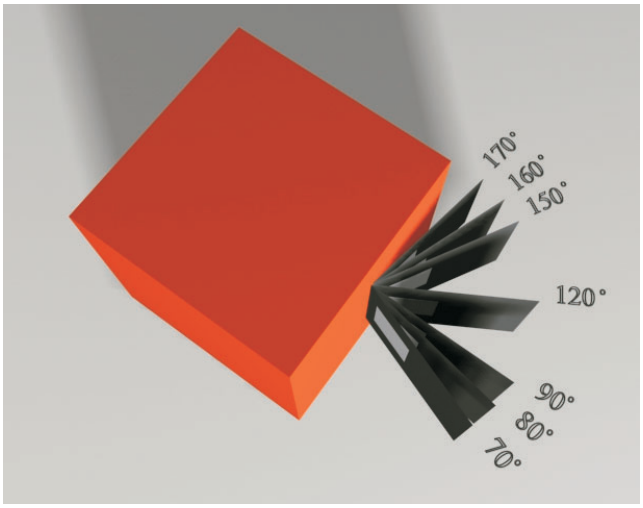

Figure 5. Orientations of the test patch. The orientation of the test patch relative to the side of cube varied randomly from trial to trial among the seven orientations shown.

We varied the orientation of the plane (and therefore the test patch) from trial to trial (see Figure 5). The simple, additional objects in the scenes were varied randomly from trial to trial. These objects could be shiny, matte, or partly shiny and matte. They were intended to provide information about the location of the punctate light source. Cube, plane, and test patch had Lambertian surface reflectance properties.

Spatial coordinate system and spatial arrangement

We used a Cartesian coordinate system with its origin at the center of the side of the cube where the test patch was also attached. The z axis was vertical and aligned with the side of the cube, the x axis was horizontal and aligned with the side of the cube, and the y axis was normal to the same side of the cube (see Figure 6). The angle between the cube's face and the test patch is denoted by τ.

Color coordinate system

In computer graphics rendering, the chromatic properties of surfaces and lights are typically specified by three numbers referred to as RGB codes. The interaction of light and surface is modeled by component-wise multiplication of the RGB codes assigned to the light and to the surface. This model of surface-light interaction is not accurate (see discussion of the "RGB heuristic" in Maloney, 1999; see also Yang & Maloney, 2001). Accurate rendering of arbitrary surfaces and lights is only possible if color codes contain more than three numbers. However, for the scenes considered in this study, we can, in fact, produce physically-accurate renderings using only RGB codes. Moreover, we can identify these codes with the red, green, and blue guns of the CRT monitors we use.

This simplification is possible because of the restricted range of lights and surfaces we employ. The primary light source in the scene is defined to be a "neutral" light and given the RGB code [1,1,1]. The primary light source is behind the observer but, were we to render and display it, it would be assigned settings proportional to [1,1,1]. When this light interacts with a Lambertian surface that has a specified RGB code, the chromaticity of the light emitted from the surface is proportional to that RGB. That is, a neutral light absorbed and re-emitted by a surface has the RGB chromaticity of the surface. We take this as the definition of neutral light, at least over the range of surfaces in our experiments. The computations of the rendering package Radiance will correctly compute such interactions of neutral light and chromatic surface. Similarly, we assign RGB codes proportional to [1,1,1] to describe the surface reflectance of neutral surfaces and, thereby, define that a neutral surface is one that does not alter the chromaticity of the light that it absorbs and re-emits. With this convention, we can denote the R, G, or B component of any light by the corresponding superscript (e.g., L^R). We will also specify the chromaticities of surfaces by the albedo terms,

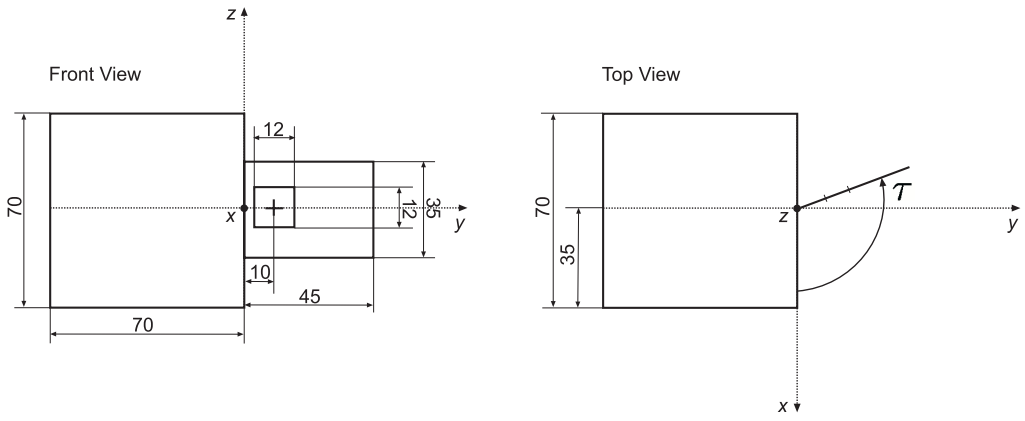

Figure 6. Coordinate system and spatial arrangement. The x axis and z axis are aligned with the side of the cube as shown; the y axis is perpendicular to the side of the cube (This same coordinate system can be seen in perspective in Figure 1). The origin of the coordinate system lies at the edge where the test patch touches the cube, at the xz side of the cube (center of the cube"s face). This drawing shows the dimensions of cube and test patch. Units are in millimeters. Note that the drawings are not completely to scale.

$\alpha^R, \alpha^G, \alpha^B$. For a neutral surface, for example, $\alpha^R = \alpha^G = \alpha^B$.

Test patch, its immediate surround and central cube

For dimensions of the plane and test patch, please see Figure 6. The size of the test patch was small enough to prevent luminance or color gradients on its surface when rendered with mutual illumination. The test patch could appear at one of seven orientations τ = {70°, 80°, 90°, 120°, 150°, 160°, 170°} after a rotation about the vertical *z axis*. The plane was fronto-parallel to the observer when τ = 120°. The dark rectangular plane on which the test patch was embedded was rendered with reflectance $\alpha^R = \alpha^G = \alpha^B = 0.01$, and the light test patch[3] with $\alpha^R = \alpha^G = \alpha^B = 0.55$. The choice of a much darker immediate surround to the test patch was in order to eliminate utilization of a simultaneous color contrast strategy by the observers (e.g., Werner & Walraven, 1982). The orange cube was rendered with reflectance $\alpha^R = 1, \alpha^G = 0.05, \alpha^B = 0$.

We rendered the entire scene, including the test patch, with the four-bounce model. However, in the beginning of a trial, the test patch was not presented to the observer with its color rendered by Radiance. Instead, we initially randomly changed the test patch's chromaticity away from that "correct" neutral point for any given trial.

Light sources

The scene was illuminated by a neutral punctate light source. This light was placed at (*x y z*) = (93.74 *cm*, 117.63 *cm*, 40 *cm*), behind and above the observer, to the right. It was sufficiently far from the cube and the test patch so that we could treat the punctate light source as collimated across the extent of the surface of the cube and the test patch.

The direction of the punctate light source can be specified by the pair of angles (ψ_P, φ_P), where ψ_P is the angle between the *x axis* and the projection of the position of the light on the *xy plane*, φ_P is the angle between the position of the light source and the *xy plane* (See Figure 7). The surface of the central cube, which is oriented toward the test patch, constitutes a secondary light source in our scene.

Apparatus

The experimental apparatus was a Wheatstone stereoscope. The left and right images were presented to the corresponding eye of the observer on two 21" Sony Trinitron Multiscan GDM-F500 monitors placed to the observer's left and right. The screens on these monitors are close to physically flat, with less than 1 mm of deviation across the surface of each monitor. Two small mirrors were placed directly in front of the observer's eyes. These mirrors reflected the images displayed on the left and right monitors upon the corresponding eye of the observer.

We tested and verified that the output of each of the monitors' guns (R, G, and B) was not appreciably affected (left less than 7%; right less than 1%) by the settings of the other two guns. These tests of additivity are available from the authors. 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 on each monitor with a Pritchard PR-650 spectrometer. The maximum luminance achievable on either screen was 114 cd/m^2. The stereoscope was contained in a box 124 *cm* on a side. The front face of the box was open and that is where the observer sat in a chin/head rest. 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.

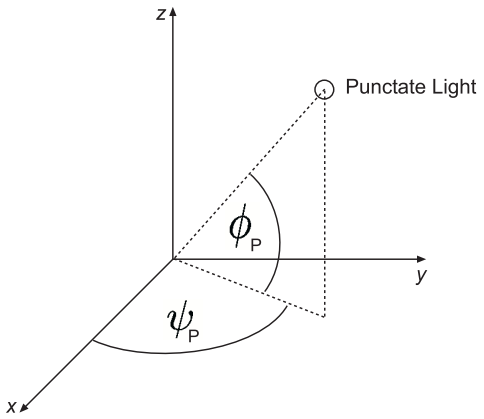

Figure 7. Punctate light source position and angles. The direction to the light source is specified by two angles (azimuth ψ_P and elevation ϕ_P) within the Cartesian coordinate system.

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* (Figure 8). To minimize any conflict between binocular disparity and accommodation depth cues, the test patches were rendered to be exactly 70 cm in front of the observer. The monocular fields of view were 55 deg × 55 deg of visual angle each. The observer's eyes were approximately at the same height as the center of the scene being viewed which was also the height of the center of the test patch.

Task

The observer was asked to adjust the color of the test patch until he or she perceived it to be achromatic. The observer was instructed to use the arrow keys of the keyboard to adjust the color in either the "blue-yellow" (up-down) or "green-red" (left-right) direction. Once the observer was satisfied with a setting, she or he hit the space bar to start the next trial (Figure 9).

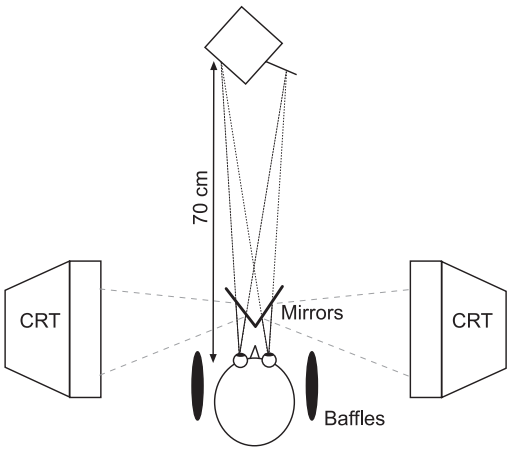

Figure 8. Wheatstone stereoscope. The left and right images of each stereo pair were displayed on two monitors placed to the left and the right of the observer. The observer viewed these images reflected in small mirrors directly in front of his or her eyes. The fused image appeared approximately 70 cm in front of the observer, the optical distance to either screen.

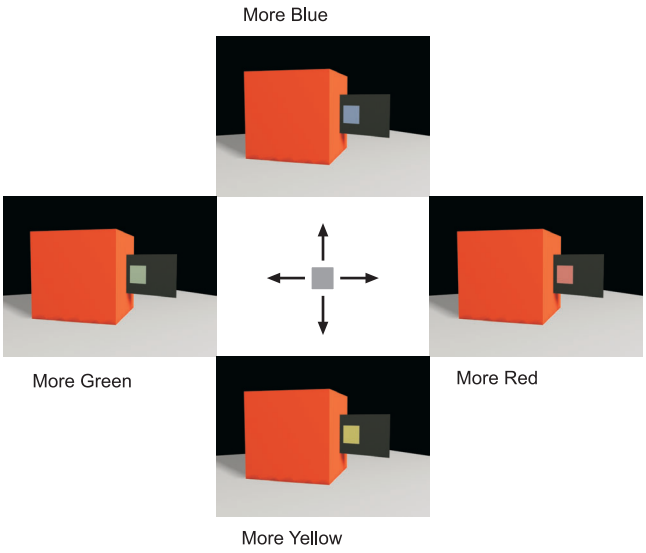

Figure 9. The task. The observer was asked to adjust the chromaticity of the test patch by pressing arrow keys until it appeared to be neither red nor green and neither blue nor yellow.

The terms "red," "blue," "green," and "yellow" refer to specific combinations of the RGB primaries, not to perceived color. Recall that a primary has the spectral power distribution of one of the guns of the monitors, linearized by a lookup table and normalized so that RGB codes [a,a,a] appear roughly neutral. Adjustments altered the intensity of the three primaries of the test patch, which we denote $\left[L^R, L^G, L^B\right]$. We constrained these three primaries so that $L^R + L^G + L^B$ was always constant. If, for example, the observer pressed the "left" key of the "left"-"right" ("green"-"red") key pair, then L^G was increased by a fixed amount δ and L^R decreased by δ so that $L^R + L^G + L^B$ remained constant. Hence, the "red"-"green" direction was simply a tradeoff between L^G and L^R. The "blue"-"yellow" settings involved a tradeoff between L^B and $L^R + L^G$. A key point is that the observer could precisely cancel the effect of the secondary light arriving from the cube by adjusting primarily L^R vs. L^G. We expected that blue-yellow settings would not change systematically with changes in test patch orientation τ and, if this is so, it would simplify the analysis of the data.

On any given trial the initial color of the test patch was randomly assigned to be within a fixed distance from the neutral point in the blue-yellow and red-green space, with luminance held constant.

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 triplets for each pixel. These triplets were translated to a 24-bit graphics code, correcting for nonlinearities in the monitors' responses by means of measured look-up tables for each monitor.

Procedure

The observers repeated each of the seven conditions 20 times. The order of presentation of the stimuli was randomized. Observers completed the trials at their own pace. There was a short break after 70 trials. The entire experiment usually took the observer less than an hour.

Observers

Six observers participated in the study. All had normal or corrected-to-normal vision. None of the observers were aware of the hypothesis under test.

Instructions to the observer

Observers were asked to familiarize themselves with the scene, and to set the test patch to be achromatic. If the observer remained uncertain about the meaning of the term "achromatic," the experimenter explained the term with words such as "neutral" and with phrases "not red," "not green," etc. Observers did not practice before starting the experimental trials.

Analysis and results

Suppose that the observer has set the test patch to be achromatic. If the test patch were rendered as a neutral surface, then the light radiating from the test patch would be proportional to the incident light, the effective illuminant which, if we ignore inter-reflections beyond one-bounce, is $E(\lambda) = \gamma_0(\tau)E_{(0)}(\lambda) + \gamma_1(\tau)E_{(1)}(\lambda)$. So, the observer's setting $[L^R, L^G, L^B]$ should be the RGB code corresponding to $E(\lambda)$ which we compute below in terms of the parameters describing the geometry of the scene. We will compare these predictions to the observer's actual settings, which, from this point on, we denote $[\hat{L}^R, \hat{L}^G, \hat{L}^B]$. In our analysis we will also allow for the possibility that the observer's subjective achromatic point could be biased.

Relative red

The light emitted by the test patch is

$$L(\lambda) = \pi^{-1}E(\lambda)S_T(\lambda). \quad (9)$$

When the test patch is neutral, $L(\lambda) = \pi^{-1}E(\lambda)\alpha_T$, and we can rewrite this identity in terms of each of the three primaries. For example, for R we have

$$L^R = \pi^{-1}E^R\alpha_T^R. \quad (10)$$

We next introduce the ideal geometric red discounting function

$$\rho = \frac{L^R}{L^R + L^G + L^B} \quad (11)$$

which is the relative amount of red that is reflected from the test patch for a given orientation τ. Using Equations 10 and 7 we obtain

$$\rho = \frac{\gamma_0 E_{(0)}^R + \gamma_1 E_{(1)}^R}{\gamma_0 E_{(0)}^{Tot} + \gamma_1 E_{(1)}^{Tot}} \quad (12)$$

where we denoted the total illuminants with $E_{(0)}^{Tot}$ and $E_{(1)}^{Tot}$. We can further simplify the above equation. First of all recall that $E_{(1)}(\lambda) = E_{(0)}(\lambda)S_C(\lambda)$. This yields $E_{(1)}^R = E_{(0)}^R S_C^R$; $E_{(1)}^G = E_{(0)}^G S_C^G$; $E_{(1)}^B = E_{(0)}^B S_C^B$, and $E_{(1)}^{Tot} = E_{(1)}^R + E_{(1)}^G + E_{(1)}^B$. Recall that the central cube was rendered with reflectance $\alpha_C^R = 1, \alpha_C^G = 0.05$, $\alpha_C^B = 0$; therefore, we can neglect α_C^G compared to α_C^R. This yields

$$\rho = \frac{\gamma_0 E_{(0)}^R + \gamma_1 E_{(0)}^R}{\gamma_0 E_{(0)}^{Tot} + \gamma_1 E_{(0)}^R} \equiv B\frac{\gamma_0 + \gamma_1}{\gamma_0 + B\gamma_1}, \quad (13)$$

where we defined $B = E_{(0)}^R / E_{(0)}^{Tot}$ whose true value is 1/3.

Now, suppose that the observer views a scene illuminated by a neutral primary punctate light and by red secondary light due to a nearby red surface, the angle between the test patch and the neighboring red surface being τ. Then we can rewrite Equation 13, emphasizing the dependence on τ as,

$$\rho(\tau) = \frac{\gamma_0(\tau) + \gamma_1(\tau)}{3\gamma_0(\tau) + \gamma_1(\tau)}. \quad (14)$$

The observer is asked to adjust the chromaticity of the test without changing the luminance until it looks achromatic, just as in our experiment. If the observer correctly discounts for the orientation of the test patch then, when the perceived color of the test patch is gray, its ideal geometric discounting function would be given by Equation 14. However, if the observer does not perfectly discount the change in orientation, and if we repeat the test for many values of τ, the observer's settings would trace a curve $\hat{\rho}(\tau)$. We refer to this plot as the *observer's geometric red discounting function*, and define it as

$$\hat{\rho}(\tau) = \frac{\hat{L}^R(\tau)}{\hat{L}^R(\tau) + \hat{L}^G(\tau) + \hat{L}^B(\tau)} \quad (15)$$

where $[\hat{L}^R(\tau), \hat{L}^G(\tau), \hat{L}^B(\tau)]$ is the observer's mean achromatic setting for a particular value of τ. The dependent variable in our study was the relative amount of "red" in the observer's setting of the color of the test patch, $\hat{\rho}$, although the term "red" here refers precisely to the chromaticity of the added secondary illuminant, which is precisely the chromaticity of the cube. We were particularly interested in whether the amount of red in the setting is affected by the orientation of the test patch (i.e., whether the observer "discounts" the angle from the perceived color of the test patch). We compared the observers setting $\hat{\rho}(\tau)$ to the prediction $\rho(\tau)$ of the one-bounce model of inter-reflection between two Lambertian surfaces derived above.

As we noted above in "Methods," we verified that the effects of a higher number of bounces on the effective illuminant would be negligible by comparing the four-bounce and one-bounce Radiance renderings. This negligible difference between one-bounce and four-bounce rendering leads to a slight (approximately 0.03 in relative red units) overall upward shift of the geometric red discounting function. We further compared numerically the ideal geometric red discounting function obtained from a one-bounce model (Equation 14) against the four-bounce Radiance renderings. This comparison yielded a MSE of 0.000289 (in normalized RGB values). Graphs of this comparison can be obtained from the authors.

Possible outcomes

It is possible that the observer might adopt a binary heuristic to compensate for mutual illumination (i.e., add a fixed chromaticity correction independent of the orientation of the test patch precisely when inter-reflection was physically possible; $\tau < 180^\circ$). Such a heuristic could permit more accurate surface color estimation without the need to compute the geometric factors (e.g., Equations 3

and 6) explicitly. It is also consistent with results of the experiment of Bloj et al. (1999). With this heuristic we would expect the relative amount of red to be a constant across the different test patch orientations as in Figure 10A (horizontal, red line). If, on the other hand, the observer correctly takes into account the orientation of the test patch when making his or her setting, we would expect the estimates to be close to the curve predicted by the model, $\rho(\tau)$, in Figure 10B.

Data

The results for all six observers are shown in Figure 11. If the observers' settings $\hat{\rho}$ were in agreement with the physically correct one-bounce model, they would fall on the curve given by Equation 14 (with true values inserted). This curve, $\rho(\tau)$, is plotted red in the data graphs.

It is apparent that the observers take the orientation of the test patch into account when discounting the mutual illumination between the cube and the test patch and that the observers' geometric discounting function is roughly a scaled and shifted copy of the geometric discounting function $\rho(\tau)$ for the one-bounce model, yet there is evident inter-observer variability. A number of observers have a geometric discounting function that is flatter than the model prediction. This "flattening" is particularly evident for angles greater than 90°. Parameters that influence the shape of the curve are discussed below. For all observers, we verified that there is, as expected, no systematic change in the *Blue/Green* ratios across angles. We do not plot these results or discuss them further.

Maximum likelihood estimation

It is possible the systematic differences we see between $\rho(\tau)$ and $\hat{\rho}(\tau)$ are the result of systematic errors in the observers' estimates of scene properties. For example, an observer may misperceive the orientation τ of the test patch. We can refit the data allowing for this possibility by adding parameters $\hat{a}$ and $\hat{b}$ where $\hat{\tau} = \hat{a} + \hat{b}\tau$. If $\hat{a}$ proved to be close to 0, and $\hat{b}$ close to 1, then we could not attribute the systematic differences between $\rho(\tau)$ and $\hat{\rho}(\tau)$ to misperception of τ. Conversely, we may be able to account for these differences as the result of a misperception of τ.

We also added a parameter that estimated the proportion of the cube that was used for integration, $\hat{w}$. The observer may underestimate or overestimate the area of integration that is relevant (Equation 6). For the simplification of our calculations, we assumed this area to be square. Note, however, that this need not necessarily be the case. Last, we included as a constant $\hat{B} = \left(\widehat{E^{R}_{(0)} / E^{Tot}_{(0)}}\right)$ that we refer to as bias in the observer's setting (see Equation 13). This bias indicated whether the observer, independent of angle τ, put too much or too little red (when compared with the ideal model) into her/his setting. The observers' geometric discounting function becomes

$$\hat{\rho} = \hat{B}\frac{\gamma_0(\hat{\tau}) + \gamma_1(\hat{\tau}, \hat{w})}{\gamma_0(\hat{\tau}) + \hat{B}\gamma_1(\hat{\tau}, \hat{w})} \quad (16)$$

and we wish to bring this family of curves into coincidence with the data by choice of the setting of the three parameters. We used maximum likelihood fitting procedures to estimate the values of the parameters for the observers' data. The estimates are shown in Table 1.

Under the assumption of the Lambertian model and given the results of our fitting procedure it appears that most observers did not integrate over the entire cube's surface, but used only a part of it ($\hat{w} < w$). Furthermore the angle τ between the test patch and cube was, except for two subjects, perceived as slightly compressed. This finding is in agreement with research that maintains that observers tend to perceive the orientation of a rectangular Lambertian patch as slightly compressed in depth (see Boyaci et al., 2003, for discussion). However, none of the six observers' estimates of τ were significantly different from the veridical values. The amount of bias varied between observers and, overall, it was relatively small (Table 1).

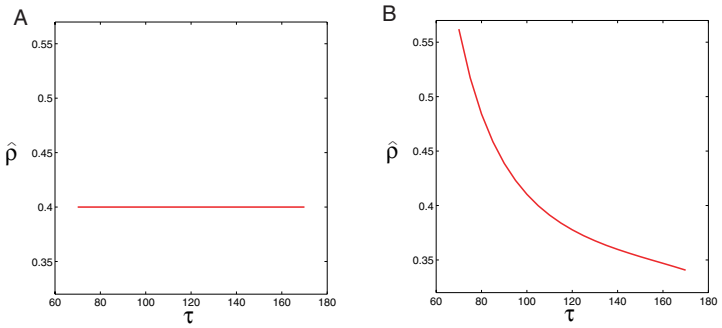

Figure 10. Hypothetical inter-reflection discounting functions. The angle τ between the test patch and the cube surface is plotted on the horizontal axis. The geometric discounting factor $\hat{\rho}$ as defined in the text is plotted on the vertical axis. If observers only added a fixed chromaticity correction, independent of the orientation of the test patch (binary heuristic), we would expect that their settings would lie on a straight line as the one shown in Figure 10A. If, however, observers correctly discount the inter-reflection between the cube and test patch their settings would fall on the red curve shown in 10B.

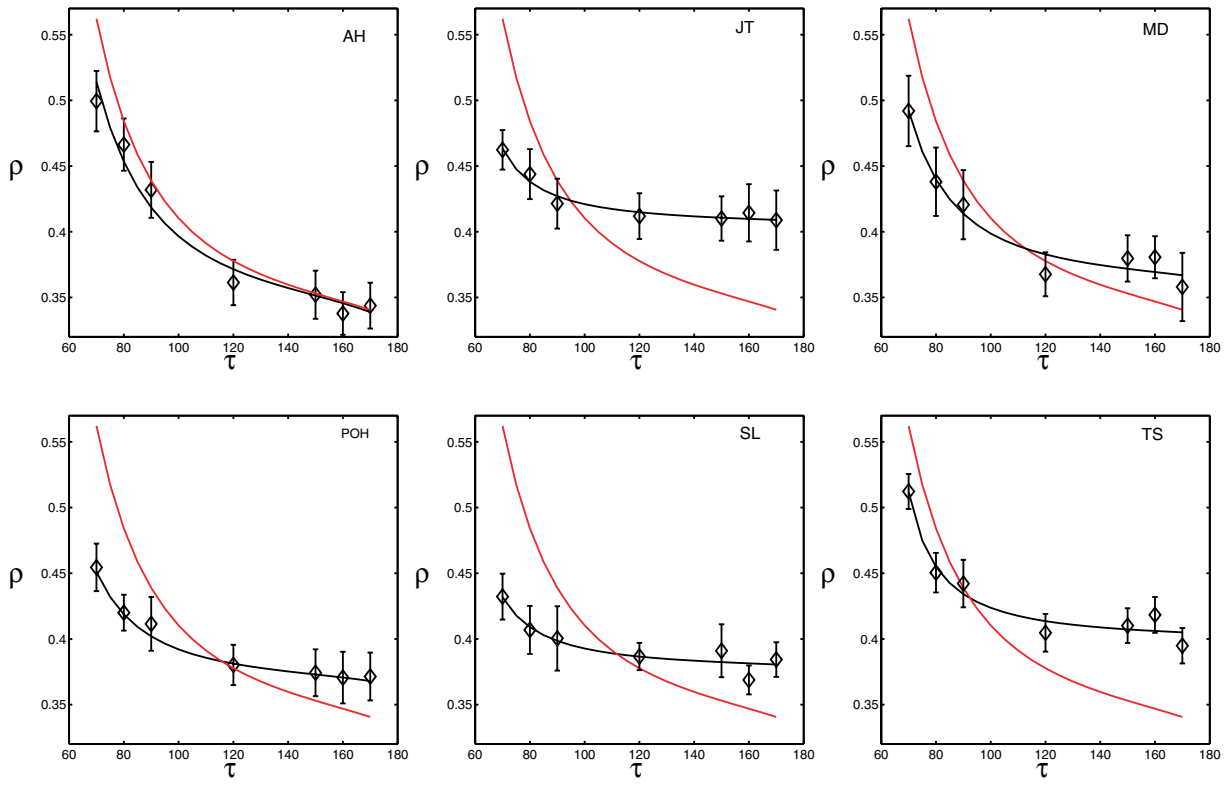

Figure 11.Results for six observers. The axes are as in Figure 10. The black-outlined diamonds show the observer's mean setting with error bars (±1 SEM). The correct geometric discounting function from Figure 10 is replotted in red. The black line through the diamonds describes the maximum likelihood fit of the model described in the text.

Hypothesis testing

We used a nested hypothesis test (Mood, Graybill, & Boes, 1974, pp.440) to test the hypothesis that the observers' geometric discounting function is constant. If it were true then the observer would not take into account the angle when discounting the mutual illumination between cube and test patch. The log likelihood of the unconstrained model λ_1 was obtained by fitting the geometric discounting function to the observers' relative red settings using the method of maximum likelihood (the four parameters described above were free to vary). In the constrained ("nested") model, we forced the geometric relative red function to be constant (red line in Figure 10A), allowing only the bias $\hat{B}$ to vary. We compared the log likelihood ratio to the relevant chi square distribution (χ^2_3). All observers' relative red settings were significantly different from a constant geometric relative red function (p<.00001).

We furthermore tested the hypothesis that the observers' estimate of the geometric discounting function is equal to the "ideal" geometric discounting function utilizing a similar nested hypothesis test. The log likelihood of the unconstrained model λ_1 was obtained as above. For our constrained model λ_0 we assigned $\hat{w}$, $\hat{a}$, and $\hat{b}$ their true values: 2.18, 0, and 1, respectively, and let $\hat{B}$ vary freely.[3] Comparing the resulting log likelihood ratio to the appropriate chi square distribution (χ^2_3) we find that for all observers the relative red settings were significantly different from the predictions of the model (p < .00001), leading to rejection of the null hypothesis.

In addition, we tested the hypothesis that the observers' estimate of given parameter ($\hat{w}$, $\hat{B}$, $\hat{a}$ and $\hat{b}$) was equal to the veridical value (2.18, 0.3333,0, and 1 respectively). The log likelihood of unconstrained model was obtained as described above. In the constrained model, we assigned the parameter in question its veridical value and let the other parameters vary freely. We compared the resulting log likelihood ratio to the chi-square distribution with one degree of freedom (χ^2_1). All p values are reported in Table 1.

Discussion

We conclude from our experimental results that observers systematically take into account the angle between the brightly colored cube and test patch when discounting mutual illumination. However, their estimates of the amount of red to be discounted with change in angle deviates significantly from the predictions of the one-bounce model of mutual illumination ("the ideal geometric discounting function"). For our choice of stimuli, the predictions of a two- or higher bounce model are little different, and we cannot explain our observers' responses by assuming that they discount more than one bounce. We consider

Observer	$\hat{w}$	$\hat{B}$	$\hat{a}$	$\hat{b}$
Veridical	**2.18**	**0.3333**	**0**	**1**
MD	**1.121***	**0.363***	**-0.000006**	**0.978**
	p<0.001	p <0.001	p =0.849	p =0.849
POH	**1.045***	**0.369***	**-0.000011**	**1.059**
	p <0.001	p <0.001	p =0.637	p =0.637
AH	**1.947***	**0.338**	**-0.000007**	**1.059**
	p <0.001	p =0.387	p =0.093	p =0.093
JT	**0.622***	**0.407***	**-0.000016**	**0.922**
	p <0.001	p <0.001	p =0.706	p =0.706
SL	**0.661***	**0.379***	**-0.000013**	**0.958**
	p <0.001	p <0.001	p =0.89	p =0.89
TS	**0.677***	**0.401***	**-0.000016**	**0.849**
	p <0.001	p <0.001	p =0.08	p =0.08

Table 1. Maximum likelihood estimates for six observers. The three parameters correspond to three possible patterns of deviations from veridical discounting of inter-reflected light. If $\hat{b}$ is not equal to 0 or $\hat{a}$ is not equal to 1, then the observer is misperceiving the orientation of the test patch with respect to the cube. (Please note that $\hat{b}$ and $\hat{a}$ were either both fixed or both varied. We constraint our ML analysis in a way such that $\hat{a}+\hat{b}\tau \leq 180$.) The parameter $\hat{w}$ is the dimension of the cube's surface over which the observer is integrating in estimating the intensity of one-bounce illumination. All but one observer markedly underestimated the correct area of integration. A deviation of the parameter $\hat{B}$ from 1/3 corresponds to a shift in neutral point, possibly due to adaptation to the stimulus across the duration of the experiment. A value greater than 1/3 indicates that the observer is effectively "red adapted." The correct values are also listed in the row labeled "veridical." An entry marked with an asterisk is significantly different from the corresponding model value (with a Bonferroni correction for 24 tests).

the possibility that the differences between observed and ideal performance are due to misperception of specific physical parameters characterizing the scenes we used.

Influence of the individual factors

Factors that influence the shape of the geometric discounting function include the area of the cube's surface that is utilized as a secondary light source (by means of $\hat{w}$), veridical perception of the surface normal of the test patch $(\hat{a},\hat{b})$ and constant over- or under-estimation of the amount of red in the scene ($\hat{B}$).

Influence of $\hat{a},\hat{b}$ **:** Based on the estimates of parameters $\hat{a}$ and $\hat{b}$, we conclude that observers perceived the orientation of the test patch nearly veridically, with a slight tendency to underestimate τ. This finding is consistent with the results of Boyaci et al. (2003), who used a similar stimulus configuration.

Influence of $\hat{B}$ **:** The observed bias in most of the observers' model fits is rather small. The bias parameter shifts the entire discounting curve up (when putting too much red in the achromatic setting) or down (this actually never occurred), it can be interpreted as a slight overestimation (or underestimation) of the amount of red in the light source that counts as neutral. This might be brought about, for example, by means of chromatic adaptation to the scene (specifically to the large bright orange cube) or may simply mean that the observer disagrees with our arbitrary choice of "neutral."

Influence of $\hat{w}$ **:** The fitted values of $\hat{w}$ indicate that most observers use less than the optimal area of the cube for their discounting function. A value less than veridical corresponds to a compression of the observer's geometric relative red function, leading to an overall flattening of the curve. If we attribute this to a failure to choose the correct area of integration, the data suggest that all observers utilize less than the relevant part of the cube.

Of course, the effect of $\hat{w}$ is confounded with any other factor that would lead the observer to underestimate the overall intensity of the secondary illumination from the cube. One possibility is that the visual system does not use the Lambertian model in computing the intensity of the secondary ("one bounce") illumination.

We chose a large, highly saturated orange cube to improve our chances of seeing any discounting of the secondary illuminant. However, it may well be that the visual system simply does not cope well with such extremes of secondary illumination. The pattern of failure across angle hints that this may be the case. In everyday situations with less saturated chromatic surfaces, inter-reflection might only have an appreciable effect at acute angles and could be negligible at obtuse angles ($\tau > 90^\circ$). In the model, the slope of the curve decreases greatly after 90°, it might be that the observer is not sensitive to this fine gradient, and anchors on a constant setting for wider angles. It is quite possible that observers may compute inter-reflection between two surfaces differently, that is, they assume different models for small and wide angles, adopting a variant of the binary heuristic for angles greater than 90°.

It would be particularly interesting if we could interpret the errors in estimated $\hat{w}$ as visual errors in choosing the limits of integration. By making the chromatic surface flat and rectangular, we set up conditions where we could compute the effective illumination with relative ease. Had we picked a curved surface or a surface with an irregular boundary instead, then our computation would have been more difficult. It is presumably this more general problem that the visual system addresses and, consequently, the problem of selecting the proper area of integration and computing the geometric factor $\gamma_1(\tau)$ is plausibly difficult and prone to error. Under this interpretation, the deviations in $\hat{w}$ are less surprising.

Can a local contrast model account for observers' settings?

Perhaps it is the case that observers used the immediate surround of the test patch for making their achromatic settings, utilizing information about local contrast and not using information about scene geometry at all. To prevent

the utilization of this cue, we rendered the immediate surround of the test patch with a very small albedo value (α=0.01). Further, evidence that local contrast cannot solely account for the achromatic shifts has also been brought by studies by several other researchers (see Delahunt, 2001; Brainard, 1998; Bloj, 1999).

To summarize, this study demonstrates that observers are qualitatively discounting an effective illuminant whose chromaticity depends upon the geometric layout of the scene. Color is not a local phenomenon but is contingent on global context, such as scene geometry and global lighting conditions. In our experiment, observers were able to make physically sensible adjustments to achieve a constant percept of surface color.

Appendix

Multiple scattering of light and secondary illumination

In a real scene, light from a primary source is scattered by the surfaces present in the scene and some of this scattered light contributes to the illumination of other surfaces. Some part of the primary visible light is reflected back and forth between the surfaces until it escapes the scene or is completely absorbed by the surfaces. Therefore, the total amount of light falling on a surface is the sum of contributions of light coming directly from the primary source and that arriving after reflected by other surfaces, possibly many times.

The flux of light emitted in a certain direction by an infinitesimal surface element $d\sigma$ around a point $\mathbf{r}$, at a given wavelength λ is (LeGrand, 1957, pp. 18ff),

$$L(\mathbf{r};\lambda) = \frac{1}{\pi} E(\mathbf{r};\lambda) S(\mathbf{r};\lambda)\, d\sigma, \tag{17}$$

where $E(\mathbf{r};\lambda)$ is the illumination (surface density of the light flux received) and $S(\mathbf{r};\lambda)$ is the surface reflectance function at the point $\mathbf{r}$. In the following derivations, wavelength λ is not displayed for simplicity. The total illumination upon the surface element at $\mathbf{r}$ satisfies the equation,

$$E(\mathbf{r}) = E_{(0)}(\mathbf{r}) + \int_{\Omega} E(\mathbf{r}') S(\mathbf{r}') G(\mathbf{r}-\mathbf{r}';\mathbf{n}(\mathbf{r}))\, d\sigma', \tag{18}$$

where $E_{(0)}(\mathbf{r})$ is the illumination due to the primary source and Ω the area of integration. The integral is over all surfaces which contribute to the illumination of the surface element at $\mathbf{r}$, and

$$G(\mathbf{r}-\mathbf{r}';\mathbf{n}(\mathbf{r})) = \frac{1}{|\mathbf{r}-\mathbf{r}'|^2} \mathbf{n}(\mathbf{r}) \cdot \frac{(\mathbf{r}-\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|} = \frac{\cos\theta}{|\mathbf{r}-\mathbf{r}'|^2}, \tag{19}$$

is a geometric factor, where the angle of incidence θ is the angle between the normal vector to the surface at the point $\mathbf{r}$ and the vector connecting the points $\mathbf{r}$ and $\mathbf{r}'$. $E_{(0)}(r)$ is equal to $I_{(0)} d^{-2} \cos\theta$ where $I_{(0)}$ is the intensity of the primary source, d is the distance from the point r to the primary source, and θ is the angle of incidence of the primary source on the infinitesimal surface element at r. The first-order ("one-bounce") approximation to Equation 18 is obtained by inserting $E_{(0)}(\mathbf{r}')$ in place of $E(\mathbf{r}')$ in the integral

$$E_{(1)}(\mathbf{r}) = E_{(0)}(\mathbf{r}) + \int_{\Omega} E_{(0)}(\mathbf{r}') S(\mathbf{r}') G(\mathbf{r}-\mathbf{r}';\mathbf{n}(\mathbf{r}))\, d\sigma'. \tag{20}$$

Naturally, a better approximation is obtained by inserting the first-order approximation from Equation 20 into the integral in Equation 18

$$E_{(2)}(\mathbf{r}) = E_{(0)}(\mathbf{r}) + \int_{\Omega} E_{(1)}(\mathbf{r}') S(\mathbf{r}') G(\mathbf{r}-\mathbf{r}';\mathbf{n}(\mathbf{r}))\, d\sigma'. \tag{21}$$

This is the 2^{nd} order or "two-bounce" approximation. Explicitly writing the second order approximation yields

$$\begin{aligned} E_{(2)}(\mathbf{r}) &= E_{(0)}(\mathbf{r}) + \int_{\Omega} S(\mathbf{r}') G(\mathbf{r}-\mathbf{r}';\mathbf{n}(\mathbf{r}))\, \mathrm{d}\sigma' \\ &\times \left[E_{(0)}(\mathbf{r}') + \int_{\Omega} E_{(0)}(\mathbf{r}'') S(\mathbf{r}'') G(\mathbf{r}'-\mathbf{r}'';\mathbf{n}(\mathbf{r}'))\, d\sigma'' \right] \\ &= E_{(0)}(\mathbf{r}) + \int_{\Omega} E_{(0)}(\mathbf{r}') S(\mathbf{r}') G(\mathbf{r}-\mathbf{r}';\mathbf{n}(\mathbf{r}))\, d\sigma' \\ &+ \int_{\Omega}\int_{\Omega} E_{(0)}(\mathbf{r}'') S(\mathbf{r}'') G(\mathbf{r}'-\mathbf{r}'';\mathbf{n}(\mathbf{r}')) S(\mathbf{r}') G(\mathbf{r}-\mathbf{r}';\mathbf{n}(\mathbf{r}))\, d\sigma' d\sigma''. \end{aligned} \tag{22}$$

as we improve the approximation by repeating the recursion, the *n-th* approximation will involve integrals taken over the region of interest once, twice,... and *n* times. As the order of approximation increases the expression gets more complicated, therefore we symbolically write it as

$$
\begin{aligned}
E(\mathbf{r}) &= E_{(0)}(\mathbf{r}) \\
&+ E_{(0)}(\mathbf{r}) \int SG \\
&+ E_{(0)}(\mathbf{r}) \iint SGSG \\
&+ \ldots \\
&+ E_{(0)}(\mathbf{r}) \underbrace{\iint \ldots \int SGSG \ldots SG}_{\text{n factors of } SG}
\end{aligned}
\tag{23}
$$

The physical significance of this expression is the following. The first term in Equation 23 is the direct illumination due to the primary source. The second term represents the response of a small area element at the point $\mathbf{r}'$ with area $d\sigma'$ around. It acts as an effective source that makes a contribution $E_{(0)}(\mathbf{r}')S(\mathbf{r}')G(\mathbf{r}-\mathbf{r}';\mathbf{n}(\mathbf{r}))d\sigma'$ to the field at another point $\mathbf{r}$. The higher order terms can also be interpreted in the same way as contributions from higher number of scatterings (up to the n-bounce term).

In the experiment, we have two nearby Lambertian surfaces that are relevant to us: the face of the orange cube, and the small gray test patch. We ignore any contribution from other surfaces in the scene. Hence the region of integration Ω involves these two surfaces only. We consider the illumination of the small achromatic test patch. In the experimental scene, the punctate light source is sufficiently far away to allow us assume that its distance is constant across the surfaces of the test patch and the cube. This allows us further simplify the equation by introducing the "*effective intensity*" of the punctate source as

$$
E_{(0)} = I_{(0)} d^{-2}. \tag{24}
$$

Then, for example, for zero-bounce we obtain

$$
E = E_{(0)} \cos\theta. \tag{25}
$$

Let us next consider the first-order approximation to the illumination of a point $\mathbf{r}$ on the test patch. Equation 25 will constitute the first term of the approximation in Equation 20. The region of integration Ω will be the face of the cube "visible" to the test patch, denoted by C. The zero-th-order illumination of a point $\mathbf{r}'$ on the cube surface is

$$
E(\mathbf{r}') = E_{(0)} \cos\theta_{CP} \tag{26}
$$

where θ_{CP} is the angle of incidence of the light from the primary source on the cube's surface at $\mathbf{r}'$. The surface reflectance of the cube is a constant throughout the face of the cube, $S(\mathbf{r}') = S_C$. The geometric factor is

$$
G = |\mathbf{r}-\mathbf{r}'|^{-2} \cos\theta_{TC}(r'). \tag{27}
$$

Note that θ_{TC} is a function of $\mathbf{r}'$ but θ_{TP} θ_{CP} are not (because we assume that the light arriving at the test patch and the cube's surface is collimated). Inserting Equations 26 and 27 into Equation 20, we obtain the first-order approximation

$$
\begin{aligned}
E(\mathbf{r}) &= E_{(0)}(\mathbf{r}) \cos\theta_{TP} \\
&+ E_{(1)}(\mathbf{r}) \int_C \cos\theta_{CP} \cos\theta_{TC}(r') |\mathbf{r}-\mathbf{r}'|^{-2} \, d\sigma
\end{aligned}
\tag{28}
$$

where we also introduced $E_{(1)} = E_{(0)} S_C$ and took it out of the integration since it is a constant. Similarly, the second-order approximation is found as follows

$$
\begin{aligned}
E(\mathbf{r}) &= E_{(0)}(\mathbf{r}) \cos\theta_{TP} \\
&+ E_{(1)}(\mathbf{r}) \int_C \cos\theta_{CP} \cos\theta_{TC}(r') |\mathbf{r}-\mathbf{r}'|^{-2} \, d\sigma_C \\
&+ E_{(2)} \int_C \left[\int_T \cos\theta_{TP} \cos\theta_{CT} \, |\mathbf{r}'-\mathbf{r}''|^{-2} \, \mathrm{d}\sigma_T \right] \\
&\times \cos\theta_{TC}(r') \, |\mathbf{r}-\mathbf{r}'|^{-2} \, d\sigma_C .
\end{aligned}
\tag{29}
$$

If we continue in this fashion, we obtain $E = \sum_{i=0}^{n} \gamma_i E_{(i)}$ as the n-th-order approximation which contains one, two,..., n bounces of the primary light before reaching the test patch. The coefficients γ_i are calculated as in Equations 28 and 29. The $E_{(i)}$ s are related by the recurrence relation

$$
\begin{aligned}
E_{(1)} &= E_{(0)} S_C; \\
E_{(i+2)} &= E_{(i)} S_C S_T
\end{aligned}
\tag{30}
$$

In the analysis of the data, we assume a first-order approximation and ignore higher order terms. Here, we will show the derivation of the one-bounce term γ_1. We will calculate one-bounce illumination only in the center of the test patch. Because the test patch is very small, we will assume that the illumination is roughly constant across its surface. Application of the law of cosines, $\cos\theta_{\mathbf{ab}} = \mathbf{a}\cdot\mathbf{b} / |\mathbf{a}||\mathbf{b}|$ yields the relevant angles of incidence in Equation 28

$$
\begin{aligned}
\cos\theta_{CP} &= \cos\varphi_P \sin\psi_P , \\
\cos\theta_{TP} &= \cos\varphi_P \sin(\tau - \psi_P), \\
\cos\theta_{TC} &= x \sin\tau \, \Delta_{CT}^{1/2}, \\
\text{where } \Delta_{\mathrm{CT}}^{1/2} &= \frac{1}{\sqrt{x^2 + z^2 + d^2 - 2xd\cos\tau}},
\end{aligned}
\tag{31}
$$

where (ψ_P, φ_P) is the direction to the punctate source as described above, d is the distance from the origin to the center of the test patch. With these in place, γ_0 and γ_1 become

$$\gamma_0 = \cos\varphi_P \sin(\tau - \psi_P)$$

$$\gamma_1 = \cos\varphi_P \sin\psi_P \sin\tau \int_0^w \int_{-w}^{w} x\, \Delta_{CT}^3 dx\, dz \tag{32}$$

where we have assumed that the integration region C is taken as a square of area $2w \times w$. Note that this integral can be solved analytically, indicating that the computation of the geometric factors of inter-reflection need not always involve explicit integration. The integration yields

$$\gamma_1 = \cos\varphi_P \sin\psi_P \left\{ \sin\tau \left(\ln \frac{w - \sqrt{2w^2 + d^2 - 2wd\cos\tau}}{w + \sqrt{2w^2 + d^2 - 2wd\cos\tau}} - \ln \frac{w - \sqrt{w^2 + d^2}}{w + \sqrt{w^2 + d^2}} \right) + 2\cos\tau \left(\arctan \frac{w(w - d\cos\tau)}{d\sin\tau\sqrt{2w^2 + d^2 - 2wd\cos\tau}} + \arctan \frac{w\cos\tau}{\sin\tau\sqrt{w^2 + d^2}} \right) \right\} \tag{33}$$

This is the equation plotted in Figure 3 as a function of τ.

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: Katja Doerschner.

Address: Department of Psychology, New York University, New York, NY, USA.

Email: kd462@nyu.edu.

Footnotes

[1] The $1/\pi$ term in Equation 6 stems from calculating the *radiant emittance in all directions* (after LeGrand, 1957).

[2] For the conditions of Figure 1 and the experiment reported here, the double integral of Equation 6 can be solved in closed form. We report these results and derive them in the "Appendix."

[3] Note that $w = 2.18d$, where d is the distance from the edge where the test patch touches the cube (the origin of the coordinate system), to the center of the test patch.

References

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. [Article] [PubMed]

Boyaci, H., Doerschner, K., & Maloney L. T. (in press). Perceived surface color in binocularly-viewed scenes with two light sources differing in chromaticity. *Journal of Vision*.

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]

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

Delahunt, P. B.(2001). An evaluation of color constancy across illumination and mutual reflection changes. Ph.D. Thesis, University of California, Santa Barbara. [Dissertation]

Endler, J. A. (1993). The color of light in forests and its implications. *Ecological Monographs*, *63*, 1027.

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

Hurlbert, A. (1998). Computational models of color constancy. In V. Walsh & J. Kulikowski (Eds.), *Perceptual constancy: Why things look as they do*. Cambridge: Cambridge University Press.

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

LeGrand, Y. (1957). *Light, colour and vision*. London: Chapman & Hall.

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: Cambridge University Press.

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

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

Werner, J. S., & Walraven, J. (1982). Effect of chromatic adaptation on the achromatic locus: The role of contrast, luminance and background color. *Vision Research*, *22*, 929-943. [PubMed]

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