Illumination estimation in three-dimensional scenes with and without specular cues

Jacqueline Leigh Snyder Department of Psychology, New York University, New York, USA

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

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

We report the results of three experiments in which observers judged the albedo of surfaces at different locations in rendered, three-dimensional scenes consisting of two rooms connected by a doorway. All surfaces composing the rooms were achromatic and Lambertian, and a gradient of illumination increased with depth. Observers made asymmetric albedo matches between a standard surface placed in the rooms at different depths along the line of sight and an adjustable surface at a fixed location. In Experiment 1, gradients of intensity on the walls, floor, and ceiling of the scene, as well as its three-dimensional structure, provided information about variations in the intensity of illumination across depth (the illumination profile). In Experiment 2, specular spheres provided an additional veridical cue to the illumination profile. We sought to determine whether observers would make use of this additional cue. They did: all observers exhibited a greater degree of lightness constancy in Experiment 2 than in Experiment 1. In Experiment 3, the specular spheres reflected an illumination profile in conflict with that signaled by the other cues in the scene. We found that observers chose albedo matches consistent with an illumination profile that was a mixture of the illumination profiles signaled by the specular spheres and by the remaining cues.

Keywords: albedo, cue combination, depth perception, illumination estimation, lightness, lightness constancy, specularity, 3-D vision

Introduction

The illumination profile

In de Hooch's (1659–1660) masterpiece of domestic tranquility, the artist rendered a scene rich with light sources and surface materials (Figure 1). Rays of sunlight illuminating the child and the room in which she stands do not also illuminate the spaniel. Less brilliant light, whose source the artist concealed from the viewer, highlights the mother's countenance. Throughout the scene, the illumination varies from place to place, generally increasing in intensity as the scene recedes in depth. The artist rendered myriad materials, from shiny to smooth, which provide the viewer with cues to the illumination. The polished checkerboard-tiled floor may be the most salient cue to changing illumination across depth.

Imagine that you could place a piece of gray paper anywhere in the scene, at any orientation. Let us assume that this piece of paper is perfectly matte, or *Lambertian,* so that it absorbs the light which falls upon it and reradiates this light uniformly in all directions (Lambert, 1760). If the surface changes location and orientation within the scene and the light falling upon it changes, the luminance radiated toward the observer will also change. From the perceptual point-of-view of the observer, how will the surface's appearance change? Will the gray surface appear darker or lighter or remain the same?

The light radiated by the Lambertian surface depends on the net intensity of the light absorbed by the surface at a given location and orientation, and the surface's albedo, or proportion of light it reflects. The light that reaches an observer's eyes from the Lambertian surface is independent of the angle created by the surface's orientation and the observer's line of sight. Lambert's Law (1760) describes how light arriving at the surface from different directions contributes to the net intensity of light absorbed.

Consider a small Lambertian surface at location $p = (x, y, z)$ with unit normal vector $\boldsymbol{\nu}$ (Figure 2). The plenoptic function of Adelson & Bergen (1991), $\Pi(p,\boldsymbol{\nu})$, specifies the amount of light arriving at point p along the direction specified by the unit vector $\boldsymbol{\nu}$. We

doi: 10.1167/5.10.8 Received April 12, 2005; published December 15, 2005 ISSN 1534-7362

can write the total light incident upon the Lambertian surface as

$$\Lambda(p,\boldsymbol{\nu}) = \int_{H(\boldsymbol{\nu})} \Pi(p,\boldsymbol{\upsilon}) < \boldsymbol{\upsilon}, \boldsymbol{\nu} > d\boldsymbol{\upsilon}, \quad (1)$$

where $<\boldsymbol{\upsilon}, \boldsymbol{\nu}>$ denotes the inner product of the two unit vectors (equal to the cosine of the angle, θ, between them) and $H(\boldsymbol{\nu})$ denotes the hemisphere around p centered on $\boldsymbol{\nu}$. In limiting the area of integration to the hemisphere $H(\boldsymbol{\nu})$, we assume that the surface is opaque: no light absorbed on one side is emitted through the other. The function $\Lambda(p,\boldsymbol{\upsilon})$ is equivalent to $\Pi(p,\boldsymbol{\upsilon})$ smoothed by convolution with the Lambertian kernel of Basri & Jacobs (2003). We refer to $\Lambda(p, \boldsymbol{\nu})$ as the *illumination profile* of the scene.

The Lambertian surface re-emits the light, $\Lambda(p, \boldsymbol{\nu})$, absorbed from all directions within the hemisphere, $H(\boldsymbol{\nu})$, after weighting it by the surface albedo, α. The intensity that arrives at the observer's eyes from the surface is proportional to

$$\mathrm{L}(p,\boldsymbol{\nu}) = \Lambda(p,\boldsymbol{\nu})\alpha. \quad (2)$$

For readers familiar with the literature on reflectance, it may be helpful to know that the illumination profile is a model of a bi-directional reflectance distribution function (BRDF). For further intuition about the interaction of illumination and albedo, we refer the reader to the more extensive review in Boyaci, Maloney, & Hersh (2003).

The ideal observer for lightness constancy

An observer who has access to the corneal illumination, $\mathrm{L}(p,\boldsymbol{\upsilon})$, and the illumination profile, $\Lambda(p,\boldsymbol{\nu})$, can compute the surface albedo, α, by discounting the illumination profile. Solving Equation 2 for α, we have

$$\alpha = \frac{\mathrm{L}(p,\boldsymbol{\nu})}{\Lambda(p,\boldsymbol{\nu})}. \quad (3)$$

For example, in Figure 1, if this ideal observer with perfect lightness constancy can accurately estimate $\Lambda(p,\boldsymbol{\nu})$ at any point p in the three-dimensional scene, then she can accurately estimate the albedo of a Lambertian achromatic surface at that point p with orientation $\boldsymbol{\nu}$.

Equivalent illumination profile (EIP)

Whether or not the observer's albedo estimation performance is ideal, we can use $\hat{\alpha}$ to compute an estimate of the observer's EIP, $\Lambda(p,\boldsymbol{\nu})$:

$$\hat{\Lambda}(p,\boldsymbol{\nu}) = \frac{\mathrm{L}(p,\boldsymbol{\nu})}{\hat{\alpha}}. \quad (4)$$

If the observer accurately estimates albedo (i.e., is lightness constant), then the observer's EIP will closely approximate the true illumination profile, $\Lambda(p,\boldsymbol{\nu})$.

Previous experimental work

There is considerable evidence that observers discount changes of illumination across a scene when adjacent regions of the scene are illuminated differently (Katz, 1935) and within smooth gradients of light (Arend, 1994). In three-dimensional scenes like Figure 1, light intensity can change dramatically in depth. There are several reports in the lightness literature suggesting that the three-dimensional interpretation of a scene can profoundly affect perceived lightness (Gilchrist, 1977, 1980, 1994; Gilchrist, Delman, & Jacobsen, 1983). Yamauchi & Uchikawa (2005) reported that coplanar information in a three-dimensional scene interacts with adjacent regions and illumination in the formation of observers' color appearance judgments.

There are also studies that compare the true illumination profile and the EIP derived from human performance. Boyaci et al. (2003) found that observers partially discounted changes in illumination intensity associated with changes in surface orientation. Boyaci, Doerschner, & Maloney (2004) found that observers partially discounted changes in chromaticity in scenes where the chromaticity of the illumination was not spatially uniform. The last two studies varied surface orientation, $\boldsymbol{\nu}$, but not surface location, p, within the scene. In this study, we address a question complementary to the work of Boyaci et al. (2003, 2004): Do observers estimate and discount $\Lambda(p,\boldsymbol{\nu})$ as a function of p with $\boldsymbol{\nu}$ held constant?

The second question we address is, under conditions where observers exhibit some amount of lightness constancy, how do they estimate the EIP? To introduce this question, we compare the outcomes of two prior studies. Schirillo & Shevell (1993) simulated scenes consisting of only two achromatic walls at different depths, and observers viewed these scenes in a Wheatstone stereoscope. A patchwork of rectangular surfaces of assorted dimensions and albedos composed the walls; researchers in vision science and computer vision often refer to these patchwork patterns as Mondrians.[1] The patchwork designs of the near and far Mondrians were identical except that the far Mondrian's surfaces were 500% more luminous than corresponding surfaces in the near Mondrian.

Observers could interpret the lightness of the far Mondrian's surfaces in one of two ways: as having lightness five times greater than the surfaces composing the near Mondrian, or as identical to the near Mondrian but illuminated by a light source five times more intense.

Figure 1. A monochromatic reproduction of Pieter de Hooch (1659–1660), *The Mother,* by permission, Gemäldegalerie, Staatliche Museen zu Berlin, Berlin, Germany. Original color photograph by Joerg P. Anders. This painting exemplifies a scene in which the illumination changes dramatically across depth.

In terms of our theoretical framework, the first interpretation is consistent with an EIP that is highly inaccurate and that will result in poor lightness constancy, while the second interpretation is consistent with an EIP accurate enough to support a high degree of lightness constancy. Schirillo & Shevell (1993) asked observers to match the lightness of test squares placed at near and far depths. (These researchers also asked observers to match brightness, but we are not concerned with these results.) Only one of their observers behaved as if the equivalent illumination at the depth of the near plane were different from the equivalent illumination at the depth of the far plane. For that observer the effect was small compared to the 5:1 change in simulated illumination. Schirillo and Shevell's results favor a no lightness constancy model with change in depth.[2]

In contrast, Mizokami, Ikeda, & Shinoda (1998) found a large effect of lightness constancy with change in depth. They measured apparent lightness for surfaces at different depths in configurations composed of two small real rooms arranged in depth. The lighting of the rooms was complex, consisting of multiple fluorescent tubes placed above both rooms, and was not directly visible to observers. Mizokami et al. placed in both rooms a number of miniature workaday objects, such as tables and chairs. Their observers viewed a test square placed at several different depths along the line of sight that also passed in depth through the center of both rooms. Using a matching task, Mizokami et al. found that observers substantially discounted the actual illumination profile at different depths in the scene.

How does one reconcile the results of these two experiments? One difference is that Mizokami et al. (1998) constructed real scenes, while Schirillo & Shevell (1993) simulated surfaces on a CRT. However, it is possible to get substantial discounting with virtual, rendered scenes (Boyaci et al., 2003).

A second evident difference is that the walls and additional objects in the scenes designed by Mizokami et al. (1998) were potential cues to the spatial distribution of the illumination. We currently understand the problem of illumination estimation as a cue combination problem (Maloney, 1999, 2002; Yang & Maloney, 2001). Some cues relevant to illumination estimation may be the uniform background of the scene, specular highlights on objects in the scene (Boyaci, Doerschner, & Maloney, in press), the surface specularity of the scene's background, cast shadows, and shading. In this study, we tested this hypothesis in scenes differing in candidate cues to the illumination profile.

Previous theoretical work

Gilchrist et al. (1999) proposed that observers segment complex scenes into illumination frameworks and discount the illumination within each framework. According

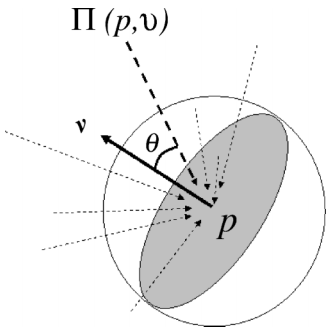

Figure 2. A schematic of the illumination profile. A two-dimensional achromatic Lambertian surface, placed into a three-dimensional scene at a point, $p = (x, y, z)$, with surface normal, ν, absorbs light arriving from all directions within a hemisphere defined by the two-dimensional plane in which the surface lies. By Lambert's Law, the proportion of light absorbed is proportional to $\cos\theta$ where θ is the angle between ν and the direction of the impinging light, ($-90° < \theta < 90°$). The amount of light that the surface absorbs from any direction, υ, is the intensity of light, $\Pi(p, \upsilon)$ arriving at point p, from direction υ weighted by the cosine of the angle, θ, between the surface normal and the direction from which the light arrives. The surface's albedo scales the total light the surface absorbs, $\Lambda(p, \upsilon)$ and the resulting light reradiates uniformly in all directions within the hemisphere.

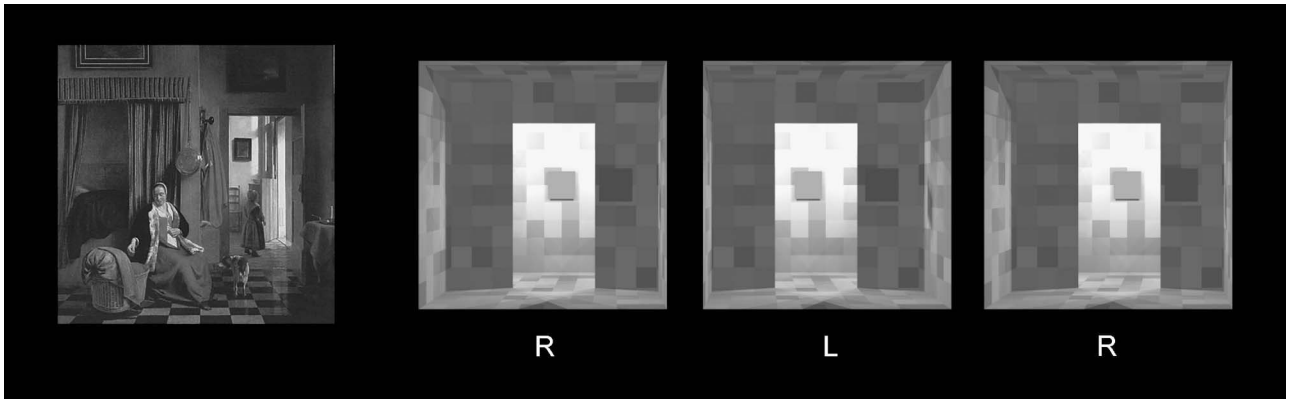

Figure 3. Portrayal of the rendered scene we used in Experiment 1 (bottom images). The reader may perceive the three-dimensionality of the scene by free- or cross-fusing left (L) and right (R) images. In this example, the standard surface is close to the far wall, although we randomized its depth across trials. The adjustable surface, a dark square to the right of the doorway, appeared in this location in all trials. We provide the thumbnail of the de Hooch painting (leftmost image) for comparison.

to Gilchrist, the rules for organizing frameworks and assigning surfaces to them are complex and not fully understood. The three-dimensional structure of a scene could also guide the segmentation of scenes into frameworks, and it is likely that there are analogous effects of three-dimensional organization on lightness perception (e.g., Bloj, Kersten, & Hurlbert, 1999).

In terms of the EIP, the observer may interpret a dramatic change in the illumination profile as a boundary between frameworks. For example, in de Hooch's scene of two rooms (Figure 1), the observer might interpret the doorway as such a boundary.

Preview

In this article, we report the results of three experiments in which observers estimated the EIP[3] in depth along the line of sight in rendered scenes, viewed binocularly. In all experiments, the illumination was akin to that realized in Figure 1—intense in the far room, dim in the near room—and gradients of intensity on the Lambertian walls, floor, and ceiling of the scene, as well as its three-dimensional structure, provided cues to the illumination.

In the first experiment, the only possible illumination cues were the nonspecular ones just mentioned. In a second experiment, we added 11 specular spheres to the scene. The spheres were a potential cue to the actual illumination profile within the scene, and we endeavored to determine whether observers' visual systems made use of them. In a third experiment, we again included 11 spheres in the scene. The spheres of Experiment 3 differed dramatically from those of Experiment 2: the former provided potential cues to illumination that was intense in the near room and dim in the far room, in conflict with the actual illumination profile reflected by the walls, floor, and ceiling. We investigated whether the contradictory cues provided by the spheres affected observers' estimates of the EIP. If an observer's EIP moved toward the conflicting illumination profile signaled by the spheres, we could conclude that the observer's visual system was using the spheres as a cue to the illumination profile.

General methods

Stimuli: Experiment 1

The stimuli were stereo pairs constructed from an achromatic, computer-rendered, virtual, three-dimensional scene. Figure 3 displays an example of a stereo pair we used in Experiment 1. The observer perceived a three-dimensional scene consisting of two, identically sized rooms, one behind the other. A doorway connected the two rooms, allowing one a partial view of the far room. We are indebted to Gilchrist (1977), who studied lightness using real rooms in this arrangement, and Mizokami et al. (1998), who further developed this paradigm.[4] We achieved a three-dimensional percept by rendering the scene twice using two viewpoints corresponding to the locations of the observer's eyes in the experimental apparatus (Figure 8; details to follow). The entire stimulus subtended 8.6 degrees of visual angle (dva).

We rendered the virtual scenes using *Radiance* (Larson & Shakespeare, 1996), an open-source software package that uses 6-point ray tracing to generate models of illuminated scenes. Given a set of surfaces and sources of illumination, *Radiance* computes a three-dimensional model of the scene, which the experimenter may render from any viewpoint. In our scene, all surfaces were Lambertian and all light sources were diffuse. The output of *Radiance* was a specification of normalized luminance at each pixel in the left and right images of the stereo pair.

Ten-by-ten arrays of square, Lambertian tiles covered the walls, ceiling, and floor of the scene. The albedos of the surfaces were drawn randomly and independently from a uniform distribution on the interval [0.25, 1.00]. We used these same randomly selected surfaces on every trial. A surface with an albedo of 1.00 is a perfect reflector: it reradiates 100% of the light incident upon it. Likewise, a surface with an albedo of 0.25 reradiates 25% of the incident light. We designed the illumination in the far room to be more intense than that in the near room. To achieve this effect, we positioned clusters of lights in the far room, while a diffuse light illuminated the entire scene. The clusters hung from the ceiling in the near corners of the far room (Figure 4a) and were concealed from the observer by the walls flanking the doorway. Each light cluster consisted of 11 closely spaced punctuate sources, ensconced in a perfectly transparent, perfectly diffuse material; the result was two extended light sources, which allowed objects in the scene to cast soft shadows onto the walls and floor.

The observer saw a standard surface along the depth axis leading through the center of the rooms and set the apparent albedo of an adjustable surface (shown as a dark square to the right of the doorway in Figure 3) to match the apparent albedo of the standard surface.

The standard surface always was orthogonal to the observer's Cyclopean line of sight, its depth varied from trial to trial, and it ranged from 1.23 to 1.64 dva on a side depending upon depth. In Figure 3, the standard surface is near the back wall. The adjustable surface was also orthogonal to the line of sight. Its depth was constant, and it was always 1.23 dva on a side.

The schematic of Figure 4a shows the relative position of items in the scene (top view) that we rendered to produce the stereo image pair. We marked with gray rectangles the seven possible locations in which the standard surface could appear. Arrows refer to the positions of the hidden light sources in the far room, the adjustable surface, and the doorway. We intend the gradient of Figure 4a to remind the reader that illumination was most intense at the far wall and decreased in the direction of the observer.

Prior to running experiments, we verified the rendered illumination profile. We replaced the standard surface with a small matte reference surface of constant albedo and rendered it within the scene. (The observer never saw the reference surface during the experiment.) We recorded the relative luminance value assigned by *Radiance* to the reference surface. We repeated this procedure with the reference surface placed at 19, equally spaced depths along the line of sight. We denote these relative luminance values by $I(d)$ for each depth d. Figure 4b plots these values,[5] normalized so that the value at the far wall is 1, as a function of relative depth. We defined the depth at the scene's edge nearest to the observer as zero and the depth at the far wall as 10. We refer to the function of illumination over depth (the solid black curve in Figure 4b) as the relative illumination profile along the

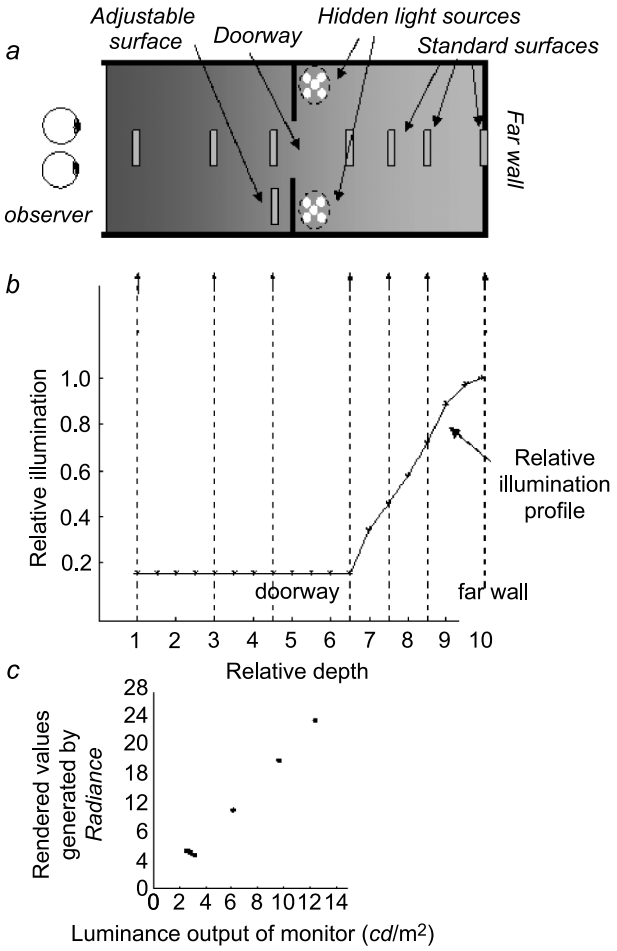

Figures 4. (a–c). The illumination profile in the scene. Panel a is a schematic of the stimuli of Experiment 1 (top view). Panel b plots the relative illumination profile, $I(d)$, of Experiment 1 (solid curve); dashed arrows label the relative depths at which standard surfaces appeared in all experiments. Panel c demonstrates verification of the rendered illumination profile (for the right image of the stereo pair): for reference surfaces at seven depths, we plotted the relative luminance values assigned by *Radiance* versus the physical luminances measured with a spectrophotometer. The points fall along a straight line through the origin, indicating that the luminance values emitted by the monitor were a fixed scaling of those assigned by *Radiance*. The results for the left image of the stereo pair were equivalent.

line of sight. Starting at the far wall, the relative illumination profile, $I(d)$, drops as we approach the observer by roughly a factor of five.[6]

We plotted the relative illumination profile of Figure 4b below the cartoon of Figure 4a in registration with it. Dashed arrows re-emphasize the seven possible relative depths at which the standard surface appeared $(d = 1, 3, 4.5, 6.5, 7.5, 8.5, 10)$ in all experiments.

To verify that the experimental hardware and software were displaying the illumination profile as desired, we measured the physical luminances (in cd/m^2) of the reference surface at each depth using a Photo Research

Figure 5. Portrayal of the rendered scene we used in Experiment 2. The reader may perceive the three-dimensionality of the scene by free- or cross-fusing left (L) and right (R) images. As in Figure 3, the standard surface is close to the far wall, although we randomized its depth across trials. The adjustable surface, a dark square to the right of the doorway, appeared in this location in all trials.

PR-650 SpectraScan spectrophotometer. We placed the spectrophotometer at the location of the observer's right eye and measured the light that would reach the eye from the center of the reference surface within the scene on the right monitor (Figure 4c). (Recall that the observer never saw the reference surfaces during the experiment.) We thrice repeated this measurement for each of seven different depths (marked in Figures 4a–b) assumed by the reference surface. We then repeated these measurements for the left eye. In Figure 4c, we plotted the measurements of the seven reference surfaces in the right eye against the relative luminance values assigned to them by *Radiance*. The points fall along a straight line through the origin, indicating that the luminance values of regions on the screen were a fixed scaling of those assigned by *Radiance*. The results for the left eye were equivalent.

Procedure

The luminance of the standard surface was constant (93.0 cd/m^2)[7] across trials, and we were interested in whether this constant-luminance surface took on different apparent albedos at different depths. We asked observers to set the lightness of the adjustable surface to match that of the standard surface presented at one of seven depths along the line of sight passing through the center of both rooms. The seven locations at which the standard surface could appear (marked locations in Figures 4a–b) were a subset of the 19 locations where we measured the luminance profile using reference surfaces (previous section). We randomized the depth at which the standard surface appeared from trial to trial.

Stimuli: Experiment 2

The stimuli of Experiment 2 were identical to those of Experiment 1 except that we added a set of 11 small, identically-sized spheres to each scene. Figure 5 displays an example stereo pair of the scene we used in Experiment 2. The spheres were of 100% specularity and perfectly reflected the ambient illumination. We chose the spheres' locations by hand, with the aim of providing the observer with a plethora of cues to the scene's illumination. We placed no sphere along the observer's line of sight to either the standard or adjustable surface. The spheres' locations were constant across trials. Each sphere subtended from 0.29 to 0.41 dva, depending upon its depth.

Stimuli: Experiment 3

The stimuli of Experiment 3 were identical to those of Experiment 2 except that the spheres reflected an illumination different from that of the scene's ambient illumination (Figure 6). The relative illumination profile reflected by the Lambertian surfaces was consistent with that of Experiments 1 and 2 (Figure 4b), in which illumination decreased from the far wall in the direction of the observer, while the specular spheres reflected a relative illumination profile that increased from the far wall in the direction of the observer. Figure 7 plots the relative illumination profile reflected by the scene's Lambertian surfaces ($I_n(d)$, solid curve) together with the relative illumination profile reflected by the specular spheres ($I_s(d)$, dotted curve).

To generate this impossible scene, we first rendered the entire scene, including the spheres, under the illumination profile we used in Experiments 1 and 2 ($I_n(d)$, solid curve). Next, we again rendered the entire scene, this time under the counterpole illumination profile, $I_s(d)$ (dotted curve). Finally, we removed the pixels comprising the spheres from the former image and replaced them with the spheres from the latter image. The resulting scene contained specular cues to $I_s(d)$ and nonspecular cues to $I_n(d)$.

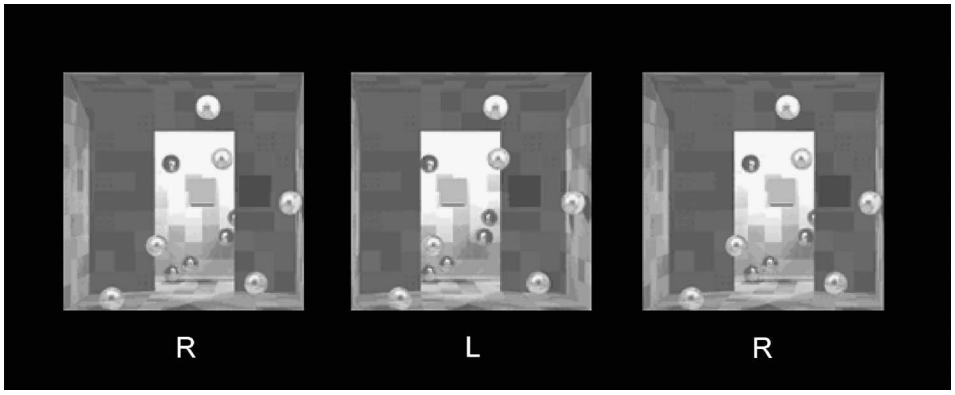

Figure 6. Portrayal of the rendered scene we used in Experiment 3. The reader may perceive the three-dimensionality of the scene by free- or cross-fusing left (L) and right (R) images. As in Figures 3 and 5, the standard surface is close to the far wall, although we randomized its depth across trials. The adjustable surface, a dark square to the right of the doorway, appeared in this location in all trials.

Apparatus

We employed a Wheatstone (1838) stereoscope as our experimental apparatus (Figure 8). Directly in front of the observer's eyes sat two half-silvered mirrors, which reflected the images displayed on the left and right monitors to the corresponding eyes of the observer. The stereoscope sat within a cubic box of dimensions 124 cm per side. One face of the box was absent, allowing the observer to look into the box. The observer sat in a chair and rested her chin in a chinrest (the HeadSpot, handmade by UHCOTech, 2004). The chinrest served to restrict the observer's head movements while permitting her eyes to move freely.

The observer viewed displays in a darkened room. To absorb stray light, black flocked paper (Edmund Scientific, 2004) coated the interior of the box. We placed additional light baffles near the observer's temples to prevent light from the monitor screens reaching the observer's eyes directly; the observer saw only the displays reflected by the mirrors.

The optical distance from the observer's eyes to the monitors was 70 cm, and the binocularly fused stimuli appeared an average of 70 cm in front of the observer, thereby minimizing possible conflicts between disparity and accommodation depth cues. The monocular fields of view were approximately 40 × 40 dva. Observers' eyes were at approximately the same height as the center of the monitor screens and the scenes being viewed.

Software

Boyaci et al. (2003) wrote the original experimental software in C, which was subsequently modified by JLS. This software controlled stimulus presentation and data collection. We used the X Window System, Version 11R6, running under Red Hat Linux 6.1 for graphical display of the stimuli.

Hardware

A Dell 410 Workstation with a Matrox G450 dual head graphics card and a special purpose graphics driver from Xi Graphics controlled both monitors. Two 21-in. Sony Trinitron Multiscan GDM-F500 monitors presented left and right images, respectively, to each eye of the observer. We positioned the monitors to the observer's left and right, as in Figure 8. The monitors' screens were planar within 1 mm of deviation.

With a Photo Research PR-650 SpectraScan spectrophotometer, we made direct measurements of the luminance values of each monitor and used these measurements to prepare lookup tables for each gun of each monitor. We used the lookup tables to correct the nonlinearities in the

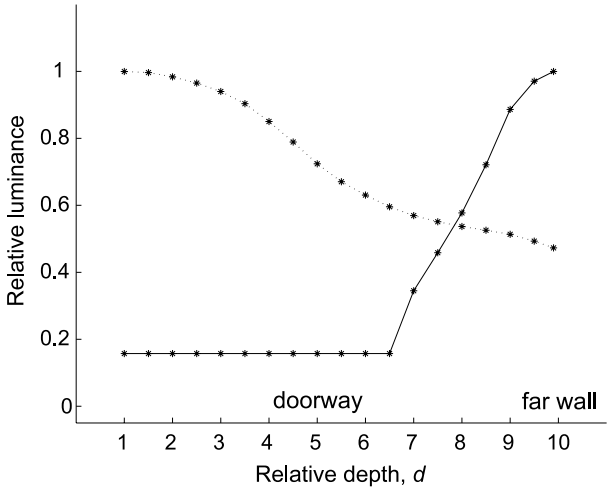

Figure 7. The relative illumination profile reflected by the Lambertian (nonspecular) surfaces, $I_N(d)$ (solid curve), together with the relative illumination profile reflected by the specular spheres, $I_S(d)$ (dotted curve), in Experiment 3. The *x*-axis is relative depth; the *y*-axis, relative luminance. The doorway is at relative depth = 5; the far wall resides at relative depth = 10.

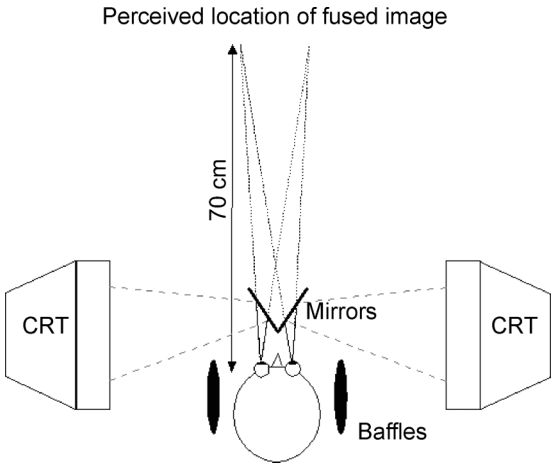

Figure 8. A schematic of our Wheatstone stereoscope. Left and right displays appeared on two monitors to the left and right of the observer. Half-silvered mirrors at 45° to the observer's line of sight reflected these displays to the observer's eyes. The observer perceived the fused display along the line of sight at a distance of approximately 70 cm, which was also the optical distance to either monitor.

gun responses and equalize the display values on the two monitors. The maximum luminance achievable on either monitor was 114 cd/m^2.

Task

We instructed observers to adjust the adjustable surface to match the standard surface in albedo, until they were convinced that the "two surfaces were cut from the same piece of paper." Observers performed this task by using the left and right buttons of a computer mouse to make the adjustable surface appear lighter and darker. In setting the adjustable surface, the observer chose 1 of 64 different values, spanning a wide range of luminance at narrow intervals. On each trial, we randomized the initial setting of the adjustable surface. In agreement with other studies (Brainard, 1998), we found that the choice of starting adjustable surface had no effect on observers' final settings.

Before beginning the experiments, we tested that observers could set veridical luminance matches in a scene containing only a standard and an adjustable surface. We then allowed observers to practice the actual task on a demonstration version of the program until they felt comfortable with the task (four trials, on average). During the experiment, observers were free to take breaks from the task whenever they wished, although none exercised this option. At each of the seven depth levels of the standard surface, the observer performed 10 trials. Although there was no time constraint, all observers completed the 70 trials in approximately 25 minutes. Observers received no feedback during the experiment.

Observers

Our five observers were graduate and postdoctoral students at the Psychology Department and Center for Neural Science of New York University. Each observer had normal or normal-to-corrected visual acuity and could binocularly fuse the stimuli with ease. All five observers participated in all experiments. All but one observer (JLS, an author) were unaware of the purpose of the experiments. We paid the naïve observers $12 per hour for their participation. Prior to selecting observers for the study, we allowed candidate observers to practice the task. We excluded from the study candidate observers who said they did not understand the instructions and therefore could not perform the task.

Analysis

Computation of normalized EIP

At each of seven depths, each observer performed 10 asymmetric matches by setting the apparent albedo of the adjustable surface to match that of the standard surface. We denote these seven luminance values by $L_1, L_2, L_3, \ldots L_7$, ordering the indices in increasing relative depth as indicated in Figure 4b. We denote the fixed luminance of the standard surface L_n, and we are interested in the following ratios:

$$\Lambda_i = \frac{L_0}{L_i}. \quad (5)$$

We can rewrite Equation 5 in terms of the apparent albedos of the surfaces by using the identity, $L_i = \hat{E}_i\hat{\alpha}$, where $\hat{\alpha}$ is perceived albedo, and $\hat{E}$ is the estimated EIP that the observer effectively discounts in arriving at an estimate of albedo. Substituting this identity into Equation 5, we have

$$\Lambda_i = \frac{\hat{E}_0\alpha}{\hat{E}_i\hat{\alpha}} = \frac{\hat{E}_0}{\hat{E}_i}, \quad (6)$$

where the second equality follows from two assumptions. The first assumption is that the observer is matching apparent albedos as instructed. The second is that the estimated EIP, $\hat{E}_i$, that the observer discounts from the adjustable surface in estimating its albedo is always the

same because the adjustable surface's position is fixed and the illumination falling on it is constant.

If we now normalize these ratios by dividing by Λ_7 (the value corresponding to the far wall), we have

$$\lambda_i = \frac{\Lambda_i}{\Lambda_7} = \frac{\hat{E}_0/\hat{E}_i}{\hat{E}_0/\hat{E}_7} = \frac{\hat{E}_7}{\hat{E}_i}, \quad (7)$$

the ratio of the EIP at the far wall to the EIP at each depth i.

We can interpret λ_i via comparison with the relative illumination profile. Suppose, on the one hand, that the observer has perfect lightness constancy. Then the EIP at each depth would be proportional to the relative illumination profile values in Figure 4b, and the values of λ_i, for $i = 1, 2, 3, \ldots 7$, plotted against depth, would fall exactly on the relative illumination profile, $I(d)$. Note that the normalized EIP, by definition, is always equal to 1 at the far wall of the scene.

Suppose, on the other hand, that these ratios, λ_i, were the same for every depth i. This outcome would indicate that the observer is simply discounting the same EIP at every depth. We therefore would conclude that the observer does not discount an illumination profile that varies in depth. In other words, while this observer is expert at matching luminance values, she does not use the information about depth provided in the scene. These perfect and no lightness constancy models define two extremes of a continuum of possible models. A third and likeliest possibility is that observers exhibit some amount of lightness constancy. Figure 9 displays a discrete sample of possible models. The dashed line at the y-axis value of 1 is the no lightness constancy model. The solid curve is the true relative illumination profile or perfect lightness constancy model. β is the scaling parameter, and the dotted curves are copies of the perfect lightness constancy model, scaled by β values ranging from 0 to 1. We note that this analysis is similar in spirit to that of Brunswik (1929). (See also Gilchrist, 2006, and Thouless, 1931.). We are approximating the observer's performance as a weighted mixture of a "no constancy" profile and a "perfect constancy" profile. One can interpret a β value as one would interpret a Brunswik ratio.

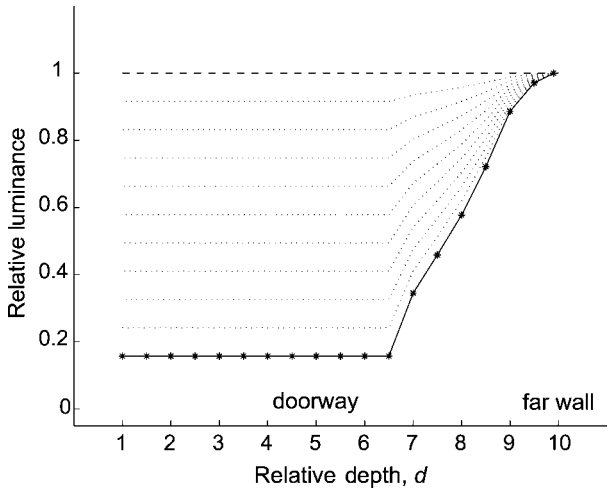

Figure 9. The continuum of lightness constancy models, from the no constancy model (dashed line) to the perfect constancy model, $I(d)$ (solid curve). The dotted curves are scaled copies of the perfect constancy model. β is the scaling parameter, which ranges from 0 (no constancy) to 1 (perfect constancy).

Results

Experiment 1

Figures 10a–e show the results of Experiment 1 for the five observers we tested. The error bars for each normalized EIP measurement correspond to ±2 standard deviations from the mean. The solid curve is the relative illumination profile of Figure 4b. The horizontal, dashed line at $y = 1$ represents the pattern expected if the observer fails completely to discount the illumination profile.

Examining the shape of the normalized EIP (open circles) for each observer in Figures 10a–e, we see that depth clearly affected observers' judgments. The effect of depth was significant, as verified by a repeated-measures ANOVA ($F = 14.610$, $p < .01$). We then performed two-tailed t tests for each observer's data for depth conditions 1–6[8] versus the null hypothesis that the observer's true normalized EIP falls on the line $y = 1$ (no discounting of the illumination profile in depth). We found that two observers' (JLS and PXV) normalized EIP settings were significantly different (with 99% confidence intervals) from the no discounting profile for all six depths tested. For the remaining three observers (RDP, SMT, and VWC), the same five of six depths tested were significantly different (with 99% confidence intervals) from the no constancy model. For these three observers, only the depth closest to the far wall was not significantly different from the no constancy model, and we attribute this result to a ceiling effect.

We next performed two-tailed t tests for each observer at each depth, testing whether the observer's settings were significantly different from the relative illumination profile, $I(d)$ (perfect lightness constancy). For all five observers, we found that all normalized EIP settings were significantly different (with 99% confidence intervals) from $I(d)$ for all six depths tested.

We conclude that depth affected observers' settings at most if not all relative depths values. All five observers shared a common pattern of normalized EIP versus relative depth. This pattern corresponded to the actual

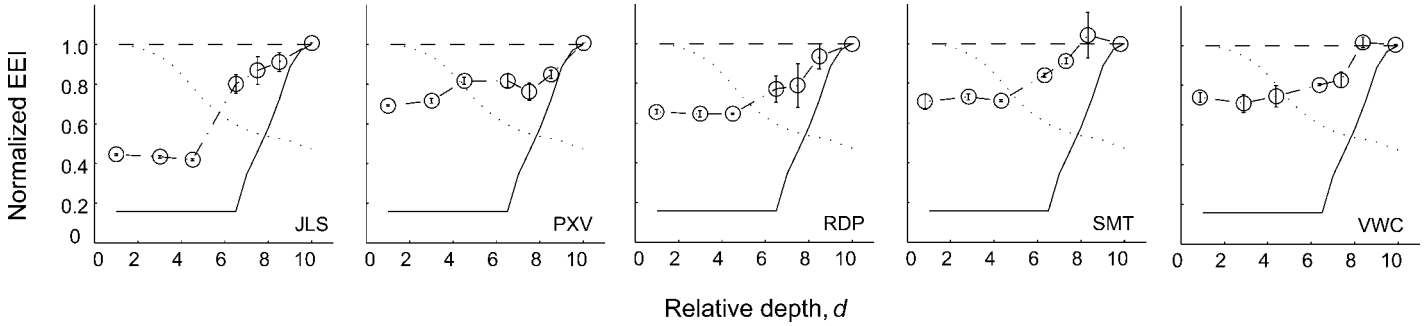

Figures 10. (a–e). Data for Experiment 1 for five observers. The solid curves are the relative illumination profile, $I(d)$, or the perfect lightness constancy model. The dashed lines are $y = 1$ or the no lightness constancy model. The open circles are normalized EIP values, and the error bars are ± 2 standard deviations.

profile of illumination in the scene, and there was a statistically significant trend for the normalized EIP to be lower in the near room. In particular, observers judged the near room to be 44–73% as illuminated as the far room.

Experiment 2

In Experiment 1, observers demonstrated some amount of illumination discounting, but their performance was far from ideal. Gradients of intensity on the Lambertian walls, floor, and ceiling conveyed limited cues about the distribution of light in the scene, which was more akin to Schirillo & Shevell's (1993) stimuli than to those of Mizokami et al. (1998). We conjectured that if observers truly were estimating and discounting the illumination in the three-dimensional scene, then giving them more information about the light—its origin, direction, and intensity—should make their estimates more veridical.

In Experiment 2, we repeated the first experiment with additional cues to the three-dimensional illumination distribution. Perfectly specular spheres provided these new cues. Yang & Maloney (2001) demonstrated that specularity provides observers with increased information about the chromaticity of the illumination. We hypothesized that adding specularities to our three-dimensional scene would improve observers' illumination estimation. We endeavored to know whether the addition of specular cues to nonspecular cues would make observers' normalized EIPs more veridical.

We analyzed the data as in Experiment 1. In Figures 11a–e, we plotted the results of Experiment 2 superimposed upon those of Experiment 1. Once again, we see that depth clearly affected observers' judgments, and all five observers discounted the illumination in the scene according to a similar profile. Visually comparing the results of Experiment 2 with those of Experiment 1, we found that the normalized EIPs of the former more closely follow the relative illumination profile, $I(d)$. In particular, observers judged the near room to be 21–51% as illuminated as the far room, whereas they judged it to be 44–73% as illuminated in Experiment 1.

Again, there was a significant effect across depth, as shown by a two-way repeated measures ANOVA ($F = 12.546$, $p < .01$). We performed two-tailed t tests for each observer's data for depth conditions 1–6, versus the null hypothesis that the observer's mean setting is not different from 1. We rejected the null hypothesis (with 99% confidence intervals) for all six depths (again, we excluded the fixed point at the far wall) for all five observers.

We next performed two-tailed t tests for each observer at each depth, testing whether the observer's settings were significantly different from the relative illumination profile, $I(d)$ (perfect lightness constancy). For three of five observers (PXV, SMT, and VWC), we found that all but one normalized EIP setting were significantly different (with 99% confidence intervals) from $I(d)$ for all six depths tested. For the remaining two observers (JLS and RDP), we found that all but two normalized EIP settings were significantly different (with 99% confidence intervals) from $I(d)$ for all six depths tested.

We conclude that depth affected observers' settings at all depths and that the resulting normalized EIP profiles followed the shape of the relative illumination profile, $I(d)$. Again, all five observers shared a common pattern of normalized EIP versus relative depth. Furthermore, the normalized EIPs of Experiment 2 more closely adhere to $I(d)$ than do those of Experiment 1. We quantify this improvement in the next section.

Comparison of Experiments 1 and 2

Figures 11a–e superimpose the results of Experiment 2 onto those of Experiment 1 for all five of the observers we tested. The open circles show Experiment 1 data; the black, filled circles, Experiment 2 data. The error bars for each normalized EIP measurement correspond to ±2 standard deviations from the mean. The solid curve is the illumination profile, $I(d)$, of Figure 4b. The horizontal, dashed line at $y = 1$ represents the pattern expected if the observer fails completely to discount the actual illumination profile. A visual analysis reveals that all

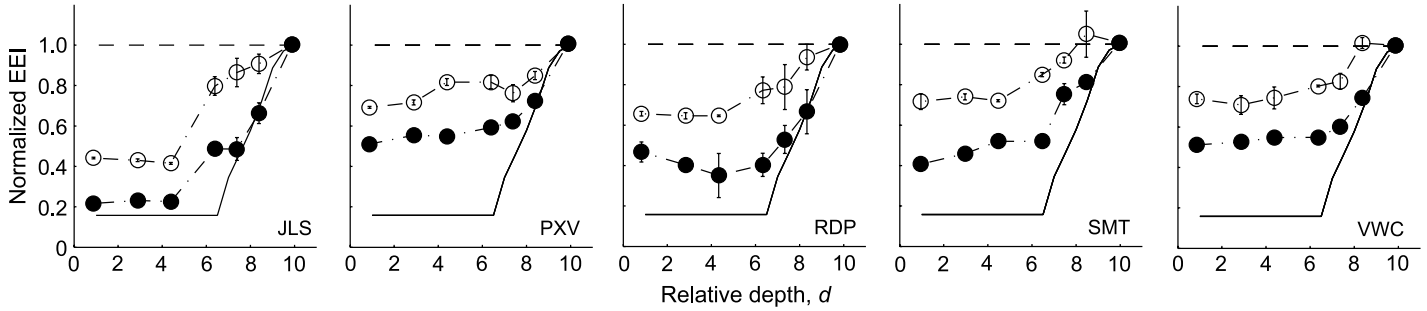

Figures 11. (a–e). Data for Experiments 1 and 2 for five observers. The solid curves are the relative illumination profile, $I(d)$, or the perfect lightness constancy model. The dashed lines are $y = 1$ or the no lightness constancy model. The open circles are normalized EIP values for Experiment 1; the black circles, the same for Experiment 2. Error bars are ±2 standard deviation.

normalized EIPs for all observers are more veridical in Experiment 2 than in Experiment 1.

We wanted to test the validity of our assumption that observers are using illumination estimates to make lightness matches. Moreover, we wanted to demonstrate numerically that observers were better able to estimate the illumination in Experiment 2 than in Experiment 1. We regressed each observer's normalized EIP profile against $I(d)$. We fit only one parameter, β, in the regression, because all profiles must have a common y-intercept equal to 1 at the far wall. We used a bootstrap estimation technique to estimate $\hat{\beta}$ and its standard error. We expected $\hat{\beta}_{exp2}$ values to exceed $\hat{\beta}_{exp1}$ values.

We also computed the R^2 statistic to summarize the variability in the data. An R^2 equal to 1 indicates that modeling the data as a scaled copy of the illumination profile can account for all of the variance in the data, whereas an R^2 equal to 0 indicates that this model can account for none of the variance in the data.

We report $\hat{\beta}$ (Figure 12) and R^2 (Table 1) values for five observers for Experiments 1 and 2. Figure 12 plots $\hat{\beta}_{exp2}$ versus $\hat{\beta}_{exp1}$ for five observers. Horizontal and vertical error bars indicating 99% confidence intervals are smaller than the black circle symbols. A line, $y = x$, aids our comparison of these estimates: points on this line evidence that $\hat{\beta}_{exp2} = \hat{\beta}_{exp1}$, points above the line evidence that $\hat{\beta}_{exp2} > \hat{\beta}_{exp1}$, and points below the line evidence that $\hat{\beta}_{exp2} < \hat{\beta}_{exp1}$. We see that the data heavily favor the conclusion that $\hat{\beta}_{exp2} > \hat{\beta}_{exp1}$ for all observers. Because $\hat{\beta}_{exp2} > \hat{\beta}_{exp1}$, we conclude that all five observers' normalized EIPs followed $I(d)$ more closely in Experiment 2 than in Experiment 1 (recall Figure 9). In addition, $R^2_{exp2} > R^2_{exp1}$ (Table 1), evidencing that observers' illumination estimates were more consistent (less variable) in Experiment 2 than Experiment 1.

Experiment 3

In the scene of Experiment 3, nonspecular cues provided information about one illumination profile ($I_n(d)$, the solid curve of Figure 7), while specular spheres reflected cues to an entirely different illumination ($I_s(d)$, the dotted curve of Figure 7). We rendered this "impossible" scene to examine the relative contributions of specular and nonspecular cues to the illumination.

If observers were combining information from both types of cue, we expected that their EIPs would resemble a mixture of the two illumination profiles in Figure 7. We wondered whether the specular or nonspecular cues would influence observers' judgments more heavily.

Figures 13a–e show Experiment 3 data for all five observers. The error bars for each normalized EIP measurement correspond to ±2 standard deviations from the mean. The solid curve is $I_n(d)$, which was reflected to the observer by the Lambertian surfaces in the scene. The dotted line is $I_s(d)$, which was reflected by the specular spheres. Instead of normalizing judgments by the observer's setting at the far wall as we did in Experiments 1 and 2, we chose as our normalization value the depth that each observer judged to have the most intense illumination. This technique insured that no normalized EIP exceeded a value of 1.

As we expected, the addition of perturbed, specular spheres to the scene resulted in illumination estimates which did not follow $I_n(d)$, the illumination profile conveyed by the nonspecular cues. The data suggest that

	R^2_{exp1}		R^2_{exp2}
JLS	.71	→	.88
PXV	.60	→	.95
RDP	.90	→	.91
SMT	.76	→	.95
VWC	.76	→	.95

Table 1. Comparison of the R^2 statistic to summarize the variability in the data of Experiments 1 and 2 for five observers. An R^2 equal to 1 indicates that modeling the data as a scaled copy of the illumination profile can account for all of the variance in the data, whereas an R^2 equal to 0 indicates that this model can account for none of the variance in the data. R^2_{exp2} exceeded R^2_{exp1} for all observers.

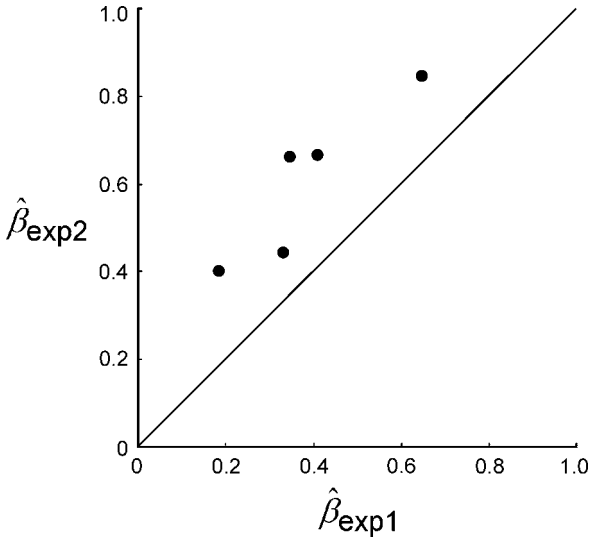

Figure 12. Comparison of $\hat{\beta}_{exp2}$ versus $\hat{\beta}_{exp1}$ for five observers. Horizontal and vertical error bars indicating 99% confidence intervals are smaller than the black circle symbols. The line $y = x$ aids our comparison of these estimates: points on this line evidence that $\hat{\beta}_{exp2} = \hat{\beta}_{exp1}$, points above the line evidence that $\hat{\beta}_{exp2} > \hat{\beta}_{exp1}$, and points below the line would evidence that $\hat{\beta}_{exp2} < \hat{\beta}_{exp1}$.

the cues to illumination provided by the specular spheres influenced observers' judgments, at least in the near room (JLS, RDP, and SMT), if not at all depths (VWC). In the next section, we attempted to determine the relative weight given to the specular or the nonspecular cues in the scene by each observer.

Cue combination: Specular and nonspecular cues

Because Experiments 1 and 2 provided the observer with cues to a single illumination gradient, $I(d) = I_n(d)$, we interpreted the data of Experiments 1 and 2 by fitting it with a univariate index of lightness constancy, β (recall Figure 9), which ranged from 0 (no constancy) to 1 (perfect constancy). In contrast, Experiment 3 presented the observer with independent cues to two, conflicting illumination gradients, $I_n(d)$ and $I_s(d)$. Therefore, it was inappropriate to fit the data of Experiment 3 with a single illumination profile. Instead, we compared the results of Experiment 3 to a weighted linear cue combination model of the form

$$I(d) = \beta_n I_n(d) + \beta_s I_s(d) \quad (8)$$

and chose the weights β_n and β_s to optimize the fit of each observers' data. Equation 8 defines a broad class of models, or possible illumination profiles, that are weighted mixtures of the relative illumination profiles $I_n(d)$ and $I_s(d)$. The model's first term is the weight of the nonspecular cues, β_n, multiplied with the illumination profile conveyed by the nonspecular cues, $I_n(d)$ (solid curve in Figure 7). The second term is the weight of the specular cues, β_s, multiplied with the illumination profile conveyed by the specular cues, $I_s(d)$ (dotted curve in Figure 7).

We estimated the weights β_n and β_s by performing a multiple regression of each observer's data, λ_i, against this model. In other words, we fit the equation

$$\lambda_i = \hat{\beta}_n I_n(d_i) + \hat{\beta}_s I_s(d_i) + \varepsilon \quad (9)$$

by finding maximum likelihood estimates $\hat{\beta}_n$ and $\hat{\beta}_s$. We assumed the error term, ε, to be Gaussian with mean zero.

Figure 14 plots $\hat{\beta}_n$ and $\hat{\beta}_s$ for five observers. Our results lead us to conjecture that each observer used a unique combination of specular and nonspecular cues to the illumination to do the task. In particular, three observers (JLS, PXV, and VWC) seem to have used primarily the specular cues, while two other observers (RDP, SMT) seem to have used the cues more or less equally.

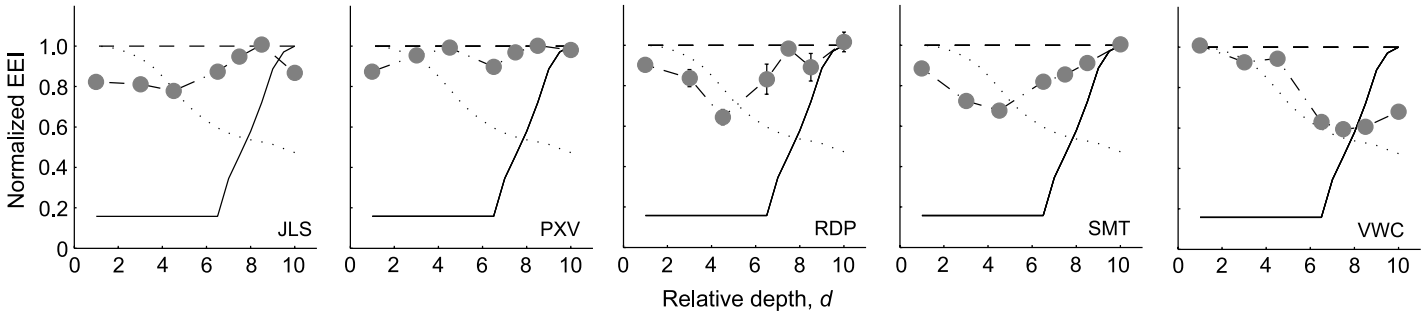

Figures 13. (a–e). Data for Experiment 3 for five observers. The solid curves are $I_n(d)$, the relative illumination profile in the scene signaled by the nonspecular cues, while the dotted lines are $I_s(d)$, the relative illumination profile signaled by the specular cues. The gray circles are normalized EIP values for Experiment 3, and the error bars are ±2 standard deviations.

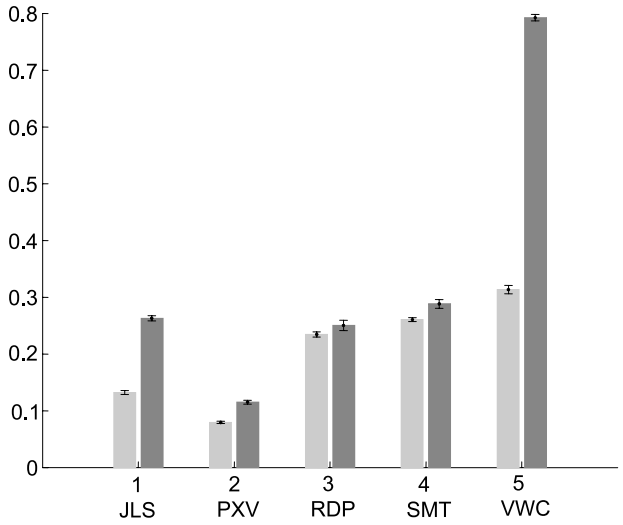

Figure 14. Plots of $\hat{\beta}_n$, the weights of the nonspecular cues, and $\hat{\beta}_s$, the weights of the specular cues, for Experiment 3 for five observers. The lighter gray bars indicate $\hat{\beta}_n$, and the darker gray bars indicate $\hat{\beta}_s$. Error bars indicate 99% confidence intervals.

We next computed R^2_{exp3} for the cue combination model for all observers and found values of .49, .33, .44, .65, and .89 for observers JLS, PXV, RDP, SMT, and VWC, respectively. While the cue combination model can partially account for observers' judgments, we note that R^2_{exp3} is less than R^2_{exp1} (and therefore also less than R^2_{exp2}) for four of five observers (all except VWC, for whom $R^2_{exp1} < R^2_{exp3} < R^2_{exp2}$). Clearly, interpreting the data as a combination of conflicting cues had a cost. A possible alternative explanation is that observers used different cues to estimate the illumination profile in the near and far rooms, respectively.

Summary

Within our rendered scenes, the illumination changed with depth. Observers were able to partially discount the illumination in each of the three scenes we used in this study. We demonstrated that multiple cues in the scene mediated observers' illumination estimation. Specifically, observers used an illumination cue based on specularity whether or not this cue signaled an illumination profile consistent with that signaled by other, nonspecular cues.

In Experiment 1, gradients on Lambertian surfaces within the scene (Figure 3) conveyed all of the available information about the illumination. Using the available information, observers exhibited a fair amount of lightness constancy: 19–65%, when we modeled observers' normalized EIPs as scaled copies ($\hat{\beta}_{exp1}$) of the relative illumination profile, or, equivalently, Thouless ratios. (For an introduction to Thouless ratios, please see Gilchrist, 2006. For a detailed explanation, please see Brunswik, 1929, and Thouless, 1931.)

In Experiment 2, the perfectly specular spheres (Figure 5) provided observers with a sample of the intensity, origin, and direction of the light sources over a large part of the scene. We found that when we combined specular cues with nonspecular cues, and all cues provided consistent information, observers' estimates of the illumination were more veridical ($\hat{\beta}_{exp2} = .41$ to .85, Thouless ratios of 41–85%) than when nonspecular cues alone were present (Figure 12). In addition, the variability of the normalized EIP decreased with the addition of specular spheres to the scene (Table 1).

When both specular and nonspecular cues were present and provided conflicting information (Experiment 3), we found that a weighted linear cue combination could partially explain observers' data. This finding suggests an interesting area for further research.

We conclude that when there are adequate cues to the spatial distribution of the illumination, the human visual system can estimate the illumination profile at arbitrary points in three-dimensional space.

Discussion

As we anticipated, our results conflict with those of Schirillo & Shevell (1993) and confirm those of Mizokami et al. (1998). We conjecture that our scenes, like those of Mizokami et al., contained more cues to the spatial profile of illumination than did the more impoverished scenes of Schirillo and Shevell. In fact, if one were to add consistent cues (illumination gradients, specular spheres) to the scenes of Schirillo and Shevell, we would expect that observers' lightness constancy would increase on the lightness matching task.

Earlier we raised the question of whether the results of Mizokami et al. (1998), who used real rooms, are applicable to studies (such as this one) that use CRTs to render three-dimensional scenes. Our results suggest a positive answer. We also raised the question of whether a full palette of cues to illumination, commensurate with those in natural scenes, must be available for observers to exhibit lightness constancy. Our results suggest a positive answer to this question too.

Our experiments showed that the more cues to the illumination we provided to observers, the greater the degree of lightness constancy they exhibited. We demonstrated empirically through a series of controlled experiments what de Hooch knew as an artist: the painter evokes a natural scene by providing the viewer with cues to the illumination. In closing, we ask the reader to consider with renewed appreciation the choices of de Hooch (1659–1660). The artist communicated *The Mother's* subtle, complex illumination by concealing the light sources and rendering their reflection from surfaces shiny,

like the checkerboard floor tiles, to matte, like the mother's skirts. Arranged in pictorial perspective, these surfaces accurately communicate the illumination profile of the scene to the observer.

Acknowledgments

This research was supported by NIH/NEI Grant EY08266 (LTM), a Human Frontiers Science Program Grant RG0109/1999-B (LTM), a NASA Graduate Student Researcher Fellowship (JLS), the RIACS Summer Student Research Program (JLS), and the New York University Graduate School of Arts and Science Dean's Dissertation Fellowship (JLS). It comprised part of the dissertation research of JLS. We thank Hüseyin Boyaci for help in developing the software used in the experiments described here. We are grateful to Hüseyin Boyaci, Marisa Carrasco, Alan Gilchrist, Brian McElree, Misha Pavel, and Qasim Zaidi for many helpful comments and suggestions.

Commercial relationships: none.
Corresponding author: Jacqueline Leigh Snyder.
Email: jacqueline.snyder@nyu.edu.
Address: New York University, Psychology Department, 6 Washington Place, Room 877-C, New York, NY 10003.

Footnotes

[1]Land & McCann (1971) coined this term, eponymous to Piet Mondrian, the early twentieth-century Dutch painter who experimented with compositions of the most basic elements: line and color.

[2]We duplicated the lightness condition of the first experiment of Schirillo & Shevell (1993) as a pilot study in designing the experiments reported here. Four naïve observers made lightness matches to stimuli of constant luminance presented at different depths. We found no effect of depth on perceived lightness of these constant luminance stimuli, substantially in agreement with the results of Schirillo and Shevell for naïve observers. We did not ask observers to estimate brightness.

[3]This terminology is consistent with that of Boyaci et al. (in press).

[4]Schirillo, Reeves, & Arend (1990) used rendered binocular scenes with a similar spatial layout.

[5]Note that the relative illumination profile (Figure 4b) has its highest value at the far wall. To those unfamiliar with the properties of Lambertian surfaces, the reason for this phenomenon may be non-intuitive. One may ask why the relative illumination profile does not achieve its highest value in the vicinity of the hidden light sources, and then drop off in the direction of the far wall. The answer lies in the fact that it is not the distance to the light source that parameterizes the relative illumination profile but instead the cosine of the angle between the vector, $\boldsymbol{v}$, of the incident light and the vector, $\boldsymbol{\nu}$, of the surface normal. At the far wall, the value of the cosine is maximized and therefore so is the value of the relative illumination profile.

[6]We designed the luminance profile to be as comparable to the 5:1 luminance ratio of Schirillo & Shevell (1993) as possible.

[7]We repeated the methods of the previous section to measure and verify these values.

[8]We excluded the point at the far wall, which we defined to equal 1.

References

Adelson, E. H., & Bergen, J. R. (1991). The plenoptic function and the elements of early vision. In M. S. Landy & J. A. Movshon (Eds.), *Computational models of early vision* (pp. 3–22). Cambridge, MA: MIT Press.

Arend, L. (1994). Surface colors, illumination, and surface geometry: Intrinsic-image models of human color perception. In A. Gilchrist (Ed.), *Lightness, brightness, and transparency* (pp. 159–213). Hillsdale, NJ: Erlbaum.

Basri, R., & Jacobs, D. W. (2003). Lambertian reflectance and linear subspaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 25*(2), 218–233.

Bloj, M. G., Kersten, D., & Hurlbert, A. C. (1999). Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature, 402*(6764), 877–879. [PubMed]

Boyaci, H., Doerschner, K., & Maloney, L. T. (2004). Perceived surface color in binocularly–viewed scenes with two light sources differing in chromaticity. *Journal of Vision, 4*(9), 664–679. http://journalofvision.org/4/9/1/, doi:10.1167/4.9.1. [PubMed] [Article]

Boyaci, H., Doerschner, K., & Maloney, L. T. (in press). Cues to an equivalent lighting model. *Journal of Vision*.

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*(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: Optics and Image Science, 15*(2), 307–325. [PubMed] [Article]

Brunswik, E. (1929). Zur entwicklung der albedowahrnehmung (On the development of the perception of albedo). *Zeitschrift für Psychologie, 109,* 40–115.

Gilchrist, A. (1994). Surface colors, illumination, and surface geometry: Intrinsic-image models of human color perception. In A. Gilchrist (Ed.), *Lightness, brightness, and transparency* (pp. 1–31). Hillsdale, NJ: Erlbaum.

Gilchrist, A. L. (1977). Perceived lightness depends on perceived spatial arrangement. *Science, 195*(4274), 185–187. [PubMed]

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

Gilchrist, A. L. (2006). *Seeing black and white.* New York: Oxford.

Gilchrist, A. L., Delman, S., & Jacobsen, A. (1983). The classification and integration of edges as critical to the perception of reflectance and illumination. *Perception & Psychophysics, 33*(5), 425–436. [PubMed]

Gilchrist, A. L., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., Spehar, B., et al. (1999). An anchoring theory of lightness perception. *Psychological Review, 106*(4), 795–834. [PubMed] [Article]

Katz, D. (1935). *The world of colour.* London: Kegan Paul, Trench, Trubner.

Lambert, J. H. (1760). *Photometria Sive Mensura et Gradibus Luminis, Colorum et Umbrae (Augsburg)*. Abridged German translation by E. Anding in Ostwald's Klassiker der exakten Wissenschaften, Nos. 31, 32, and 33 (Leipzig).

Land, E. H., & McCann, J. J. (1971). Lightness and Retinex theory. *Journal of the Optical Society of America, 61,* 1–11. [PubMed]

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

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. 387–416). Cambridge: 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., & Schirillo, J. A. (2002). Color constancy, lightness constancy, and the articulation hypothesis. *Perception, 31,* 135–139. [PubMed]

Mizokami, Y., Ikeda, M., & Shinoda, H. (1998). Lightness change as perceived in relation to the size of recognized visual space of illumination. *Optical Review, 5,* 315–319.

Rutherford, M. D., & Brainard, D. H. (2002). Lightness constancy: A direct test of the illumination-estimation hypothesis. *Psychological Science, 13*(2), 142–149. [PubMed] [Article]

Schirillo, J. A., Reeves, A., & Arend, L. E. (1990). Perceived lightness, but not brightness, of achromatic surfaces depends on perceived depth information. *Perception & Psychophysics, 48,* 82–90. [PubMed]

Schirillo, J. A., & Shevell, S. K. (1993). Lightness and brightness judgments of coplanar retinally noncontiguous surfaces. *Journal of the Optical Society of America A, 10*(12), 2442–2452. [PubMed]

Thouless, R. H. (1931). Phenomenal regression to the real object, Parts 1 & 2. *British Journal of Psychology, 21,* 339–359, and *22,* 1–30.

Wheatstone, C. (1838). On some remarkable, and hitherto unobserved, phenomena of binocular vision. *Philosophical Transactions of the Royal Society of London, 33,* 371–394.

Yamauchi, Y., & Uchikawa, K. (2005). Depth information affects the judgment of the surface-color mode appearance. *Journal of Vision, 5*(6), 515–524, http://journalofvision.org/5/6/3/, doi:10.1167/5.6.3. [PubMed] [Article]

Yang, J. N., & Maloney, L. T. (2001). Illuminant cues in surface color perception: Tests of three candidate cues. *Vision Research, 41*(20), 2581–2600. [PubMed]