Color Perception of 3D Objects: Constancy with Respect To Variation of Surface Gloss

Bei Xiao* David H. Brainard#

University of Pennsylvania

ABSTRACT

What determines the color appearance of real objects viewed under natural conditions? The light reflected from different locations on a single object can vary enormously. This variation is enhanced when the material properties of the object are changed from matte to glossy. Yet humans have no trouble assigning a color name to most things. We studied how people perceive the color of spheres in complex scenes. Observers viewed graphics simulations of a three-dimensional scene containing two spheres, test and match. The observer's task was to adjust the match sphere until its color appearance was the same as that of the test sphere. The match sphere was always matte, and observers varied its color by changing the simulated spectral reflectance function. The surface gloss of the test spheres was varied across conditions. The data show that for fixed test sphere body reflectance, color appearance depends on surface gloss. This effect is small, however, compared to the variation that would be expected if observers simply matched the average of the light reflected from the test.

CR Categories: J.4 [Social and Behavioral Sciences]: Psychology—; J.2 [Physical Sciences and Engineering]: Physics––; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, Shading, Shadowing, and Texture.

Keywords: color, visual psychophysics, surface gloss, human perception and performance, illumination

* email: beixiao@mail.med.upenn.edu

email: brainard@psych.upenn.edu

1 Introduction

A salient perceptual feature of many objects is that they have a color. The blueness of the mugs shown in Figure 1, for example, is readily perceived. Yet a closer look at the figure reveals large variation in the spectral properties of the individual locations on the mugs' surfaces. Some of the variation arises because the light reflected by the object's surface depends on the direction of the incident and reflected light, and some arises because of the geometrical pattern of the illumination. Thus the shape and pose of an object made from a single homogenous material can interact with the geometrical properties of the illumination to produce large changes in the light reflected from different locations on the object's surface. This means that perceiving an object as having a unified color requires extracting and integrating color information from multiple object locations.

In addition, objects can be made of a variety of materials. One way that materials vary, which we will focus on in this paper, is surface gloss. The two mugs shown in Figure 1 share the same diffuse reflectance or body color, but differ in their glossiness. The presence of gloss increases the variation of reflected light across the second mug's surface. In addition to being interested in how color is assigned to three-dimensional objects, we are interested in the effect of surface gloss on color appearance.

Previous studies of object color perception have focused on the appearance of flat matte surfaces under varying illumination [Breneman 1987; Brainard and Wandell 1992; Brainard 1998]. Recently, there have been studies of how the scene geometry

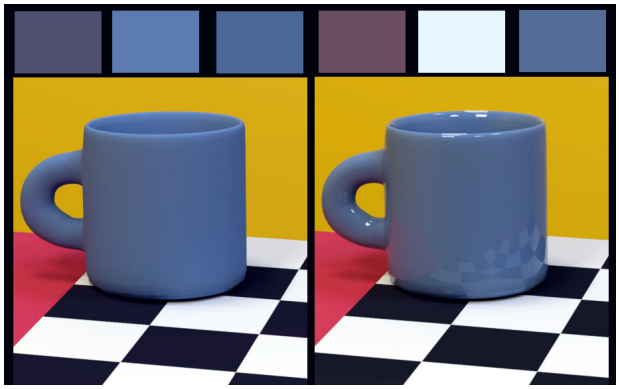

Figure 1. Reflected light varies across the surface of three-dimensional objects. The left panel shows an image of a matte mug in a synthetic scene, while the right panel shows a glossy mug in the same scene. The diffuse component of the reflectance of the two mugs is identical. The colored squares at the top of each panel show the color of three individual pixels from each mug. The samples above each panel are taken from corresponding locations on the two mugs. There is large variation in the spectrum of the light reflected from different regions of each mug, and this variation is larger for the glossy mug.

affects the perception of lightness and color of such surfaces [Kraft and Brainard 1999; Ripamonti et al. 2004; Yang and Maloney 2001; Boyaci et al. 2003; Doerschner et al. 2004; Boyaci et al. 2004]. In these experiments, the observer does not need to integrate a range of test object chromaticities and luminannces, as the reflected light is essentially uniform across the flat object. There has also been increasing interest in exploring how people perceive object material properties such as surface gloss [Hunter and Harold 1987; Fleming et al. 2003; Fleming et al. 2004; Pellacini et al. 2000; Obein et al. 2004]. Little is known, however, about how we perceive the color of three-dimensional objects made of different surface materials. In this paper, we report psychophysical experiments designed to explore this question.

In our experiments, observers were asked to match the color appearance of a match sphere to that of a test sphere. The match sphere was always matte, while the test sphere was varied from matte to glossy. The data reject the simple hypothesis that the color matches depend only on the spatial average of the light reflected from the test sphere. The data also show that observers exhibit some stability of object color perception in the face of variation of surface gloss.

2 Method

2.1 Stimuli

2.1.1 Surface reflectance and scene contents

Observers viewed objects in a rendered scene, as in Figure 2. The rendered room was specified to have dimensions 36 cm (width) × 40 cm (height) × 28 cm (depth). The front wall of the rendered room was opaque but contained an aperture through which the contents of the room could be seen. The size of the aperture determined the size of the visible stimulus image, which was 21.5 cm (width) × 24.5 cm (height). The test sphere was on the left of the rendered room and the match sphere was on the right. The spheres were located approximately in the middle of the room and their specified diameters were 2 cm. The viewing position used for rendering was at the same height as the center of the spheres and was 76.4 cm from the center of the spheres. In the experiments, the monitors were positioned at this distance from the observer.

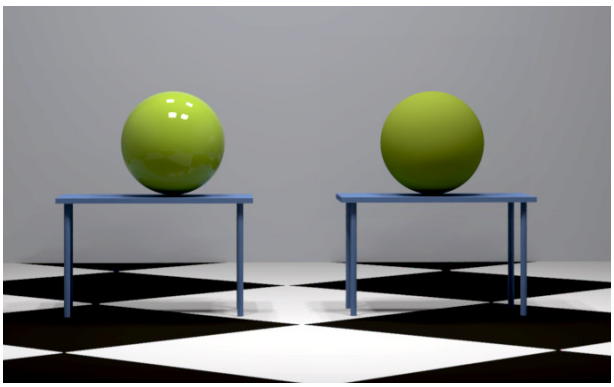

Figure 2. Experimental Scene. The figure shows the central portion of the rendered image. The scene contains two spheres. The sphere on the left is the test sphere. Its surface reflectance properties were varied from trial to trial. The sphere on the right is the match sphere. It was always matte, and observers adjusted its diffuse reflectance component to match the color appearance of the test sphere.

The surface reflectance of all objects in the scene conformed to the isotropic version of Ward light reflection model [Ward 1992]. This model represents surface reflectance as the sum of two components, diffuse and specular:

$$\rho(\theta_i,\phi_i,\theta_o,\phi_o,\lambda)=\frac{\rho_d(\lambda)}{\pi}+\rho_s\cdot\frac{\exp\left[-\tan^2\delta/\alpha^2\right]}{4\pi\alpha^2\sqrt{\cos\theta_i\cos\theta_o}} \quad (1)$$

where $\rho(\theta_i,\phi_i,\theta_o,\phi_o,\lambda)$ is the surface BRDF, $\rho_d(\lambda)$ is the diffuse reflectance, which depends on wavelength, ρ_s governs the strength of the specular component, and α is a roughness parameter that describes the spread of the specular highlight. The angles θ_i and ϕ_i describe the direction of a light ray incident to the object, while θ_o and ϕ_o describe the direction of the reflected light. The quantity δ is computed from the four angles as described in Ward [Ward 1992].

We varied $\rho_d(\lambda)$ to control the body color of objects, and varied (ρ_s,α) to control surface gloss. Two different body colors and five different glossinesses were used for the test sphere. The body colors were chosen to have the diffuse reflectance spectra of two of the squares on the Macbeth Color Checker. The top row of Figure 3 shows the five purple test spheres used. The material parameters are listed in Table 1. The same materials were used for the yellow-green tests. The match sphere was always matte and its body color was adjusted by the observer. The rest of the objects in the scene were held fixed throughout the experiment and were as shown in Figure 2. The simulated scene was illuminated by four area lights and a diffuse illuminant. The four area lights were located on the ceiling of the rendered room. All of the simulated illuminants had the spectrum of CIE D65 [CIE 1986.

Conditions	ρ_s	α
Matte	0.00	0.00
Condition A	0.08	0.02
Condition B	0.12	0.18
Condition C	0.02	0.00
Condition D	0.18	0.12

Table 1. Test Sphere Parameters. This table provides the BRDF parameters that were used to render each of the test spheres.

2.1.2 Image generation

The scenes were modeled using the Maya (Alias, Inc., San Rafael, CA) software tools. Model scenes were then exported to custom MATLAB software (Mathworks, Inc., Natick, MA). This software associated full spectra with each illuminant in the modeled scene, and parameters $\rho_d(\lambda)$, ρ_s, and α with each object. Spectra were represented using 10 nm sampling between 400 and 700 nm (31 wavelength samples). The MATLAB software converted the scene representation to a format appropriate for the RADIANCE renderer [Ward 1994] and invoked RADIANCE 31 times (once for each wavelength). This process resulted in a 31-plane hyperspectral image of the scene.

This was converted to a three-plane LMS representation by computing at each pixel the excitations that would be produced in the human L-, M-, and S-cones. The LMS image was tone-mapped by truncating the pixel luminances at five times the mean luminance. It was then converted for display on the computer monitor using standard methods [Brainard et al. 2002] together with measurements of the monitor's phosphor emission spectra and gamma functions (PR-650 spectral radiometer, Photo Research, Inc., Chatsworth, CA). As part of the display conversion, images were scaled to occupy the luminance range available on the monitor.

2.2 Apparatus

Stimuli were presented to the observers stereoscopically, using two 21 in. CRT monitors (Model No. VCDTS 23852-2M, Viewsonic, Walnut, CA). Images delivered to the left and right eye were rendered separately using viewpoints separated horizontally by 6.3 cm. The size of each stereo image displayed on the corresponding monitor was 24.5 cm (width) × 21.5 cm (height), corresponding to the size of the aperture in the front wall of the rendered room. The monitors were driven at a 75 Hz refresh rate at 1152 by 870 spatial resolution, and with 14-bits of intensity control for each color channel (BITS++, Cambridge Research Systems, Rochester, England).

2.3 Procedure

The observer's task was to adjust the diffuse component of the match sphere until its color appearance was the same as that of the test sphere. The observer controlled the match sphere using a game pad, which provided adjustments along the CIELAB L* (lightness), a* (red/green), and b* (blue/yellow) coordinates of the rendered sphere. These coordinates were taken as those computed from the average LMS values of all pixels on the match sphere, and software transformed between CIELAB, LMS, and the underlying reflectance representation.

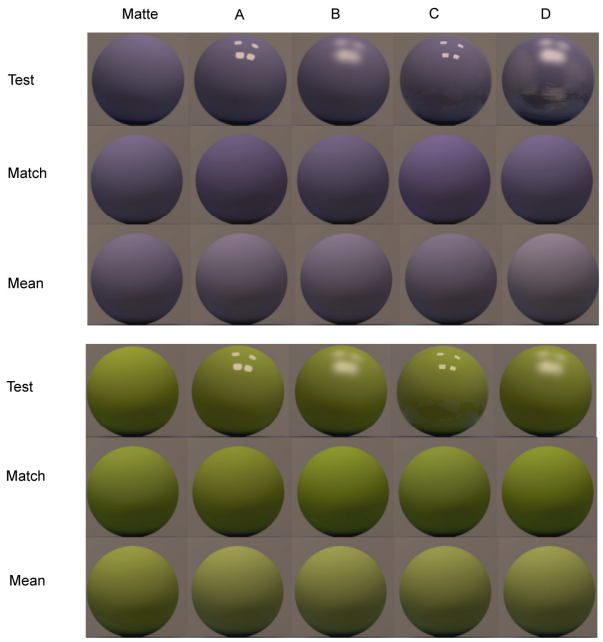

Figure 3. Pictorial representation of the data. Top panel: The upper row shows the purple test spheres. All have the same diffuse reflectance (body color) and each has a distinctive surface gloss (matte, conditions A-D). The second row shows observer HTS's mean matched spheres for the corresponding purple test spheres. The third row shows matte spheres with the same average LMS values as the corresponding test spheres. Bottom panel: HTS's data for yellow-green test spheres, same format as the top panel. Note that the figure is reproduced without calibration, so that the pictorial representation is only illustrative.

Because it was not feasible to re-render the entire displayed images in real time, an approximation was used. We pre-rendered four images for each eye, containing red, green, blue, and white match spheres. These images were linearly combined using weights computed from the desired LMS coordinates of the match sphere and those of the precomputed images. We verified by comparison with a directly rendered image that this method provided a good approximation - the mean error (taken over all approximated pixels) was 1%, while the maximum error was 3%.

As they performed the match, observers had several adjustment step sizes available. When the observer was satisfied with a match, he or she indicated this with a button press. Each match began with a randomly chosen color for the match sphere. Observers made three matches for each of the 10 conditions (2 diffuse reflectances, 5 material properties).

2.4 Observers

Seven observers completed the experiment. All had normal color vision, 20/20 corrected acuity, and normal stereopsis. One observer was the first author, one was a member of the lab familiar with the experiment, and five were naïve paid observers. Before the data were collected, observers were asked to complete practice trials. The test spheres used in the practice trials had diffuse reflectances different from those used in the actual experiment. Observers repeated each of the 6 practice conditions (two diffuse reflectances, three material properties) 3 times.

3 Results

3.1 Pictorial representation of the data

One way to represent the data is to show pictures of test and match spheres that observers judged equal in color. Figure 3 shows data from one observer in this form. We averaged the LMS coordinates of the three matches set by this observer and produced the match sphere corresponding to the result. In each panel, the five test spheres are shown in the top row while the five corresponding matched spheres are shown in the second row. Overall, the matched spheres were quite similar. For the purple test spheres, however, there are some systematic differences: as the test spheres become glossier (left to right), the matched spheres become more saturated. This suggests that there is an effect of surface gloss on this observer's object color perception, when the diffuse reflectance is held fixed. This effect is less obvious for the yellow-green spheres.

3.2 Quantitative representation of the data

To produce a quantitative representation of the data, we calculated the spatial average of the LMS values across both test and matched spheres and transformed them to CIELAB coordinates. Figure 4 shows the data from two observers for the purple tests in this format. The filled symbols correspond to the test spheres shown in Figure 3 (top panel, first row). The open symbols in the top two panels correspond to the matched spheres (top panel, second row). We can make a few observations from this

representation. First, there is variation in the matches with the change of gloss of the tests. Second, this variation differs between the observers: HTS's data is more spread than KEE's. Third, the variation in the tests does not appear to predict directly the variation in the matches for either observer.

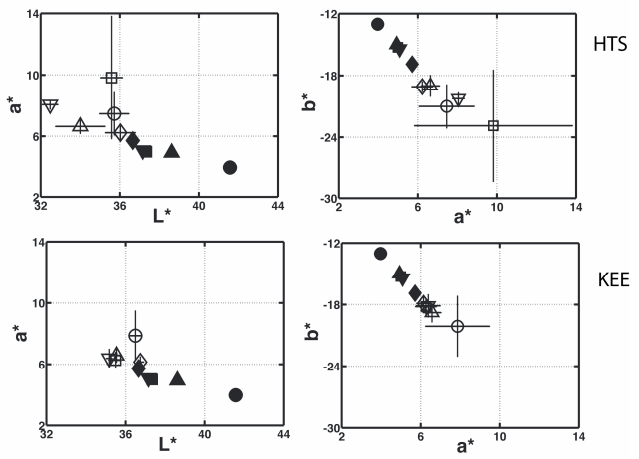

Figure 4. CIE Lab data representation. The top panels show data for HTS for the purple tests, the bottom panels show data for KEE. Filled symbols show the spatial average of the test spheres, while open symbols show the spatial average of the matched spheres. Left panels plot a* vs. L* while right panels plot b* vs. a*. Diamonds: matte test; downward triangles: condition A (see Figure 3); upward triangles: condition B; squares: condition C; circles: condition D.

3.3 Dimensionality reduction

It is difficult to visualize data in three dimensions. In exploratory analyses, we discovered that the test and matched spheres clustered near a plane in the three-dimensional LMS space. To take advantage of this fact, we constructed a new coordinate system to represent the data for each test color. The direction of the first axis, which we call the diffuse direction, is given by the LMS coordinates of the matte test sphere. Variation along this direction corresponds to a change in the overall intensity of the diffuse component. The direction of the second axis, which we call the specular direction, is given by the difference in LMS coordinates between the glossiest test sphere (condition D) and the matte test sphere. Increasing values along this dimension corresponds closely to adding the LMS coordinates of the illuminant to those of the matte test sphere. The set of test spheres differs almost entirely along this dimension, and thus variation along this dimension may be visualized by the set of matte spheres shown in the third row of each panel of Figure 3. A third axis in the transformed space is taken as perpendicular to the first two. As shown in Figure 5, all of the test and matched spheres had a value of essentially zero along this third axis. This allows us to represent the data by the coordinates of the test and matched spheres on the diffuse and specular axes of the transformed space. Note that the transformed representation is specific to each test sphere color. The matte test sphere always has coordinates (1, 0, 0) in the transformed space, while the condition D test sphere always has coordinates (1, 1, 0).

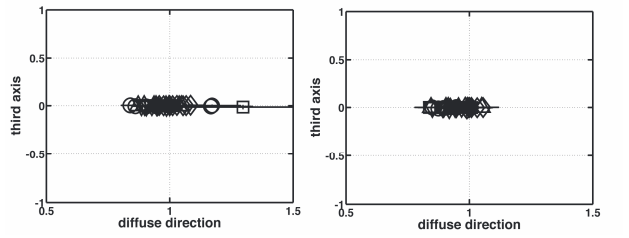

Figure 5. Data dimensionality in the transformed space. This figure shows all observers' matches and the tests plotted in the transformed space. The left panel plots the third axis against the diffuse direction for the purple test condition. The right panel is in the same format but is for the yellow-green test condition.

The advantages of the transformed representation are a) that it allows two-dimensional visualization of the data, b) that its axes have meaningful physical interpretation, and c) that it allows us to examine data from each test sphere color against a comparable physical metric defined by the matte test sphere. Figure 6 replots the data for HTS and KEE in the new space. Comparison with Figure 4 shows that the same features of the data are visible.

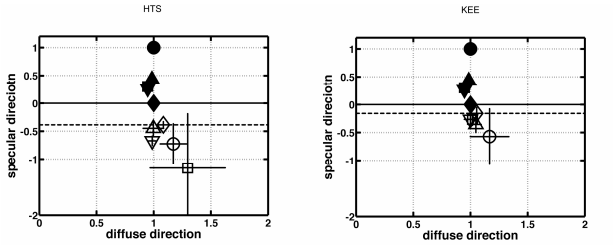

Figure 6. Transformed data representation. The left panel shows data for HTS for the purple tests plotted in the transformed space, the right panel shows data for KEE. Both panels plot specular direction against diffuse direction. The symbols follow the same format as in Figure 4. The horizontal dashed line shows the value of the match to the matte test along the specular direction.

It is useful to separate two sorts of effects in the data. First, the match to the matte test sphere does not have the same coordinates as the matte test. Since these two spheres are made from the same material, this represents an effect arising from some asymmetry in the location of the two spheres within the scene. This effect is not due to material properties and is not the focus of this paper. Second, the matched spheres do not all have the same coordinates as each other. This is an effect of material property on the matches. To separate these two effects, we can recenter the matches by shifting them so that the match for the matte test sphere overlays the matte test at (1, 0) in the new coordinate system. Figure 7 shows the data, recentered in this way for all observers, for both test spheres. We can see in Figure 7 that the spread in the matched spheres does not correspond to the spread in the test spheres. Indeed, the data appear to rule out one simple

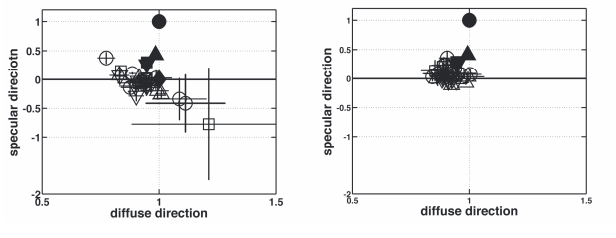

Figure 7. Transformed and recentered data representation for all observers. Left panel: tests and all seven observers' matches for the purple condition. Each observer's matches have been recentered as described in the text. Right panel: tests and all seven observers' matches for the yellow-green condition.

hypothesis: that observers match the color of two spheres by matching the spatial average of their LMS coordinates. If this were the case, the recentered test and matched spheres should overlay but they do not. For both test sphere colors, the tests are arrayed along the vertical axis. For the purple tests, the matches trace out a diagonal line, while for the yellow-green tests, the matches cluster near the point (1, 0).

3.4 Material influence index

To quantify the stability of color perception across changes in surface gloss, we computed a material influence index (MII). This index measures the distance between each match to a glossy test sphere and the match to the matte test sphere, relative to the distance between the two test spheres. We computed distances in the two-dimensional representation of Figure 7, after the recentering process described above. The formula for the material influence index is:

$$MII = \left[|mg - mm| / |tg - tm| \right] \quad (2)$$

where *mg* represents the coordinates of the match to the glossy test, *mm* represents the match to the matte test, *tg* represents the glossy test, and *tm* represents the matte test. Because of the recentering, *mm* = *tm* = (1, 0). A small value of the MII indicates a match that is stable with respect to the change in surface gloss, with 0 indicating complete stability. A high value of the MII means the match has changed considerably with the change in surface gloss. If the match shifted the same amount as the test, then the MII would take on a value of 1. Figure 8 shows histograms of the MII for all the observers' matches. Both histograms are shifted towards 0 from 1, indicating that in general the matches are not perturbed as much as the tests with changes in material properties. The mean index for the purple condition was 0.45, while the mean index for the yellow-green condition was 0.32.

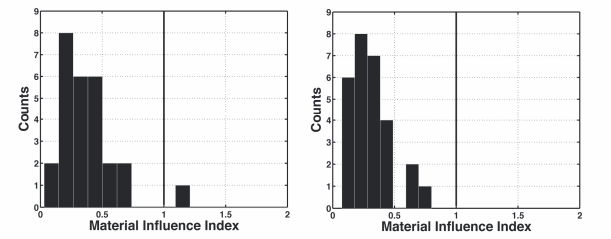

Figure 8. Histograms of material influence indices. Left panel: histogram of material influence indices obtained from the matches for the purple condition. Right panel: corresponding histogram for the yellow-green condition.

4 Discussion

This paper provides a starting point for understanding how the visual system integrates information across three-dimensional objects to determine color appearance. Such understanding is crucial if we are to generalize theories of color vision to allow prediction of the appearance of actual objects in real scenes. For such scenes, the light reflected from different locations on an object made of a uniform material can differ considerably (Figure 1). We used a matching method to study how varying surface gloss affects the color appearance of simulated test spheres, when the diffuse reflectance of the test spheres was held fixed. We report two main findings. First, changing surface gloss affects color appearance. Second, the magnitude of this effect is smaller than predicted from the change in the average reflectance of the test sphere. Consequently, the procedure used by the visual system to compute object color appearance is not a simple averaging of the light reflected across the image of the object.

Given that a simple averaging hypothesis cannot account for the data, how might we approach developing an adequate model? One possibility is some variant of the averaging hypothesis. One can imagine other image-based models that compute simple statistics from the image of the object and use these to predict color appearance, without reference to the three-dimensional structure of the scene. On the other hand, one can also consider scene-based or inverse-optics models. For example, perhaps observers estimate the geometry of the light source and the BRDF of the spheres. From the latter, they then compute color from the estimated diffuse reflectance component. Models of this general sort have been successful for understanding the color appearance of matte surfaces across a variety of manipulations [Bloj et al. 1999; Boyaci et al. 2003; Bloj et al. 2004; Brainard et al. 1997; Doerschner et al. 2004]. To develop such a model we can make use of advances in inverse-rendering algorithms for image processing and graphics applications [Ramamoorthi and Hanrahan 2001]. Our goal in extending the research will be to formulate models of both types and develop a body of empirical data that can be used to evaluate them.

References

BLOJ, A., KERSTEN, D. and HURLBERT, A. C. 1999. Perception of three-dimensional shape influences colour perception through mutual illumination. Nature, 402, 877-879.

BLOJ, M., RIPAMONTI, C., MITHA, K., GREENWALD, S., HAUCK, R. and BRAINARD, D. H. 2004. An equivalent illuminant model for the effect of surface slant on perceived lightness. Journal of Vision, 4, 735-746.

BOYACI, H., DOERSCHNER, K. and MALONEY, L. T. 2004. Perceived surface color in binocularly viewed scenes with two light sources differing in chromaticity. Journal of Vision, 4, 664-679.

BOYACI, H., MALONEY, L. T. and HERSH, S. 2003. The effect of perceived surface orientation on perceived surface albedo in binocularly viewed scenes. Journal of Vision, 3, 541-553.

BRAINARD, D. H. 1998. Color constancy in the nearly natural image. 2. achromatic loci. Journal of the Optical Society of America A, 15, 307-325.

BRAINARD, D. H., BRUNT, W. A. and SPEIGLE, J. M. 1997. Color constancy in the nearly natural image. 1. asymmetric matches. Journal of the Optical Society of America A, 14, 2091-2110.

BRAINARD, D. H., PELLI, D. G. and ROBSON, T. 2002. Display characterization. In Encylopedia of Imaging Science and Technology. HORNAK, J., Wiley, 72-188.

BRAINARD, D. H. and WANDELL, B. A. 1992. Asymmetric color-matching: how color appearance depends on the illuminant. Journal of the Optical Society of America A, 9, 1433-1448.

BRENEMAN, E. J. 1987. Corresponding chromaticities for different states of adaptation to complex visual fields. Journal of the Optical Society of America A, 4, 1115-1129.

DOERSCHNER, K., BOYACI, H. and MALONEY, L. T. 2004. Human observers compensate for secondary illumination originating in nearby chromatic surfaces. Journal of Vision, 4, 92-105.

FLEMING, R. W., DROR, R. O. and ADELSON, E. H. 2003. Real-world illumination and the perception of surface reflectance properties. Journal of Vision, 3, 347-68.

FLEMING, R. W., TORRALBA, A. and ADELSON, E. H. 2004. Specular reflections and the perception of shape. Journal of Vision, 4, 798-820.

HUNTER, R. S. and HAROLD, R. W. 1987. The Measurement of Appearance. Wiley, New York.

KRAFT, J. M. and BRAINARD, D. H. 1999. Mechanisms of color constancy under nearly natural viewing. Proc. Nat. Acad. Sci. USA, 96, 307-312.

OBEIN, G., KNOBLAUCH, K. and VIÉNOT, F. 2004. Difference scaling of gloss: nonlinearity, binocularity, and constancy. Journal of Vision, 4 711-720.

PELLACINI, F., FERWERDA, J. A. and GREENBERG, D. P. 2000. Toward a psychophysically-based light reflection model for image synthesis. Computer Graphics, 34, 55-64.

RAMAMOORTHI, R. and HANRAHAN, P. 2001 A signal-processing framework for inverse rendering. ACM SIGGRAPH 2001, ACM Press/ACM SIGGRAPH, 117-108.

RIPAMONTI, C., BLOJ, M., HAUCK, R., MITHA, K., GREENWALD, S., MALONEY, S. I. and BRAINARD, D. H. 2004. Measurements of the effect of surface slant on perceived lightness. Journal of Vision, 4, 747-763.

WARD, G. J. 1992 Measuring and modeling anisotropic reflection. Computer Graphics, 26, 265-72.

WARD, G. J. 1994 The RADIANCE lighting simulation and rendering system. Proceedings of ACM SIGGRAPH 1994, ACM Press/ACM SIGGRAPH, 459-472.

YANG, J. N. and MALONEY, L. T. 2001. Illuminant cues in surface color perception: tests of three candidate cues. Vision Research, 41, 2581-2600.