

Surface color perception in three-dimensional scenes

HUSEYIN BOYACI,¹ KATJA DOERSCHNER,² JACQUELINE L. SNYDER,²
AND LAURENCE T. MALONEY^{2,3}

¹Department of Psychology, University of Minnesota, Minneapolis, Minnesota

²Department of Psychology, New York University, New York, New York

³Center for Neural Science New York University, New York, New York

(RECEIVED March 10, 2006; ACCEPTED March 10, 2006)

Abstract

Researchers studying surface color perception have typically used stimuli that consist of a small number of matte patches (real or simulated) embedded in a plane perpendicular to the line of sight (a “Mondrian,” Land & McCann, 1971). Reliable estimation of the color of a matte surface is a difficult if not impossible computational problem in such limited scenes (Maloney, 1999). In more realistic, three-dimensional scenes the difficulty of the problem increases, in part, because the effective illumination incident on the surface (the light field) now depends on surface orientation and location. We review recent work in multiple laboratories that examines (1) the degree to which the human visual system discounts the light field in judging matte surface lightness and color and (2) what illuminant cues the visual system uses in estimating the flow of light in a scene.

Keywords: Surface color perception, Binocular perception, Light field, Plenoptic function, Color constancy, Lightness constancy

Introduction

The spectral power distribution of light emitted by the sun is almost invariant and the range of daylight we experience results from sunlight interacting with earth’s atmosphere (Henderson, 1977). The resulting spectral distribution of daylight across the sky is typically spatially inhomogeneous and constantly changing (Lee & Hernández-Andrés, 2005a, 2005b). Furthermore, objects in the scene create shadows or act as secondary light sources, adding complexity to the *light field* (Gershun, 1936/1939) or *plenoptic function* (Adelson & Bergen, 1991) that describes the spectral power distribution of light arriving from every direction at every point in the scene.¹

When the light field is inhomogeneous, the light absorbed and re-radiated by a matte smooth surface patch can vary markedly with patch orientation or location in the scene.

In Fig. 1, for example, we illustrate the wide range of light emitted by identical rectangular achromatic matte surfaces at many orientations, illuminated by a distant neutral punctate light source.

Any visual system designed to estimate surface color properties (including lightness) are confronted with a new problem to solve with each change of surface orientation or location.

To arrive at a stable estimate of surface reflectance the visual system has to discount the effects of the light field on the patch. At first glance, discounting spatial variation in the light field is an under-constrained problem (Adelson & Pentland, 1996; Belhumeur et al., 1999; Dror et al., 2004) and in many scenes it, in fact, is.

The Mondrian singularity

Previous research in color vision has typically avoided addressing the problems introduced by spatial and spectral variation in the light field by choice of scenes. These scenes, consisting of flat, co-planar matte surfaces were referred to as Mondrians (Land & McCann, 1971). In such scenes, observers can accurately make a variety of judgments concerning surface color and lightness. Arend and Spehar (1993a, 1993b) showed that observers were able to estimate lightness of a matte surface embedded in a two-dimensional Mondrian nearly constantly. Foster and Nascimento (1994; Nascimento & Foster, 2000) showed that observers can reliably distinguish whether the change in appearance of a Mondrian is due to an “illumination” change or a reflectance change, and that this could be explained by a model based on cone-excitation ratios. Bäuml (1999) showed that observers are capable of constant estimation of color of surfaces in Mondrians following changes in illumination and showed that his results could well be accounted by using the von Kries principle, which is a simple linear transformation of

Address correspondence and reprint request to: Huseyin Boyaci, University of Minnesota at Twin Cities, Department of Psychology, 75 East River Rd., 219, Minneapolis, MN 55455, USA. E-mail: boyac003@umn.edu


¹The light field, as defined in Gershun (1936/1939), is radiance as a function of position and direction in regions of space free of occluders. We use the term ‘light field’ in a more general sense to refer to the spectral power distribution of light arriving at every point in the scene. It captures what a radiospectrophotometer placed at each point in the scene pointing in all possible directions would record.




Fig. 1. The Effect of Orientation: identical rectangular matte patches at different locations and orientations, rendered under a distant neutral punctate source placed along the line of sight. The luminance of each patch is proportional to the cosine of the angle between the surface normal and the direction to the punctate source (Lambert's Law; Haralick & Shapiro, 1993).

cone responses. However, those studies need to be extended for several reasons.

First, there is no obvious way to generalize those results to normal viewing conditions of our ever changing 3-dimensional world. In the flat, two-dimensional world of Mondrians no matter how complex the light field is, the light emitted by a surface contains essentially no information about the spatio-spectral distribution of the light incident on the surface (Maloney, 1999). A matte surface absorbs light from all directions in a hemisphere centered on its surface normal and then re-emits uniformly in all directions a fraction of the total light absorbed. That is, a matte surface “forgets” where the light came from. In Fig. 2A, we illustrate this point: the rectangular regions of a painting by Piet Mondrian have been replaced by Lambertian idealizations rendered under a particular light field generated by a combination of neutral punctate and diffuse sources. Given only Fig. 2A, there is no way to tell that the light field used in rendering the scene is as advertised. The lighting could be composed of zero, one or more distant punctate or extended sources, or it could be just diffuse. The sources could differ in chromaticity as long as the net light impinging on the surfaces is neutral. For Mondrian stimuli, when light sources are distant, there is essentially no information about the spatial and spectral distribution of the light field. Mondrian stimuli exist in a kind of singularity where remarkably little information about the light field is available. We call this phenomenon the *Mondrian singularity*. In Fig. 2B, we have rotated some of the matte surfaces in the Mondrian of Fig. 2A and added a specular sphere, thereby



(A)



(B)

Fig. 2. The “Mondrian Singularity.” (A) A rendered scene whose spatial structure resembles *Composition with Red, Yellow, Blue and Black*, by Piet Mondrian, 1921. A combination of punctate and diffuse light sources illuminate the scene. Viewing this scene, the observer has little or no information about the spatial distribution of light sources in the scene. See text. (B) A rendered scene based on the same painting but with several of the rectangular regions rotated and/or displaced and a specular sphere added. The observer now has cues to the distribution and chromaticity of the light sources present.

revealing more information about the light sources used in rendering Fig. 2A and 2B.

Another reason why it is important to consider a wider range of stimuli in evaluating human color perception is that three-dimensional scenes can convey considerable information about the light field in a scene, as suggested by a comparison of Fig. 2A and 2B. Maloney (1999) noted that there are potential “cues to the illuminant” in three-dimensional scenes that signal illuminant chromaticity. Here we also consider recent work directed to determining what cues signal how the intensity and chromaticity of illumination incident on a matte surface varies with surface orientation and location in three-dimensional scenes.

Illuminant cues

Surface color perception in Mondrian scenes is an intrinsically difficult problem. In order to estimate surface color accurately, the visual system must estimate the net intensity and chromaticity of the light incident on the Mondrian. The typical approach taken is to develop simple measures of central tendency, variance, and covariance of the photoreceptor excitations and use them as a basis for estimating light intensity and chromaticity. “Gray world” algorithms (for reviews see Hurlbert, 1998; Maloney, 1999), for example, use the mean chromaticity of the scene as an estimate of the chromaticity of the light. Mausfeld and Andres (2002) have conjectured that means, variances and covariances contain all of the information used by the visual system in estimating surface color. Golz and MacLeod (2002; MacLeod & Golz, 2003) concluded that correlations between the chromaticities and luminance values of surfaces contained useful information about the chromaticity of the effectively-uniform illumination of a Mondrian scene, but this conclusion has been challenged by recent work (Ciurea & Funt, 2004; Granzier et al., 2005). These measurements, based on simple moments (mean, variance, covariance) of distributions, eliminate what little spatial structure is present in the Mondrian. It is not clear that simple moments derived from Mondrian scenes convey any information about the chromaticity of the illuminant or its intensity (see Maloney, 1999), and they convey no information about spatial and spectral inhomogeneities in the light field.

When a scene is not restricted to flat, co-planar matte surfaces arranged in a Mondrian, more information about the chromaticity of the illuminant (Maloney, 1999; Yang & Maloney, 2001) and the spatial and spectral distribution of the light field may be available to the observer. Consider, for example, Fig. 2B. In this particular scene, the observer may use cast shadows to estimate that the direction of the light is from the left and slightly above. We emphasize that any deviation from “flat” or “matte” in an otherwise Mondrian stimulus could disclose information about the light field, and we refer to these sources of information as *illuminant cues* (Kaiser & Boynton, 1996; Maloney, 1999; Yang & Maloney, 2001; Pont & Koenderink, 2003, 2004; Koenderink & van Doorn, 1996; Koenderink et al., 2004). Illuminant cues, by definition, carry information about the illuminant. It is possible to develop algorithms that recover the light field from such cues, because they are available from single or multiple images of a given object or scene (see e.g., Hara et al., 2005 for an overview) but currently such algorithms depend on restrictive assumptions about the scene and its illumination. These algorithms are based on the physics of image formation, and, when they succeed, we can be sure they carry the desired information. The relevant question concerning an illuminant cue is whether it is used in human vision.

In this article we review recent work from a small number of research groups concerning how biological visual systems extract information about surfaces (albedo, color) in scenes outside the Mondrian singularity. We describe in more detail two sets of experiments, the first testing whether human observers can compensate for changes in surface orientation and examining what illuminant cues they may be using, and the second examining whether human observers can compensate for changes in surface location in scenes with a strong illuminant gradient in depth. These results suggest that human color vision is well-equipped to solve these apparently more complex problems in surface color perception outside the Mondrian singularity (for a review see Maloney et al., 2005).

Lightness and color perception with changes in orientation

Boyaci et al. (2003) investigated how human observers compensate for changes in surface orientation in binocularly-viewed, computer-rendered scenes illuminated by a combination of neutral punctate and diffuse light sources. The simulated punctate source was sufficiently far from the rendered scene that it could be treated as a collimated source. The observer’s task was to match the lightness (perceived albedo) of a test surface within the scene to a nearby lightness scale. The orientation of the test patch with respect to the punctate source was varied, and the question of interest was whether observers would compensate for test patch orientation, partially or fully. Previous work had found little or no compensation (Hochberg & Beck, 1954; Epstein, 1961; Flock & Freedberg, 1970; Redding & Lester, 1980; see Boyaci et al., 2003 for details).

The methods and stimuli used are similar to those of Boyaci et al. (2004), which we present in more detail later in this study. In contrast to previous researchers, Boyaci et al. (2003) found that observers substantially compensated for changes in test patch orientation. Ripamonti et al. (2004) drew the same conclusions using scenes of similar design composed of actual surfaces (not computer-rendered) viewed under a combination of collimated and diffuse light sources. The conclusion of both studies is that the visual system compensates in part for changes in surface orientation in scenes, in which the lighting model consisted of a combination of a diffuse and a collimated source.

Boyaci et al. (2004) examined judgments of surface color in a similar experiment. The lighting model consisted of a distant punctate yellow light source (“sun”) and a diffuse blue light source (“sky”). The test surface was illuminated by a mixture of the two that depended on the orientation of the test surface and the lighting model. The observer’s task was to set the test patch to be achromatic (achromatic setting task). To do so, the observer first needed to estimate the blue–yellow balance of light incident on a test patch, which was, itself, part of the spatial organization of the scene. Next, the observer needed to set the chromaticity of the light emitted by the surface to be consistent with that of an achromatic surface.

The punctate light source was simulated to be behind the observer at elevation 30° and at azimuth -15° (on the observer’s left) or 15° (on the observer’s right). The location of light source remained constant during an experimental block. On every trial each of four naïve observers was presented with a rendered scene and asked to adjust a test surface to be achromatic. Scenes were rendered as a stereo image pair and viewed binocularly. A typical scene is shown in Fig. 3. The test patch was always in the center



Fig. 3. A scene from Boyaci et al. (2004). Observers viewed rendered scenes binocularly. The left two images permit crossed binocular fusion, the right two images, uncrossed. The scenes were rendered with a combination of yellow punctate and blue diffuse light sources. The punctate source was always behind the observer, to his left on half the trials, and to his right on the remainder. The orientation of the test surface varied in azimuth and elevation from trial to trial. The observer’s task was to set the test surface in the center of the scene to be achromatic.

and at the front of the scene, and additional objects were provided as possible cues to the light field.

From trial to trial, test patch orientation was varied in either azimuth or elevation, but not both. The dependent measure of interest was the relative blue intensity Δ^B in the observer’s achromatic setting (for details see Boyaci et al., 2004). Intuitively, as the angle between the normal to the test surface and the direction to the yellow punctate light source increases, the observer should

make achromatic settings that are “bluer.” Boyaci et al. (2004) derived setting predictions for an ideal observer who made achromatic settings that were color constant, always picking the setting consistent with a test surface that was achromatic. These setting predictions are plotted in Fig. 4A. There are two plots, one for the punctate light on the observer’s left and one for the light on the observer’s right. In each plot, the relative blue intensity Δ^B is plotted versus the azimuth of the test surface (solid curve) and

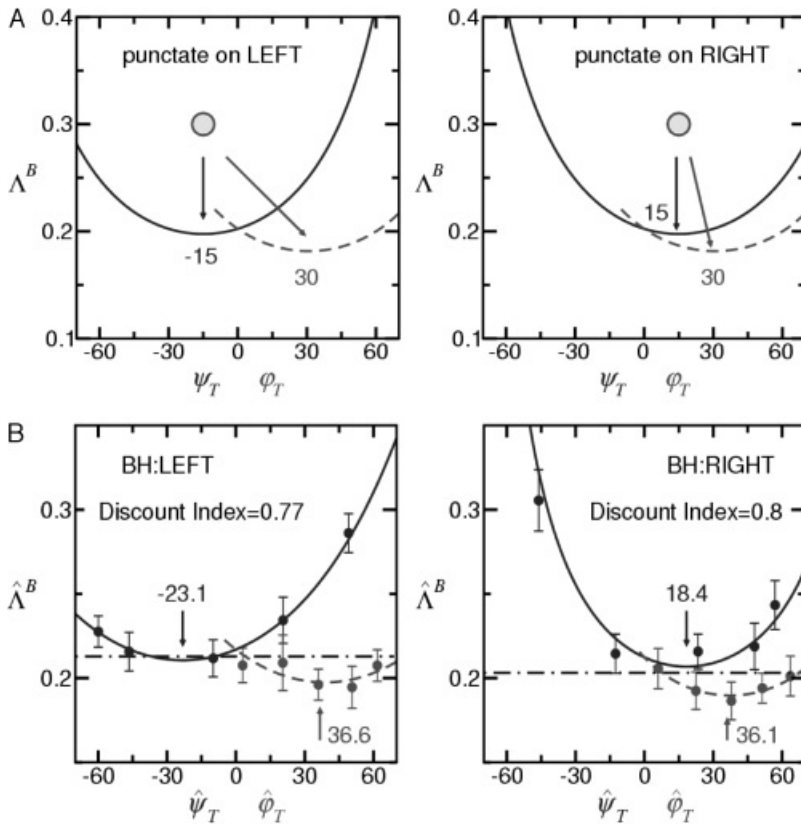


Fig. 4. Achromatic Setting Results from Boyaci et al. (2004). The dependent variable was the amount of blue (blue/total, or relative blue) in the observer’s achromatic setting. (A) The settings for an ideal observer who perfectly compensates for changes in test patch orientation and punctate source position. The left graph contains a plot of settings for trials where the punctate source was at 30° elevation, -15° azimuth (above and behind the observer, to his left). The right graph contains a plot of settings for trials where the punctate source was at 30° elevation, 15° azimuth (above and behind the observer, to his right). The horizontal axis for both graphs is used to plot either the azimuth or elevation of the test surface. The vertical axis is the relative blue in the observer’s settings. The solid curve in each graph signifies the settings of the ideal observer that compensates for changes in test surface azimuth. The dashed curve in each graph signifies the settings of the ideal observer that compensates for changes in test surface elevation. Note that both curves reach a minimum when the test surface is closest in azimuth or elevation to the “yellow” punctate source. (B) Settings of one observer from Boyaci et al. (2004). The format is identical to that of (A) The lines through the data are based on an equivalent light model not described here (see Boyaci et al., 2004).

versus the elevation (dashed curve). It is important to note that each curve reaches a minimum when the test patch's orientation matches the direction of the yellow punctate light source.

The results are shown in Fig. 4B for the subject closest to the ideal observer. All four subjects substantially discounted the effect of changes in orientation. Boyaci et al. (2004) were able to recover crude estimates of light source direction from each observer's achromatic settings by estimating and plotting the minima of the four curves (Fig. 5). There are four estimates of azimuth (one for each observer) for the light source on the left and four for the light source on the right. There are eight corresponding estimates of elevation. With one exception, the eight estimates of elevation are within 10 degrees of the true values

The outcome of the experiment in Boyaci et al. (2004) together with the results of Boyaci et al. (2003), and Ripamonti et al. (2004) imply that the observer's visual system effectively develops an

equivalent lighting model (Boyaci et al., 2003) in a scene and uses this model in estimating albedo and surface color. In order to do so, the visual system must use the cues present within the scene itself.

In a more recent experiment (Boyaci et al., 2006) we examined three possible "cues to the lighting model" that were present in the scenes described earlier: cast shadows, surface shading, and specular highlights. We asked the observers to judge the lightness of a central test patch at different orientations embedded in scenes that contained various cues to lighting model. The methods and stimuli were similar to those in Boyaci et al. (2003) and the first experiment described earlier. We compared four conditions: all cues present condition, where all three cues were present in the scene (Fig. 6A); cast shadows only condition (Fig. 6B); shading only condition (Fig. 6C); specular highlights only condition (Fig. 6D).




Fig. 5. Estimates of Punctate Source Direction from Boyaci et al. (2004). For each observer, Boyaci et al. (2004) estimated the minimum of the achromatic setting curves for each subject (Fig. 4) and interpreted these as estimates of punctate light source direction. The true value is plotted as a dashed line, the observer's estimates as solid lines. (A) Azimuth estimates, punctate source at -15° azimuth. (B) Azimuth estimates, punctate source at 15° azimuth. (C.) Elevation estimates.




Fig. 6. Candidate cues to the illuminant. Stimuli used in Boyaci et al. (2006) (A) All cues present, (B) Cast shadows only condition, (C) Shading only condition, (D) Specular highlights only condition.

In Fig. 7 we show the results from one naive observer. This observer corrected for the test patch orientation in all four cue conditions, suggesting that she used each candidate cue in isolation. We also performed a reliability analysis to address to what extent observers combined the cues when all three cues were present (all cues condition). The results of that analysis indicated that the reliability of observers' settings in all cues condition were higher than the best individual cue ("effective cue combination"), nevertheless smaller than the reliability predicted by optimal cue combination rules.

In the next section we describe two experiments where the orientation of the test surface never changes. Instead, there is a strong gradient of illumination in depth within the scene and the test surface is moved from a dimly lit region to a brightly lit region. The only difference between the two experiments is the presence of specular surfaces that serve as candidate illumination cues. A comparison of performance in the two experiments reveals whether observers use these specular cues to estimate the light field.

Lightness perception with changes in location

In indoor scenes, the light field can vary markedly with location as walls serve to block or reflect incident light. The celebrated experiments of Gilchrist (1977, 1980; see also Kardos, 1934) demonstrate that the visual system partially compensates for light fields that vary across space. Gilchrist et al. (1999) proposed that observers segment complex scenes into illumination *frameworks* (Katz, 1935; Gilchrist & Annan, 2002) and discount the illumination (light field) within each framework. 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 color perception (e.g., Bloj et al., 1999; Doerschner et al., 2004).

Ikeda et al. (1998) examined lightness perception in scenes comprising two small rooms arranged in depth with a doorway between them, patterned after Gilchrist (1977). The lighting of the




Fig. 7. Results from Boyaci et al. (2006). One naive observer's data is plotted same as in Fig. 4. (A) All cues present, (B) Cast shadows only condition, (C) Shading only condition, (D) Specular highlights only condition. In all conditions this observer discounted the test patch orientation for her lightness estimates.

rooms was complex, consisting of multiple fluorescent tubes placed above both rooms, and the observer could not see these light sources. The intensity of light incident on a test surface placed along the line of sight through the doorway varied with depth. Ikeda et al. (1998) measured apparent lightness for surfaces at different depths. Their observers viewed a test square placed at several different depths along the line of sight and passing in depth through the center of both rooms. The observers' task was to match the test square to a lightness scale. Ikeda et al. found that observers substantially discounted the actual illumination profile at different depths in the scene.

We next describe two experiments by Snyder et al. (2005) using rendered scenes similar in design to that of Gilchrist (1977) and Ikeda et al. (1998). All scenes were presented binocularly and consisted of two rooms arranged along the line of sight with walls composed of random, achromatic Mondrians. A top view of the simulated scenes is shown in Fig. 8A. The far room was lit by two light sources not visible to the observer. The near room was lit by diffuse light only. The test surface (called a standard surface from this point forward) varied in depth from trial to trial as shown. The observer adjusted an adjustable surface in the near room until the standard and adjustable surfaces seemed to "be cut from the same piece of paper." In Experiment 2, we also added a candidate cue to the spatial distribution of the illumination: 11 specular spheres placed at random in the scene (but never in front of either the standard or adjustable surfaces). The relative luminance of the light (with respect to the back wall of the far room) is plotted in Fig. 8B. It varied by roughly a factor of five from far room to near room. An example of a scene (for Experiment 2) is shown in Fig. 8. The scenes for Experiment 1 were similar but lacked specular spheres.

The results of Snyder et al. (2005) for five subjects, four naïve and one (JLS) an author, are shown in Fig. 10. In both experiments, Snyder et al. estimated the ratio of the luminance of the standard surface to that of the adjustable surface (relative luminance) at each location in the room. If observers were lightness constant, these settings would follow the relative illumination profile in

Fig. 8B, which is re-plotted with the results of each subject in Fig. 10. The horizontal dashed line corresponds to the settings made if the observers are simply matching luminance. The results for Experiment 1 are plotted with hollow circles, those for Experiment 2 with filled circles.

Snyder et al. concluded that all observers significantly discounted the gradient of illumination in depth in both experiments and that their degree of constancy significantly improved with the addition of the specular spheres.

Conclusion

The world in which we live is three-dimensional, and claims about the usefulness of visual information should be based on performance in three-dimensional environments. Many researchers in color vision have limited their scope to the realm we call the Mondrian singularity: experimental conditions that are very different from the world in which we live. Such studies have consequently yielded limited results. A fruitful alternative is to examine human color perception in 3-dimensional scenes that contain cues to the light field.

In this article, we briefly reviewed recent work by researchers evaluating surface color, lightness perception, and constancy in three-dimensional scenes with moderately complex lighting models, and we presented two recent studies in detail. The implication of this research is that the human visual system can compensate for spatially and spectrally inhomogeneous light fields. In the discussion we found that performance is affected by the availability of specular illuminant cues that signal the light field (Snyder et al., 2005, Boyaci et al., 2006) and that are not available in Mondrian scenes.

The results of Snyder et al. are particularly interesting. The stimuli are presented binocularly and, if we view either of the binocular images in isolation, we find that the only change in the stimulus from trial to trial is small shifts to the left and right against an otherwise constant background (Fig. 9). If we attempted to explain the perceived lightness of the test surface in terms of its

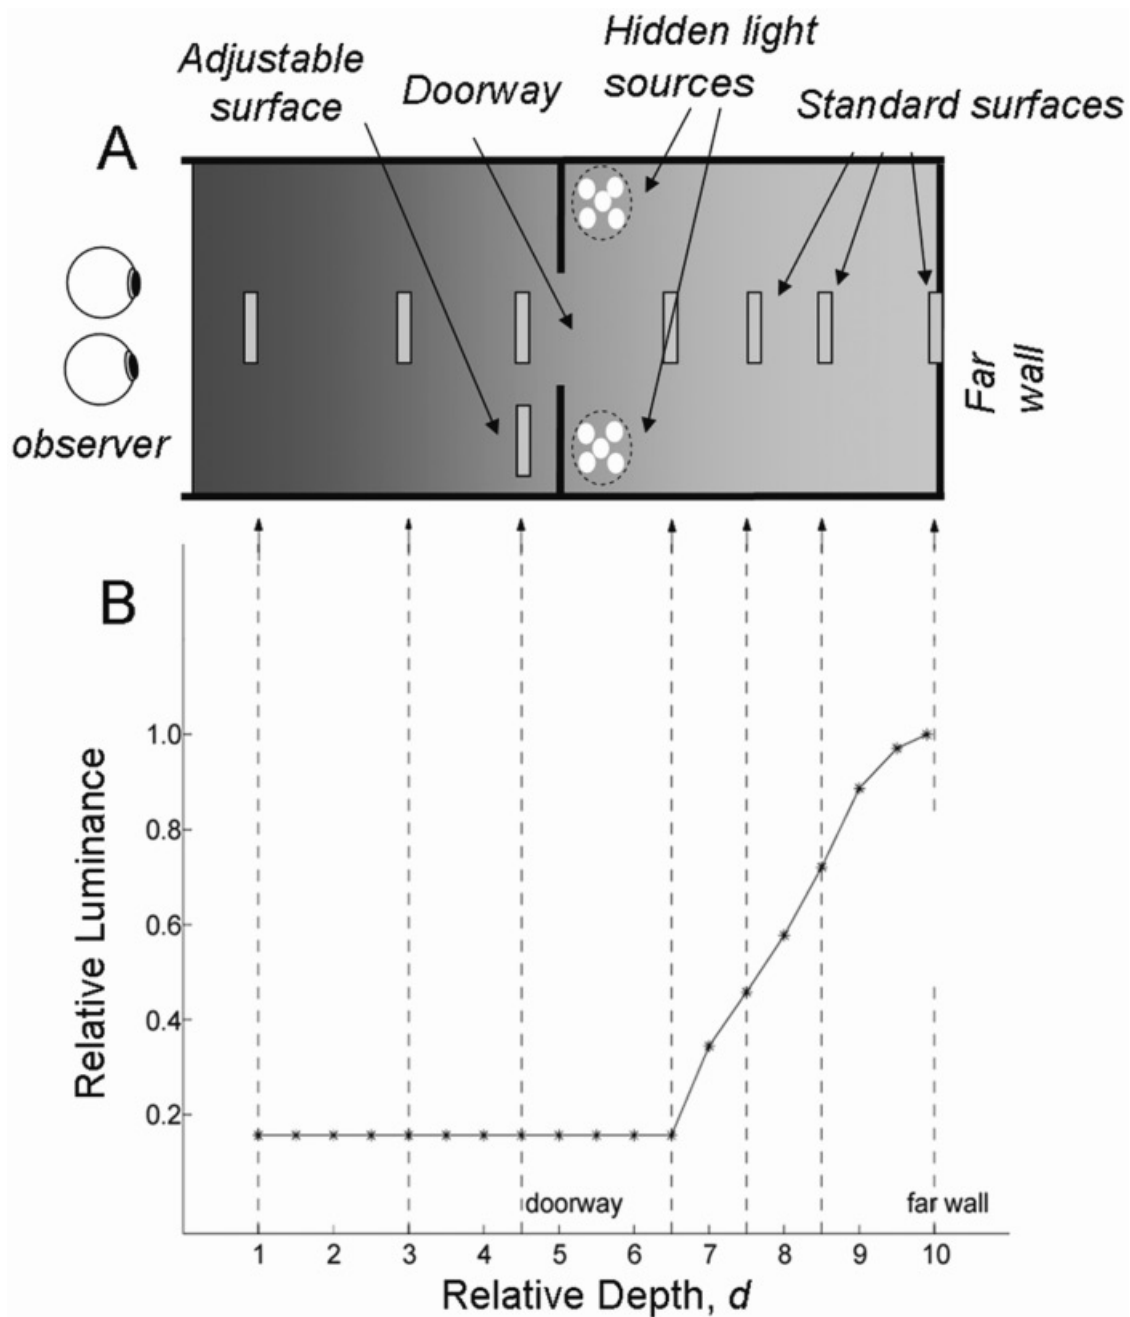


Fig. 8. Schematic of Scenes used in Snyder et al. (2005). (A) A schematic top view of the scenes used in Snyder et al. (2005). (B) The actual relative illumination profile. The intensity of light incident on a matte surface perpendicular to the observer's line of sight as a function of depth.

immediate surround, then we could only predict that there would be little or no trial to trial variation. Yet, we find large changes in perceived lightness (Fig. 10) as a function of depth. The binocular disparity cues that lead to altered perception of lightness are not present in either image alone. These results are consistent with those of Ikeda et al. (1998). Color perception in 3-dimensional scenes cannot readily be predicted given only the results of experiments on Mondrian scenes.

The work described leaves open several major questions. First of all, given that human observers can compensate for the light

field introduced by the relatively simple lighting models (punctate and diffuse, such as those used in the experiments described here), what are the limits to the human ability to discount complex lighting models? In other words, how complex can the illumination in a scene become before human observers can no longer compensate for the effects of spatial and spectral inhomogeneities in the light field? Second, what cues to the spatial and spectral distribution of light in a scene does the human visual system use, and how are they estimated? Reframing the problem of illuminant estimation in terms of combination of veridical cues to the light




Fig. 9. A Scene from Snyder et al. (2005). Observers viewed rendered scenes binocularly. The left two images permit crossed binocular fusion, the right two images, uncrossed. The scenes were rendered with a combination of punctate and diffuse light sources. The punctate sources were in the far room behind the wall containing the doorway. The standard patch was in the center of the scene and moved from trial to trial along the observer’s line of sight in depth. The adjustable patch was next to the doorway on the right. The observer adjusted the luminance of this patch until its lightness (perceived albedo) matched the lightness of the test patch.

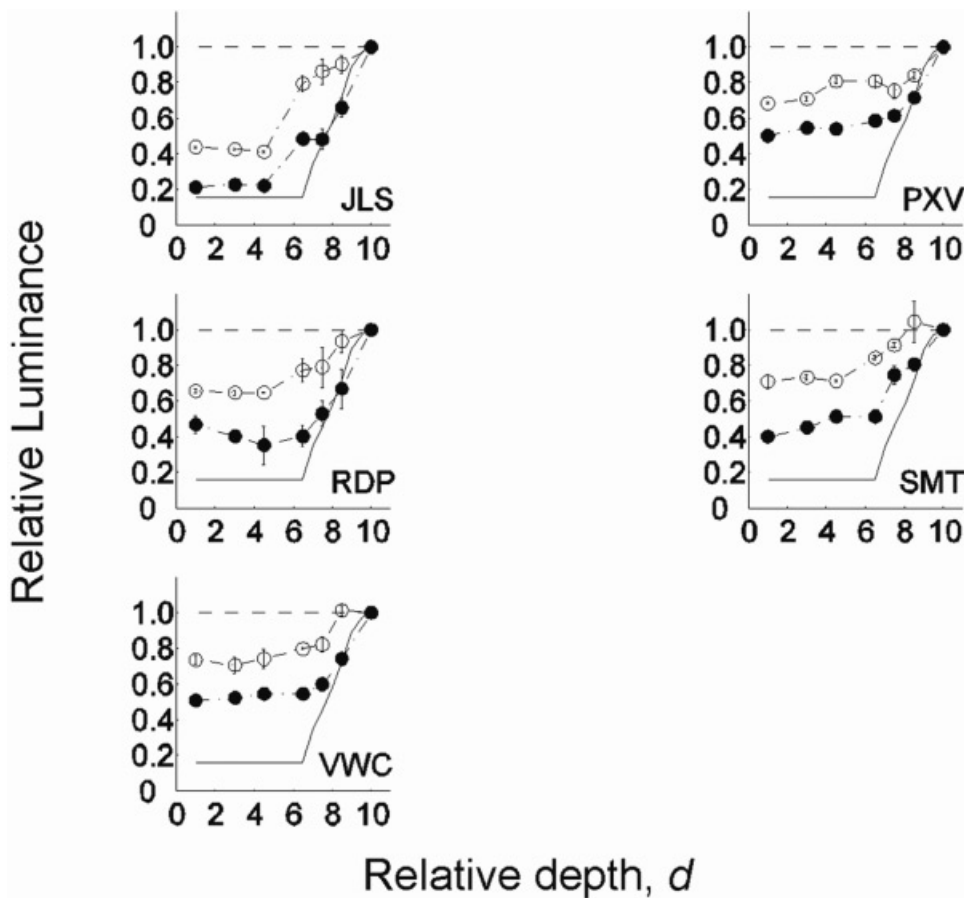


Fig. 10. Results from Snyder et al. (2005). The relative luminance of the observers’ lightness matches is plotted as a function of depth with specular spheres (solid circles) and without (open circles). The actual relative illumination profile is also included as a solid curve. An observer with perfect lightness constancy would have settings on this line. The horizontal line signifies settings for an observer with no lightness constancy (luminance matching). The observers partially discounted the actual gradient of light intensity with and without the specular spheres. With the specular spheres, their performance is markedly closer to that of a lightness constant observer, and this result suggests that the spheres act as cues to spatial variations in the light field.

field opens new and potentially valuable directions for research (Maloney, 1999).

In this review we have focused on surface color perception. It would also be of interest to see how human observers judge color relations (Foster & Nascimento, 1994; Nascimento & Foster, 2000) between surfaces free to differ in orientation and location in the kinds of scenes used in the experiments presented here. Equally, it is of interest to assess how judgments of material properties (Fleming et al., 2003) vary in such scenes, and conversely, how observers' estimates of spatial characteristics of the light field (direction, diffuseness) depend on the material of objects (te Pas & Pont, 2005).

From the broadest perspective, a full description of human color perception requires that we examine how the human visual system operates in fully three-dimensional scenes with adequate cues to the illuminant. Understanding human color perception in the Mondrian singularity remains an extremely important research topic, and work in this area contributes to our understanding of visual perception. The work described here serves to complement and extend the large existing body of literature on surface color perception.

Acknowledgments

This research was funded in part by Grant EY08266 from the National Eye Institute of the National Institutes of Health and a NASA Graduate Student Researcher Fellowship to JLS. This work was presented as an invited address at the 2005 Conference of the International Colour Vision Society, Lyon, France.

References

- ADELSON, E.H. & BERGEN, J.R. (1991). The plenoptic function and the elements of early vision. In *Computational Models of Visual Processing*, eds. LANDY, M.S. & MOVSHON, J.A., pp. 3–20. Cambridge, MA: MIT Press.
- ADELSON, E.H. & PENTLAND, A.P. (1996). The perception of shading and reflectance. In *Perception as Bayesian Inference*, eds. KNILL, D. & RICHARDS, W., pp. 409–423. New York: Cambridge University Press.
- AREND, L.E. & SPEHAR, B. (1993a). Lightness, brightness, and brightness contrast: 1. illuminance variation. *Perception & Psychophysics* **54**, 446–456.
- AREND, L.E. & SPEHAR, B. (1993b). Lightness, brightness, and brightness contrast: 2. reflectance variation. *Perception & Psychophysics* **54**, 457–468.
- BÄUML, K.-H. (1999). Simultaneous color constancy: how surface color perception varies with the illuminant. *Vision Research* **39**, 1531–1550.
- BELHUMEUR, P.N., KRIEGMAN, D. & YUILLE, A. (1999). The bas-relief ambiguity. *International Journal of Computer Vision* **35**(1), 33–44.
- BLOI, M.G., KERSTEN, D. & HURLBERT, A.C. (1999). Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature* **402**, 877–879.
- BOYACI, H., DOERSCHNER, K. & MALONEY, L.T. (2004). Perceived surface color in binocularly-viewed scenes with two light sources differing in chromaticity. *Journal of Vision* **4**, 664–679.
- BOYACI, H., DOERSCHNER, K. & MALONEY, L.T. (2006). Cues to an equivalent lighting model. *Journal of Vision* **6**, 106–118.
- BOYACI, H., MALONEY, L.T. & HERSH, S. (2003). The effect of perceived surface orientation on perceived surface albedo in three-dimensional scenes. *Journal of Vision* **3**, 541–553.
- CIUREA, F. & FUNT, B. (2004). Failure of luminance-redness correlation for illuminant estimation. *Proceedings Twelfth Color Imaging Conference*, pp. 42–46.
- DOERSCHNER, K., BOYACI, H. & MALONEY, L.T. (2004). Human observers compensate for secondary illumination originating in nearby chromatic surfaces. *Journal of Vision* **4**, 92–105.
- DROR, R.O., WILLSKY, A. & ADELSON, E.H. (2004). Statistical characterization of real-world illumination. *Journal of Vision* **4**, 821–837.
- EPSTEIN, W. (1961). Phenomenal orientation and perceived achromatic color. *Journal of Psychology* **52**, 51–53.
- FLEMING, R.W., DROR, R.O. & ADELSON, E.H. (2003). Real-world illumination and the perception of surface reflectance properties. *Journal of Vision* **3**, 347–368.
- FLOCK, H.R. & FREEDBERG, E. (1970). Perceived angle of incidence and achromatic surface color. *Perception & Psychophysics* **8**, 251–256.
- FOSTER, D.H. & NASCIMENTO, S.M.C. (1994). Relational colour constancy from invariant cone-excitation ratios. *Proceedings of the Royal Society of London B* **257**, 115–121.
- GERSHUN, A. (1936/1939). *Svetovoe Pole* (English: The Light Field), Moscow, 1936. Translated by P. Moon and G. Timoshenko (1939) in *Journal of Mathematics and Physics* **18**, 51–151.
- GILCHRIST, A.L. (1977). Perceived lightness depends on spatial arrangement. *Science* **195**, 185–187.
- GILCHRIST, A.L. (1980). When does perceived lightness depend on perceived spatial arrangement? *Perception & Psychophysics* **28**, 527–538.
- GILCHRIST, A.L. & ANNAN, JR., A. (2002). Articulation effects in lightness: historical background and theoretical implications. *Perception* **31**, 141–150.
- GILCHRIST, A.L., KOSSYFIDIS, C., BONATO, F., AGOSTINI, T., CATALIOTTI, J., LI, X.J., SPEHAR, B., ANNAN, V. & ECONOMOU, E. (1999). An anchoring theory of lightness perception. *Psychological Review* **106**, 795–834.
- GOLZ, J. & MACLEOD, D.I.A. (2002). Influence of scene statistics on colour constancy. *Nature* **415**, 637–640.
- GRANZIER, J.J.M., BRENNER, E., CORNELISSEN, F.W. & SMEETS, J.B.J. (2005). Luminance–color correlation is not used to estimate the color of the illumination. *Journal of Vision* **5**, 20–27.
- HARA, K., NISHINO, K. & IKEUCHI, K. (2005). Light source position and reflectance estimation from a single view without the distant illumination assumption. *IEEE Transactions on Pattern and Machine Intelligence* **27**, 493–505.
- HARALICK, R.M. & SHAPIRO, L.G. (1993). *Computer and robot vision*, Vol. 2. Reading, MA: Addison-Wesley.
- HENDERSON, S.T. (1977). *Daylight and its spectrum*, 2nd Ed. Bristol, UK: Adam Hilger.
- HOCHBERG, J.E. & BECK, J. (1954). Apparent spatial arrangements and perceived brightness. *Journal of Experimental Psychology* **47**, 263–266.
- HURLBERT, A.C. (1998). Computational models of colour constancy. In *Perceptual Constancy: Why things look as they do*, eds. WALSH, V. & KULIKOWSKI, J., pp. 283–322. Cambridge, UK: Cambridge University Press.
- IKEDA, M., SHINODA, H. & MIZOKAMI, Y. (1998). Three dimensionality of the recognized visual space of illumination proved by hidden illumination. *Optical Review* **5**, 200–205.
- KAISER, P.K. & BOYNTON, R.M. (1996). *Human color vision*, 2nd Ed. Washington, DC: Optical Society of America.
- KARDOS, L. (1934). Ding und Schatten; Eine experimentelle Untersuchung über die Grundlagen des Farbsehens. *Zeitschrift für Psychologie und Physiologie der Sinnesorgane, Ergänzungsband* **23**, Leipzig, Germany: Verlag von J.A. Barth. (Edited by SCHUMANN, F., JAENSCH, E.R. & KROH, O.).
- KATZ, D. (1935). *The World of Colour*. London: Kegan, Paul, Trench, Trubner & Co.
- KOENDERINK, J.J. & VAN DOORN, A.J. (1996). Illuminance texture due to surface mesostructure. *Journal of the Optical Society of America A* **13**, 452–463.
- KOENDERINK, J.J., VAN DOORN, A.J. & PONT, S.C. (2004). Light direction from shad(ow)ed random Gaussian surfaces. *Perception* **33** (12), 1403–1404, special issue: Shadows and Illumination II.
- LAND, E.H. & MCCANN, J.J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America* **61**, 1–11.
- LEE, JR., R.L. & HERNÁNDEZ-ANDRÉS, J. (2005a). Short-term variability of overcast brightness. *Applied Optics* **44**, 5704–5711.
- LEE, JR., R.L. & HERNÁNDEZ-ANDRÉS, J. (2005b). Colors of the daytime overcast sky. *Applied Optics* **44**, 5712–5722.
- MACLEOD, D.I.A. & GOLZ, J. (2003). A computational analysis of colour constancy. In *Colour Perception: Mind and the Physical World*, eds. MAUSFELD, R. & HEYER, D., pp. 205–242. Oxford, UK: Oxford University Press.
- MALONEY, L.T. (1999). Physics-based approaches to modeling surface color perception. In *Color Vision: From Genes to Perception*, eds.

- GEGENFURTNER, K.R. & SHARPE, L.T., pp. 387–422. Cambridge, UK: Cambridge University Press.
- MALONEY, L.T., BOYACI, H. & DOERSCHNER, K. (2005). Surface color perception as an inverse problem in biological vision. *Proceedings of the SPIE-IS & T Electronic Imaging* **5674**, 15–26.
- MAUSFELD, R. & ANDRES, J. (2002). Second-order statistics of colour codes modulate transformations that effectuate varying degrees of scene invariance and illumination invariance. *Perception* **31**, 209–224.
- NASCIMENTO, S.M.C. & FOSTER, D.H. (2000). Relational color constancy in achromatic and isoluminant images. *Journal of the Optical Society of America A-Optics Image Science and Vision* **17**, 225–231.
- PONT, S.C. & KOENDERINK, J.J. (2003). Illuminance flow. In *Computer Analysis of Images and Patterns*, eds. PETKOV, N. & WETSEBERG, M.A., pp. 90–97. Berlin: Springer-Verlag.
- PONT, S.C. & KOENDERINK, J.J. (2004). Surface illuminance flow. *Proceedings Second International Symposium on 3D Data Processing Visualization and Transmission*, eds. ALOIMONOS, Y. & TAUBIN, G.
- REDDING, G.M. & LESTER, C.F. (1980). Achromatic color matching as a function of apparent test orientation, test and background luminance, and lightness or brightness instructions. *Perception & Psychophysics* **27**, 557–563.
- RIPAMONTI, C., BLOI, M., HAUCK, R., KIRAN, K., GREENWALD, S., MALONEY, S.I. & BRAINARD, D.H. (2004). Measurements of the effect of surface slant on perceived lightness. *Journal of Vision* **4**, 747–763.
- SNYDER, J.L., DOERSCHNER, K. & MALONEY, L.T. (2005). Illumination estimation in three-dimensional scenes with and without specular cues. *Journal of Vision* **5**, 863–877.
- TE PAS, S.F. & PONT, S.C. (2005). Comparison of material and illumination discrimination performance for real rough, real smooth and computer generated smooth spheres. In *Proceedings of the 2nd Symposium on Applied Perception in Graphics and Visualization* (A Coruña, Spain, August 26–28, 2005). APGV '05, vol. 95. ACM Press, New York, NY, 75–81.
- YANG, J.N. & MALONEY, L.T. (2001). Illuminant cues in surface color perception: Tests of three candidate cues. *Vision Research* **41**, 2581–2600.