

Research Article

LIGHTNESS CONSTANCY: A Direct Test of the Illumination-Estimation Hypothesis

M.D. Rutherford and D.H. Brainard

University of California, Santa Barbara

Abstract—Many models of color constancy assume that the visual system estimates the scene illuminant and uses this estimate to determine an object's color appearance. A version of this illumination-estimation hypothesis, in which the illuminant estimate is associated with the explicitly perceived illuminant, was tested. Observers made appearance matches between two experimental chambers. Observers adjusted the illumination in one chamber to match that in the other and then adjusted a test patch in one chamber to match the surface lightness of a patch in the other. The illumination-estimation hypothesis, as formulated here, predicted that after both matches the luminances of the light reflected from the test patches would be identical. The data contradict this prediction. A second experiment showed that manipulating the immediate surround of a test patch can affect perceived lightness without affecting perceived illumination. This finding also falsifies the illumination-estimation hypothesis.

Color vision helps observers identify objects and their properties. Because the light reflected to the eye confounds information about an object's surface properties and the illuminant, color is useful for identification only if visual processing produces a *color-constant* psychological representation that depends primarily on the object. Under many (but not all) conditions, human vision exhibits excellent color constancy. It remains to be understood how color constancy comes about and why it sometimes fails.

The computational problem of how a visual system could achieve color constancy has been examined extensively (see Hurlbert, 1998; Maloney, 1999). The majority of candidate algorithms share a common two-step framework. First, the image data are processed to yield an estimate of the illuminant. Then, this estimate is used to correct the light reflected from each image location to yield a representation that is approximately illuminant independent. The various algorithms differ in exactly how the estimates are obtained.

Algorithms that estimate the illuminant can potentially explain why color constancy is good for some scenes (the algorithm accurately estimates the illuminant) and poor for others (the algorithm inaccurately estimates the illuminant). Attempts to test computational algorithms as theories of human color vision have generally examined predictions derived from specific algorithms (e.g., Kraft & Brainard, 1999; Maloney & Yang, in press; McCann, McKee, & Taylor, 1976; Valberg & Lange-Malecki, 1990). We wanted to step back and ask whether the two-step framework itself provides a reasonable description of how human color vision works. Following Maloney and Yang (in press), we refer to the idea that it does as the *illumination-estimation hypothesis*.¹

Address correspondence to M.D. Rutherford, Department of Psychology, University of Denver, 2155 South Race St., Denver, CO 80208; e-mail: mdrutherford@hotmail.com.

1. Other researchers have referred to this idea as the *albedo hypothesis* (Beck, 1972) or *taking-into-account hypothesis* (Epstein, 1973). We prefer the more descriptive *illumination-estimation hypothesis*.

If this hypothesis is correct, researchers can confidently use the computational literature as a guide to focus experiments on human vision. If it is false, then either a closer theoretical treatment of the link between computation and human performance is required or else the computational work does not provide a useful foundation for understanding human vision.

The illumination-estimation hypothesis has a history considerably longer than that of the computational literature. Several early authors felt that an estimate (perhaps unconscious) of the illuminant should be a central feature in any theory of color constancy (e.g., Katz, 1935; Koffka, 1935; Woodworth, 1938). Over the years, a number of studies have investigated the perception of illumination (e.g., Beck, 1959; Gilchrist & Jacobsen, 1984; Hurlbert, 1989; Kardos, 1928), and a few have explicitly tested the illumination-estimation hypothesis (Beck, 1961; Kozaki & Noguchi, 1976; Logvinenko & Menshikova, 1994; Noguchi & Kozaki, 1985; Oyama, 1968). Nonetheless, the status of the hypothesis remains uncertain.

THEORY

It is generally accepted that color percepts can be described in terms of three underlying perceptual dimensions: lightness, redness-greenness, and blueness-yellowness. In this article, we restrict attention to lightness and consider stimuli with varying luminances but constant chromaticity. Consider scenes consisting of matte, flat coplanar objects that are uniformly and diffusely illuminated. Given the restriction to isochromatic stimuli, object reflectance at each scene location is described by a single number, s , that specifies the proportion of incident illumination reflected to the observer. The scene illuminant is also specified by a single number, i . It is convenient to define i as the luminance of the light that would be reflected by an object having a reflectance of 1. We refer to i as the *luminance of the illuminant*.² This convention allows the luminance of the light reflected to the observer, e , to be expressed by $e = is$.

The illumination-estimation hypothesis is that the visual system forms an estimate of the illuminant, i' , and this estimate is used in processing the light reflected to the eye (e) to determine an estimate of the surface reflectance, s' . Because it is clear that in addition to judging lightness, observers can make perceptual judgments of the illuminant (e.g., Gilchrist, 1988; Hurlbert, 1989; Jameson & Hurvich, 1989; Katz, 1935; Woodworth, 1938; Zaidi, 1998), a natural expression of the illumination-estimation hypothesis is that perceived illumination is determined by i' and perceived surface lightness is determined by s' . This is the form of the hypothesis tested here. An alternative, not tested here, postulates that an illuminant estimate plays a causal role in

2. The use of a single luminance to describe the physical properties of the illuminant is valid for the restricted class of scenes considered here. The theory would have to be elaborated for scenes with more geometrical complexity.




Fig. 1. Schematic of the experimental apparatus. An observer sat between two experimental chambers and viewed them in alternation through a reduction screen. A rotating mirror controlled which chamber was visible at any given time. The aperture in the reduction screen was shuttered, so that the rotation of the mirror was not visible to the observer. The observer's view of the two chambers is shown in Figure 2.

surface perception, but this estimate is dissociated from the perceived illuminant (see Discussion).

A simple form of the illuminant-estimation hypothesis postulates (e.g., Beck, 1972) that i' and s' are linked to the stimulus through the equation

$$s' = e/i'. \quad (1)$$

Equation 1 has been tested empirically and rejected (Kozaki & Noguchi, 1976; Logvinenko & Menshikova, 1994; Noguchi & Kozaki, 1985; Oyama, 1968). However, it expresses just one of many possible ways that perceived illuminant could influence perceived reflectance. A more general formulation is

$$s' = f(e, i'), \quad (2)$$

where for any i' , the relation between s' and e is one-to-one. Equation 2 is the expression of the illumination-estimation hypothesis considered in this article.³ Suppose that we have two scenes, A and B, for which the perceived illumination is identical: $i'_A = i'_B$. Equation 2 predicts that if we identify surfaces in the two scenes that have the same perceived lightness, the light reflected from these surfaces to the eye must be the same: $e_A = e_B$.

A thought experiment suggests conditions in which the illumination-estimation hypothesis is likely to fail. Consider the classic simultaneous-contrast display, in which a small test region of fixed luminance has a different apparent lightness depending on whether it is seen against a high-luminance or low-luminance background. If, as introspection suggests, observers judge the illumination to be percep-

tually uniform across the two backgrounds, then this effect would falsify the illumination-estimation hypothesis. The experiments described here exploited this insight to provide a quantitative test of the illumination-estimation hypothesis.

EXPERIMENT 1

Method

Overview

Observers viewed two experimental chambers in alternation. On each trial, the observer first adjusted the illumination in one chamber so that it appeared to be the same as the illumination in the other. After equating the perceived illuminants, the observer adjusted the luminance of a test patch in one chamber so that its surface lightness matched that of a corresponding patch in the other chamber. Thus, at the end of each trial, both the perceived illuminant and the perceived surface lightness of one location matched across the two chambers. The illumination-estimation hypothesis predicted that after the matches, the light reaching the eye from the matched locations of the two chambers would be identical.

Apparatus

The experimental apparatus is shown schematically in Figure 1. The surfaces in the standard chamber were of high or medium reflectance. The surfaces in the match chamber were of medium or low reflectance. A computer-controlled mirror allowed observers to view the chambers in alternation. Observers viewed one chamber at a time monocularly (right eye) through a black reduction screen. A chin rest sta-

3. See Foley's (1972) treatment of the size-distance-invariance hypothesis.

Lightness Constancy

bilized head position. The aperture in the screen was 14 1/4 in. from the observer's eye and subtended a visual angle of 18° (horizontal) by 14° (vertical). The chambers were viewed for 4 s each in continuous alternation, with a 1-s interval between views. During this interval, the aperture was occluded.

Each chamber was 36 in. deep by 31 in. wide. The ceiling, which was out of view of the observer, was 35 1/2 in. above the floor and comprised two layers of diffuser paper. The chambers were lined with matte cardboard that could be changed by the experimenter in order to change reflectance. In the rear third of each chamber was an array of objects, creating a naturalistic scene. Figure 2 shows what each chamber looked like from the observer's point of view.

The illumination in each chamber was produced from above by a bank of six computer-controlled stage lamps. Individual lamps had either red, green, or blue filters. By varying the lamp intensities, it was possible to change the luminance of the illuminant while holding its chromaticity constant. Each chamber also contained a test patch located on its back wall (see Fig. 2). This patch subtended a visual angle of 0.88° (horizontal) by 1.50° (vertical) and consisted of an LCD flat-panel monitor diffused by a gray gelatin filter and plastic diffuser. The patch looked like a surface, but its apparent reflectance could be adjusted by varying the monitor's luminance. Calibration details are available in Rutherford (2000).

General procedure

On each trial, the luminance of the illuminant and simulated reflectance of the test patch in the standard chamber were set at one of four levels (illuminant: 8, 19, 30, or 40 cd/m²; reflectance: .38, .46, .53, or .60). All levels of illuminant were paired with all levels of reflectance. On each trial, observers first adjusted the illumination in the match chamber so that it appeared the same as the illumination in the stan-

dard chamber. They were instructed to match the overall illuminations in the chambers. Next, observers adjusted the luminance of the test patch in the match chamber so that it appeared to "be cut out of the same piece of paper" (Arend & Reeves, 1986) as the test patch in the standard chamber. Both adjustments were accomplished using a joystick. After setting each match, observers indicated whether it was satisfactory. They were sometimes unable to obtain satisfactory matches. In most cases when this occurred, the best match was at the high or low limit of the range that could be produced by the apparatus. Data were analyzed only for trials on which observers were satisfied with both the illuminant and the surface matches.

At the end of each session, the settings for each match were replayed, and a spectroradiometer, placed at the observer's position, was used to measure the stimuli. Illuminants were assessed by measuring the light reflected from a reflectance standard at the location of the test patch. The surface reflectance of each match was obtained by dividing the luminance from the test patch by the luminance of the illuminant.

Observers

There were 7 observers: 2 females in their early 20s, 1 female in her early 30s, and 4 males in their early 20s. All observers except the first author (M.D.R.) were naive and were paid \$10 for participating in each session. Each observer participated in two sessions.

Preliminary experiments

Preliminary experiments were conducted under conditions in which the standard and match chambers contained identical objects. The results verified that (a) observers could set veridical illuminant and surface matches and (b) there was no asymmetry between the two chambers (Rutherford, 2000).



Fig. 2. Stimuli for Experiment 1. This figure shows each chamber from the observer's point of view when the two chambers were illuminated identically. The surfaces in the standard chamber (left) were all of high or medium reflectance (maximum reflectance .82, minimum reflectance .28). The surfaces in the match chamber (right) were all of medium or low reflectance (maximum reflectance .30, minimum reflectance .02). The spatial layout of the objects was nearly identical in the two chambers except for a left-right reversal. The small rectangular patch visible in the back wall of each chamber is the test patch. The monitors that served as the test patches were turned off when the pictures were taken. These regions have been further darkened here, using an image-processing program, to make the locations of the patches more salient.

Results

Illuminant matches

The data obtained for 1 observer are shown in the left panel of Figure 3. The data would lie on the dashed diagonal if the matched illuminant were physically the same as the standard illuminant. The data are above the line, indicating that the change in the reflectance range across the two chambers induced a bias in the illuminant matches.

The raw data may be summarized with the slope of the best-fitting line through the data. The summary slopes for all 7 observers are shown in the right panel of Figure 3. All observers showed a bias in their illuminant matches; the average slope was 1.84. The average difference between the standard and match illuminant for the 7 observers was significantly different from zero, $t(6) = 10.16$, $p < .0001$, two-tailed.

Surface matches

Surface-match data for 1 observer are shown in the left panel of Figure 4. The data would lie on the dashed diagonal if the matches were veridical. The data are all below the line; this observer does not show perfect color constancy.

Surface-match slopes for the 7 observers are shown in the right panel of Figure 4. The average slope for all observers was 0.38. The average difference between the standard and the match reflectance for the 7 observers was significantly different from zero, $t(6) = 24.88$, $p < .0001$, two-tailed.

Surface matches expressed as proximal luminance

The critical test of the illumination-estimation hypothesis is whether the physical luminance of the light reaching the eye from the test patch

in the two chambers was the same after both the illuminant and the surface reflectance were perceptually matched. All luminance data for 1 observer are shown in the left panel of Figure 5. The data would lie on the dashed diagonal if the physical luminances measured at the test patches in the two chambers were the same. Most of the data are below the line. All luminance slopes for the 7 observers are shown in the right panel of Figure 5. The average slope for all observers was 0.67. The average difference between the standard and match luminance for each observer was significantly different from zero, $t(6) = 4.82$, $p < .005$, two-tailed. The data falsify the illumination-estimation hypothesis as formulated earlier.

EXPERIMENT 2

Experiment 2 was designed to test whether reflectance matches can be manipulated independently of illuminant matches. Following Equation 2, if there are experimental manipulations that affect perceived surface reflectance without affecting perceived illumination, it cannot be the case that perceived illumination uniquely determines perceived reflectance for a given proximal stimulus.

Method

There were two differences between Experiments 1 and 2. First, the gray and black cardboard halves of the back wall were swapped in the match chamber, changing the immediate surround of the test patch without altering the range of luminances in the overall image. Second, a different set of test reflectances (.17, .31, .46, and .60) was chosen to keep the range of observers' matches within the possibilities provided by the apparatus. Otherwise, the same procedures were used. The same observers from Experiment 1 participated in Experiment 2.




Fig. 3. Illuminant-matching data from Experiment 1. The left panel shows the illuminant matches for 1 observer. The dashed line indicates the diagonal, on which the data would lie if the matched illuminant was physically the same as the standard illuminant. The solid line shows the slope for this observer. The right panel shows the slopes of the illuminant matches for all the observers.




Fig. 4. Surface-matching data from Experiment 1. The left panel shows the surface matches for 1 observer. The dashed line indicates the diagonal, on which the data would lie if the matched reflectance was the same as the standard reflectance. The solid line shows the slope for this observer. The right panel shows the slopes of the surface matches for all the observers.

Results

All illuminant-match slopes for the 7 observers are shown in the left panel of Figure 6. The average slope for all observers was 1.86, not significantly different from the average slope of 1.84 obtained in

Experiment 1, $t(6) = 0.146$, n.s., paired two-tailed t test. In contrast, the surface matches were affected by the manipulation. The right panel of Figure 6 shows the slopes obtained when the surface matches were expressed in terms of reflectance. The average slope for all ob-




Fig. 5. Surface-matching data in terms of proximal luminance. The left panel shows the surface matches for 1 observer for Experiment 1. The illumination-estimation hypothesis predicts that the matched illuminance will be physically the same as the standard luminance, so that the data would lie along the positive diagonal (dashed line). The solid line shows the slope for this observer. The right panel shows the slopes of the surface matches, expressed in terms of proximal luminance, for all the observers. The illumination-estimation hypothesis predicts that these slopes should be 1.

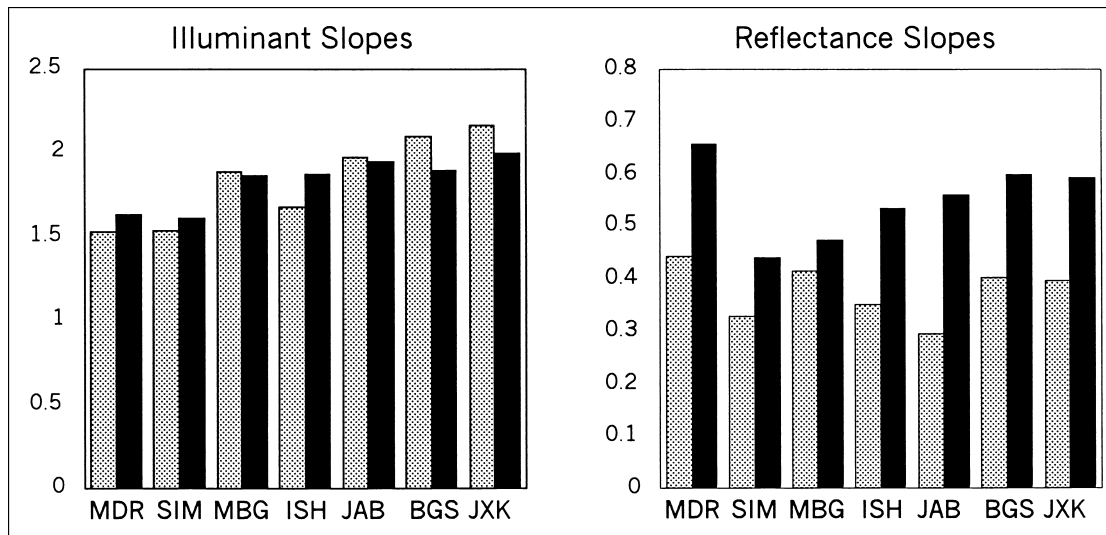


Fig. 6. Results from Experiment 2. In the left panel, the black bars show the slopes of the illuminant matches for all the observers. In the right panel, the black bars show the slopes of the surface matches, expressed in terms of reflectance, for all the observers. In both panels, the corresponding slopes obtained in Experiment 1 are shown by gray bars.

servers was 0.55 in Experiment 2, significantly different from the average of 0.38 obtained in Experiment 1, $t(6) = 6.79$, $p < .001$, paired two-tailed t test.

Discussion

Comparison of the data from Experiments 1 and 2 shows that changing the surround of the test patch in the match chamber affected surface matches without affecting illuminant matches. This result falsifies the illumination-estimation hypothesis, because this hypothesis does not allow the relationship between proximal stimulus and perceived reflectance to change under conditions in which perceived illumination is constant.

Note that the results of Experiment 2, taken in isolation, do not falsify the illumination-estimation hypothesis. Given that it is possible to manipulate perceived surface lightness independently of perceived illumination, it is not surprising that it is possible to find stimulus configurations in which the data are consistent with the illumination-estimation hypothesis. We do not know how common such configurations are for natural viewing.

GENERAL DISCUSSION

Effect of Instructions and Task

In these experiments, observers were instructed to make an overall illumination judgment. This instruction seemed sensible because the chambers appeared uniformly illuminated (to us). To control for the possibility that a judgment made about the illumination incident on the test patch would be different from an overall judgment of illumination in the chamber, we recruited 4 additional naive observers and repeated the experiments, asking observers to judge the illumination at

the test patch. The results were essentially identical to those in the main experiments (Rutherford, 2000).

We also wondered whether the ordering of judgments in our experiment might have influenced the results. To investigate this possibility, we added a second joystick to the apparatus, and observers were allowed to adjust illuminant and surface lightness simultaneously, with instructions to interleave the two. The results were indistinguishable from those in the main experiments (Rutherford, 2000). In this replication, the roles of illuminant and surface judgments were symmetric, so the experiment addresses variants of the illumination-estimation hypothesis that tightly couple perceived illumination and perceived surface lightness but do not assign a causal role to the perception of the illuminant (see Brainard & Freeman, 1997; Koffka, 1935).

Our conclusions rest on the assumption that perceptual matching adequately measures perceived illumination and surface lightness. In space perception, the same physical quantity has been found to have multiple psychological representations (e.g., Bhalla & Proffitt, 1999; but see Philbeck & Loomis, 1996). Only a few studies have examined this question in surface-color perception, and the results have been mixed (see Arend & Reeves, 1986; Gilchrist et al., 1999, p. 826; Speigle & Brainard, 1996a, 1999). We are unaware of any research that has investigated whether different methods of measuring illumination perception lead to consistent conclusions.

There is a difference in the information provided by illuminant matching versus direct judgments of whether the illuminant is the same at two locations. The explicitly perceived illuminant may be a low-resolution readout of a high-precision representation that is available for computing surface lightness. The illumination-estimation hypothesis cannot be rejected via a simple judgment of whether the illuminant is the same at the two locations—affirmative judgments might reflect a lack of precision in readout of the underlying representation rather than identity of the perceived illuminant. A matching paradigm, in contrast, provides information about the precision of the representation through the variability of the matches.

Constancy

Our observers did not show good constancy, either for illumination or for surface matches. However, these results are not inconsistent with the good constancy revealed in other experiments. We manipulated the range of surface reflectances between the two chambers, reducing the validity of a number of cues often taken to support constancy, so it is not surprising that constancy was poor for our conditions. The degree of constancy found empirically depends critically on the scene manipulations employed (Kraft & Brainard, 1999).

Poor constancy per se is not inconsistent with the illumination-estimation hypothesis. Indeed, the variation in constancy across viewing conditions may be explained within the context of the hypothesis by supposing that what varies across the conditions is how accurately the visual system estimates the illuminant. This modeling approach has shown some promise for understanding human color constancy (e.g., Brainard, Brunt, & Speigle, 1997; Brainard, Kraft, & Longère, in press; Maloney & Yang, in press; also Speigle & Brainard, 1996b).

One can move beyond the question of constancy by examining hypotheses that might explain the observed perceptual matches. We considered whether the data from Experiment 1 could be explained by any of the four following simple hypotheses: that observers matched (a) the highest image luminance, (b) the lowest image luminance, (c) the mean image luminance, or (d) the standard deviation of the image luminances. The obtained illuminant-matching data were inconsistent with all four of these hypotheses.⁴

The surface-matching data differed significantly between the two experiments. Because the only difference between the two experiments was that the locations of the background surfaces in the match chamber were interchanged, the surface matches cannot be explained by any theory that does not take spatial arrangement into account. In addition, the data of Experiment 2 falsify the idea that the lightness of a surface is determined solely by the contrast between that surface and its immediate surround. The standard and match surfaces were surrounded by background surfaces of the same reflectance, so if local contrast determined lightness, the reflectance slopes plotted in Figure 6 would be 1.

The Illumination-Estimation Hypothesis

Previous studies of the relation between perceived illumination and perceived lightness have come to mixed conclusions (Beck, 1961, 1972; Kozaki & Noguchi, 1976; Logvinenko & Menshikova, 1994; Noguchi & Kozaki, 1985; Oyama, 1968). The differences in the results of these studies may be due to differences in stimulus conditions or methodology. Our study employed spatially rich three-dimensional scenes, had observers match both illuminants and surfaces on every trial, and was tailored to provide a sharp test of the illumination-estimation hypothesis in its general form.

Our results falsify the idea that perceived illumination is the sole variable that governs the transformation of the luminance of the light reflected to the eye into perceived surface lightness, at least as assessed by matching. Our results challenge the straightforward application of computational theories of color constancy as process models of

human color vision, at least if one wishes to associate computed illuminant estimates with the explicitly perceived illumination.

Our results do not contradict the view that the perception of illumination is important (e.g., Jameson & Hurvich, 1989; Zaidi, 1998) or that it is involved with the perception of surface lightness (e.g., Gilchrist, 1988; Gilchrist et al., 1999; Helmholtz, 1866; Katz, 1935; Koffka, 1935; Schirillo & Shevell, 2000; Woodworth, 1938). Rather, they imply that if an explanatory role is assigned to the explicitly perceived illuminant, one or more additional variables must be taken into account. Our results suggest that the local surround of the surface being judged is one such variable. Thus, one could consider a model of the general form

$$s' = f(e, i', l), \quad (3)$$

where the variable l characterizes the local surround. If a model consistent with Equation 3 accounts for performance across a wide range of viewing conditions, then it would be sensible to return to computational theory for guidance about how i' is related to the image and to use i' as an explanatory variable in the prediction of surface lightness. If, however, each further experiment requires the addition of another explanatory variable, then it would probably be best to develop separate theories of how perceived illumination and perceived lightness are related to the visual stimulus, an approach advocated by Beck (1972). Along these lines, some recent theories of surface appearance (e.g., Adelson, 1999; Gilchrist et al., 1999) employ neither perceived nor estimated illuminants in an explanatory capacity. Note, however, that Gilchrist et al. (1999) regarded the absence of perceived illuminant as a gap in their theory (p. 829).

Another alternative is to retain the use of illuminant estimates to explain surface appearance but to disavow the link between these estimates and the explicitly perceived illumination. Helmholtz (1866) suggested that color constancy does not involve conscious perception of the illuminant but that an illuminant estimate is implicitly registered. As Gilchrist et al. (1999) emphasized, the challenge for this modeling approach is to develop a computationally based theory that can account both for cases in which constancy is good and for cases in which it is not.

Finally, one could attempt to understand our results while preserving the elegance of the illumination-estimation hypothesis by replacing the physical luminance e in Equation 2 with a perceived quantity e' , which might be taken to represent the perceived luminance reaching the eye. This general approach was taken by Gogel (1990) in his theory of space perception. We have reservations about this approach, however, because it dissociates the theory from the physical stimulus and seems not to facilitate a functional description of human performance.

Acknowledgments—The experiments reported here formed part of the first author's dissertation research. We thank E. Adelson, J. Foley, A. Gilchrist, J. Kraft, J. Loomis, Z. Pizlo, and three anonymous reviewers for useful discussions or comments. Support was provided by National Institutes of Health Grant EY 10016.

REFERENCES

- Adelson, E.H. (1999). Lightness perception and lightness illusions. In M. Gazzaniga (Ed.), *The new cognitive neurosciences* (2nd ed., pp. 339–351). Cambridge, MA: MIT Press.

4. Assessment of image statistics was based on radiometric measurements (Photo Research PR-650) or calibrated monochromatic images (Photometrics PXL camera with 500-nm, 550-nm, and 600-nm filters).

- Arend, L.E., & Reeves, A. (1986). Simultaneous color constancy. *Journal of the Optical Society of America A*, 3, 1743–1751.
- Beck, J. (1959). Stimulus correlates for the judged illumination of a surface. *Journal of Experimental Psychology*, 58, 267–274.
- Beck, J. (1961). Judgments of surface illumination and lightness. *Journal of Experimental Psychology*, 61, 368–373.
- Beck, J. (1972). *Surface color perception*. Ithaca, NY: Cornell University Press.
- Bhalla, M., & Proffitt, D.R. (1999). Visual-motor recalibration in geographical slant perception. *Journal of Experimental Psychology: Human Perception and Performance*, 25, 1076–1096.
- Brainard, D.H., Brunt, W.A., & 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., & Freeman, W.T. (1997). Bayesian color constancy. *Journal of the Optical Society of America A*, 14, 1393–1411.
- Brainard, D.H., Kraft, J.M., & Longère, P. (in press). Color constancy: Developing empirical tests of computational models. In R. Mausfeld & D. Heyer (Eds.), *Colour perception: From light to object*. Oxford, England: Oxford University Press.
- Epstein, W. (1973). The process of “taking-into-account” in visual perception. *Perception*, 2, 267–285.
- Foley, J.M. (1972). The size-distance relation and intrinsic geometry of visual space: Implications for processing. *Vision Research*, 12, 323–332.
- Gilchrist, A., & Jacobsen, A. (1984). Perception of lightness and illumination in a world of one reflectance. *Perception*, 13, 5–19.
- Gilchrist, A., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., Spehar, B., Annan, V., & Economou, E. (1999). An anchoring theory of lightness perception. *Psychological Review*, 106, 795–834.
- Gilchrist, A.L. (1988). Lightness contrast and failures of constancy: A common explanation. *Perception & Psychophysics*, 43, 415–424.
- Gogel, W.C. (1990). A theory of phenomenal geometry and its applications. *Perception & Psychophysics*, 48, 105–123.
- Helmholtz, H., von. (1866). *Treatise on physiological optics* (2nd ed., J.P.C. Southall, Trans.). New York: Dover.
- Hurlbert, A.C. (1989). *Cues to the color of the illuminant*. Unpublished doctoral dissertation, Massachusetts Institute of Technology, Cambridge, MA.
- Hurlbert, A.C. (1998). Computational models of color constancy. In V. Walsh & J. Kulikowski (Eds.), *Perceptual constancy: Why things look as they do* (pp. 283–322). Cambridge, England: Cambridge University Press.
- Jameson, D., & Hurvich, L. (1989). Essay concerning color constancy. *Annual Review of Psychology*, 40, 1–22.
- Kardos, L. (1928). Dingfarbenwahrnehmung und Duplizitätstheorie. *Zeitschrift für Psychologie*, 108, 240–314.
- Katz, D. (1935). *The world of colour* (R.B. MacLeod & C.W. Fox, Trans.). London: Kegan, Paul, Trench Truber & Co.
- Koffka, K. (1935). *Principles of Gestalt psychology*. New York: Harcourt, Brace and Company.
- Kozaki, A., & Noguchi, K. (1976). The relationship between perceived surface-lightness and perceived illumination. *Psychological Research*, 39, 1–16.
- Kraft, J.M., & Brainard, D.H. (1999). Mechanisms of color constancy under nearly natural viewing. *Proceedings of the National Academy of Sciences, USA*, 96, 307–312.
- Logvinenko, A., & Menshikova, G. (1994). Trade-off between achromatic colour and perceived illumination as revealed by the use of pseudoscopic inversion of apparent depth. *Perception*, 23, 1007–1023.
- Maloney, L.T. (1999). Physics-based approaches to modeling surface color perception. In K.T. Gegenfurtner & L.T. Sharpe (Eds.), *Color vision: From genes to perception* (pp. 387–416). Cambridge, England: Cambridge University Press.
- Maloney, L.T., & Yang, J.N. (in press). The illuminant estimation hypothesis and surface color perception. In R. Mausfeld & D. Heyer (Eds.), *Colour perception: From light to object*. Oxford, England: Oxford University Press.
- McCann, J.J., McKee, S.P., & Taylor, T.H. (1976). Quantitative studies in retinex theory: A comparison between theoretical predictions and observer responses to the ‘Color Mondrian’ experiments. *Vision Research*, 16, 445–458.
- Noguchi, K., & Kozaki, A. (1985). Perceptual scission of surface lightness and illumination: An examination of the Gelb effect. *Psychological Research*, 47, 19–25.
- Oyama, T. (1968). Stimulus determinants of brightness constancy and the perception of illumination. *Japanese Psychological Research*, 10, 146–155.
- Philbeck, J.W., & Loomis, J.M. (1996). A comparison of two indicators of perceived egocentric distance under full-cue and reduced-cue conditions. *Journal of Experimental Psychology: Human Perception and Performance*, 23, 72–85.
- Rutherford, M.D. (2000). *The role of illumination perception in color constancy*. Unpublished doctoral dissertation, University of California, Santa Barbara. Available at <http://color.psych.ucsb.edu/rutherford/thesis.pdf>
- Schirillo, J.A., & Shevell, S.K. (2000). Role of perceptual organization in chromatic induction. *Journal of the Optical Society of America A*, 17, 244–254.
- Speigle, J.M., & Brainard, D.H. (1996a, November). *Is color constancy task independent?* Paper presented at the Society for Imaging Science and Technology/Society for Information Display Color Imaging Conference: Color Science, Systems, and Applications, Scottsdale, AZ.
- Speigle, J.M., & Brainard, D.H. (1996b). Luminosity thresholds: Effects of test chromaticity and ambient illumination. *Journal of the Optical Society of America A*, 13, 436–451.
- Speigle, J.M., & Brainard, D.H. (1999). Predicting color from gray: The relationship between achromatic adjustment and asymmetric matching. *Journal of the Optical Society of America A*, 16, 2370–2376.
- Valberg, A., & Lange-Malecki, B. (1990). “Color constancy” in Mondrian patterns: A partial cancellation of physical chromaticity shifts by simultaneous contrast. *Vision Research*, 30, 371–380.
- Woodworth, R.S. (1938). *Experimental psychology*. London: Methuen.
- Zaidi, Q. (1998). Identification of illuminant and object colors: Heuristic-based algorithms. *Journal of the Optical Society of America*, 15, 1767–1776.

(RECEIVED 8/9/00; REVISION ACCEPTED 6/5/01)