

Surface color perception and equivalent illumination models

David H. Brainard

Department of Psychology, University of Pennsylvania,
Pennsylvania, PA, USA



Laurence T. Maloney

Department of Psychology, Center for Neural Science,
New York University, USA



Vision provides information about the properties and identity of objects. The ease with which we perceive object properties belies the difficulty of the underlying information-processing task. In the case of object color, retinal information about object reflectance is confounded with information about the illumination as well as about the object's shape and pose. There is no obvious rule that allows transformation of the retinal image to a color representation that depends primarily on object surface reflectance. Under many circumstances, however, object color appearance is remarkably stable across scenes in which the object is viewed. Here, we review a line of experiments and theory that aim to understand how the visual system stabilizes object color appearance. Our emphasis is on models derived from explicit analysis of the computational problem of estimating the physical properties of illuminants and surfaces from the retinal image, and experiments that test these models. We argue that this approach has considerable promise for allowing generalization from simplified laboratory experiments to richer scenes that more closely approximate natural viewing. We discuss the relation between the work we review and other theoretical approaches available in the literature.

Keywords: object surface perception, surface color perception, surface lightness perception, equivalent illumination model, color constancy, lightness constancy, computational models

Citation: Brainard, D. H., & Maloney, L. T. (2011). Surface color perception and equivalent illumination models. *Journal of Vision*, 11(5):1, 1–18, <http://www.journalofvision.org/content/11/5/1>, doi:10.1167/11.5.1.

Introduction

Vision provides information about the properties and identity of objects—a glance at [Figure 1](#), for example, leads effortlessly to the impression of a blue mug. The ease with which we make this judgment, however, belies the difficulty of the information-processing task that accomplishes it. In the case of the mug's color, examination of the image pixel by pixel reveals enormous variation in the spectral properties of the reflected light across the mug's surface. This variation occurs because the reflected light depends not only on the intrinsic reflectance of the mug but also on the geometric properties of the illumination, the object's shape, and the viewpoint of the observer. Because of these factors, there is no obvious rule that allows transformation of the left and right eye retinal images of an object to a stable perceptual representation of its surface reflectance properties. More generally, object recognition is difficult precisely because the same object or class of objects corresponds to a vast number of images. Processing the retinal images to produce representations that are *invariant* to object-extrinsic variation is a fundamental task of visual information processing (DiCarlo & Cox, 2007; Rust & Stocker, 2010; Ullman, 1989).

The same considerations that make object perception a difficult information-processing problem also make it

difficult to study. There are simply too many image configurations to allow direct measurement of what we see for every possible image. To make progress, it is necessary to choose a restricted domain that enables laboratory study, but to formulate theories that can be readily generalized to more complex domains. In this paper, we review work from our laboratories that illustrates how computational models clarify the perception of object surface properties.

The rest of this review is organized as follows. First, we introduce a simplified description of a subset of possible scenes, which we refer to as *flat–matte–diffuse* conditions, and specify the computational problem of recovering descriptions of illuminant and surface properties from photoreceptor responses for these scenes. Although such recovery cannot be performed with complete accuracy, we show how the structure of the computational analysis leads to a class of candidate models for human performance. We call these *equivalent illumination models* (EIMs). We show that such models provide a compact description of how changing the spectrum of the illumination affects perceived color appearance for human observers.

We then turn to more complex viewing conditions, ones that include variations in depth, and show that the same principles that allowed model development for flat–matte–diffuse conditions also allow generation of models for this case. Indeed, the major theme of this article is that the principles underlying equivalent illumination models provide




Figure 1. Image formation. Individual locations on a coffee mug can reflect very different spectra. The uniform patches at the top of the figure show image values from the locations indicated by the arrows. The mug itself, however, is readily perceived as homogeneous blue. After Figure 1 of Xiao and Brainard (2008).

a fruitful recipe for generalization as we extend our measurements and analysis to progressively more realistic viewing conditions.

Finally, we close with a more general discussion of the relation between the computational approach developed here and approaches based on consideration of evolution and visual mechanism.

Flat–matte–diffuse

The most commonly employed laboratory model for studying object surface perception consists of a collection of flat–matte objects arranged in a plane, illuminated by spatially diffuse illumination (Figure 2; Arend & Reeves, 1986; Brainard, 1998; Brainard & Wandell, 1992; Helson & Jeffers, 1940; McCann, McKee, & Taylor, 1976). We refer to such stimulus configurations as flat–matte–diffuse.

Under these viewing conditions, a simple imaging model describes the interaction between object surface reflectance, the illuminant, and the retinal image. A matte surface’s reflectance is characterized by its surface reflectance function, $S(\lambda)$, which specifies the fraction of incident illumination reflected at each wavelength in the visible spectrum. The spatially diffuse illuminant is characterized by its spectral power distribution, $E(\lambda)$, which specifies the illuminant power at each wavelength.

The spectrum of the light reflected from the surface, $C(\lambda)$, is given as the wavelength-by-wavelength product:

$$C(\lambda) = E(\lambda)S(\lambda). \quad (1)$$

$C(\lambda)$ is proportional to the light that reaches the human retina and is called the *color signal* (Buchsbaum, 1980). For our purposes, we can set the constant of proportionality to 1. The color signal is encoded by the excitations of three classes of cone photoreceptors present in a trichromatic human retina. These are the L, M, and S cones and we denote the excitation of cones in each class by the symbol ρ_k , $k = 1, 2, 3$, where the subscript k indexes cone class. The excitation of a cone is computed from the color signal as

$$\rho_k = \int C(\lambda)R_k(\lambda)d\lambda, \quad k = 1, 2, 3, \quad (2)$$

where $R_k(\lambda)$ is the spectral sensitivity of the k^{th} cone class. For a typical trichromatic human observer, k ranges from one to three. The *cone excitation vector* $\rho = (\rho_1, \rho_2, \rho_3)$




Figure 2. Flat–matte–diffuse conditions. A collection of flat–matte surfaces is illuminated by a diffuse light source with spectral power distribution $E(\lambda)$. The color signal $C^j(\lambda)$ reflected from the j th surface is given by $C^j(\lambda) = E(\lambda)S^j(\lambda)$, where $S^j(\lambda)$ is the surface reflectance function of that surface. The cone excitation vector elicited by $C^j(\lambda)$ is ρ^j . The set of ρ^j across the retinal image, $\{\rho^1, \dots, \rho^n\}$, is the information available to computational algorithms that seek to estimate illuminant and surface properties under flat–matte–diffuse conditions. It is also the information available to the human visual system for producing the perceptual representation that is color appearance. When we consider the properties (either physical or perceptual) of a specific surface (denoted the j th surface), we often refer to it as the *test surface* and refer to the collection of surfaces making up the rest of the scene together with the illuminant as the *scene context*.

provides the information about surface reflectance and illumination available at one retinal location, corresponding to one surface patch in the flat–matte–diffuse environment. We will use superscripts to distinguish cone excitation vectors corresponding to different retinal locations (or to surfaces at different locations in the scene). For flat–matte–diffuse conditions, the cone excitation vectors $\{\rho^1, \dots, \rho^n\}$ across all retinal locations carry the information available to the visual system about illuminant and surface properties.

Naturally occurring illuminant spectral power distributions and surface reflectance functions exhibit considerable regularity (Cohen, 1964; DiCarlo & Wandell, 2000; Jaaskelainen, Parkkinen, & Toyooka, 1990; Judd, MacAdam, & Wyszecki, 1964; Maloney, 1986). Contemporary methods for estimating illuminant and surface properties from cone excitation vectors (or from camera sensor responses) incorporate models of this regularity (Maloney, 1999). Rather than specifying the full spectral function for each surface and illuminant, one can instead use a low-dimensional vector (see Maloney, 1986, 1999). For surfaces, we denote the resulting vector by σ and for illuminants, ε . We refer to these as the *surface* and *illuminant coordinates*. Any choice of surface coordinates σ specifies a particular surface reflectance function, $S_\sigma(\lambda)$, within the constrained class of reflectance spectra, and any choice of illuminant coordinates ε specifies a particular illuminant spectral power distribution, $E_\varepsilon(\lambda)$, within the constrained class of illuminant spectral power distributions.

Once we have selected surface and illuminant coordinates, we can rewrite Equations 1 and 2 schematically as

$$\rho = \langle \varepsilon | \sigma \rangle. \quad (3)$$

We refer to Equation 3 as the *rendering equation* for a single surface and illuminant, since it converts a description of physical *scene parameters*, here ε and σ , into the information available to the visual system, the cone excitation vector ρ .¹ The key point preserved by the abstract notation of Equation 3 is that the initial retinal information available to the visual system depends both on the illuminant and on the surface.

Equation 3 is appropriate for flat–matte–diffuse conditions. For these conditions, it is not necessary to include any specification of object shape, material, illuminant geometry, or observer viewpoint. Later, as we extend to richer viewing conditions and incorporate knowledge of the human visual system, we will generalize the rendering equation.

Although flat–matte–diffuse scenes are simple, Equation 3 makes clear that they incorporate the invariance problem described in the Introduction section. Even when constraints on natural illuminants and surfaces are imposed, the same surface reflectance, parameterized by σ , can produce different retinal excitations ρ as the illumination, parameterized by ε , varies. Indeed, for realistic constraints on

natural spectra there are typically many combinations of ε and σ that produce the same retinal excitations ρ . Starting with only photoreceptor excitations ρ^j at a set of locations (indexed by j) across the retina, estimating surface reflectance at each location in the scene under conditions where the scene illumination is unknown is an underdetermined problem. At the same time, Equation 3 satisfies an important constraint that we refer to as *surface–illuminant duality*.

Surface–illuminant duality

From Equation 3, we see that given the coordinates of illumination ε and surface σ we can compute the information available to the visual system ρ . However, for typical choices of surface and illuminant coordinate systems, we can do more. The following two properties emerge from analyses of the type of low-dimensional linear model surface–illuminant coordinate systems used in many computational analyses of color constancy and are valid for flat–matte–diffuse conditions (D’Zmura & Iverson, 1993a, 1993b; Maloney, 1984). Such coordinate systems provide accurate approximations to natural surfaces and daylights (Cohen, 1964; DiCarlo & Wandell, 2000; Jaaskelainen et al., 1990; Judd et al., 1964; Maloney, 1986).

Surface–Illuminant Duality Property 1: Given the illumination coordinates ε and retinal information ρ^j corresponding to any illuminated surface, we can solve for that surface’s coordinates σ^j .

Consequently, there are a wide range of color constancy algorithms that first estimate the coordinates of the scene illuminant and then use this information in conjunction with retinal information to estimate surface reflectance (for a review, see Maloney, 1999). These are referred to as two-stage algorithms (D’Zmura & Iverson, 1993a). If the estimate of the illuminant coordinates is correct, then so are the estimates of the surface coordinates.

Moreover, we can reverse the roles of illuminant and surface.

Surface–Illuminant Duality Property 2: If we know the coordinates of a sufficient number of surfaces $\{\sigma^j; j = 1, \dots, n\}$ and the corresponding retinal information $\{\rho^j; j = 1, \dots, n\}$, then we can solve for ε . The number of surfaces needed depends on the complexity of the lighting model.

The key idea embodied in both properties is that, given the illuminant and retinal excitations, we know the surfaces and vice versa. Demonstrating the precise conditions on choice of surface–illuminant coordinates for which these properties hold exactly is a mathematical problem (D’Zmura & Iverson, 1993a, 1993b, 1994; Maloney, 1984) that need not concern us here.

For the illuminant and surface coordinates that we employ for flat–matte–diffuse conditions, the requisite number of surfaces for Property 2 to hold is one ($n = 1$). In

our consideration of a wider range of possible scenes later in the paper, information about more than one surface is required to determine the illuminant coordinates from the retinal information and surface coordinates.

Human color constancy

We turn now to consider human color vision and to understand how we can exploit surface–illuminant duality to develop and test models of human performance. *A priori*, computational estimates of the coordinates of illuminants and surfaces need have little or nothing to do with human color vision. After all, our subjective experience of color is not in the form of spectral functions or coordinate vectors. Rather, we associate a percept of color appearance with object surfaces. This percept is often described in terms of its “hue,” “colorfulness,” and “lightness,” and these terms do not immediately connect to the constructs used in computations. A natural link between perception of color and computational estimation arises, however, when we consider stability of representation across scene changes.

As Helmholtz emphasized over a century ago, an important value of color appearance is to represent the physical properties of objects and to aid in identification: “Colors have their greatest significance for us in so far as they are properties of bodies and can be used as marks of identification of bodies (Helmholtz, 1896, p. 286).” For color appearance to be useful in this manner, the appearance of any given object must remain stable across changes in the scene in which it is viewed. To the extent that visual processing assigns the same surface color percept to a given physical surface, independent of the illuminant and the other surfaces in the scene, we say that the visual system is *color constant*.

A long history of research on color constancy for flat–matte–diffuse and closely related viewing conditions has provided us with a mature empirical characterization within the bounds of these conditions (Arend & Reeves, 1986; Brainard, 1998; Brainard, Brunt, & Speigle, 1997; Brainard & Wandell, 1992; Delahunt & Brainard, 2004; Granzier, Brenner, Cornelissen, & Smeets, 2005; Helson & Jeffers, 1940; Katz, 1935; Kraft & Brainard, 1999; McCann et al., 1976; Olkkonen, Hansen, & Gegenfurtner, 2009). This work may be summarized by two broad empirical generalizations (Brainard, 2004, 2009). First, when the collection of surface reflectances in the scene is held fixed and only the spectral power distribution of the illuminant is varied, human color constancy is often quite good (see, e.g., Brainard, 1998). Maloney (1984) referred to this case as *scene color constancy*, to distinguish it from the more general case where both the illumination and surrounding surfaces change. Gilchrist (2006) refers to it as *illumination-independent*

color constancy and we follow his terminology here. A visual system can be illumination-independent color constant without being color constant in general.

Because of the conceptual prominence of illuminant changes in the formulation of the constancy problem, the bulk of research on human color constancy has focused on illumination-independent color constancy. The well-known retinex algorithms of Land (1986) and Land and McCann (1971; see Brainard & Wandell, 1986), for example, achieve approximate illumination-independent constancy; the work on “relational color constancy” by Foster and Nascimento (1994) also focuses on this special case.

However, if both the surfaces surrounding a test surface and the illumination change (as when an object is moved to a different scene under different lighting), human color constancy can be greatly reduced (see, e.g., Brainard, 1998; Kraft & Brainard, 1999). What such results drive home is that we must seek models that predict in detail how surface color appearance depends on both the illumination of the scene and on the surrounding surfaces in the scene.

Equivalent illumination models

For flat–matte–diffuse conditions, surface–illuminant duality means that if the visual system had an estimate of the coordinates of the scene illuminant, then it would be straightforward to use to estimate the coordinates of the surface reflectance function at each location in the scene from the cone excitation vectors—and *vice versa*. This observation suggests an approach to modeling object surface perception, which we call the *equivalent illumination model (EIM)* approach (Brainard et al., 1997; Brainard & Wandell, 1991; Brainard, Wandell, & Chichilnisky, 1993; Speigle & Brainard, 1996). The idea is illustrated in Figure 3. The EIM approach supposes that visual processing proceeds in the same general two-stage fashion as many computational surface estimation algorithms. First, the visual system uses the cone excitation vectors from all surfaces in the scene to form an estimate of the illuminant coordinates, $\hat{\epsilon}$. We refer to this estimate as the *equivalent illuminant*. Second, the parameters $\hat{\epsilon}$ are used to set the transformation between the color signal corresponding to each surface and that surface’s perceived color. Thus, in this approach, color appearance at each location j can be thought of as a function of an implicit estimate of object surface coordinates, $\hat{\sigma}^j$. These, in turn, are generated by processing the color signal in a manner that depends on the equivalent illuminant.

Of critical importance is that we do not assume that the equivalent illuminant is correct or even close to correct. Rather, the equivalent illuminant $\hat{\epsilon}$ is a state variable that, although not directly observable, characterizes how the color signal is transformed to color appearance for a particular viewing context. Thus, when $\hat{\epsilon}$ deviates from the coordinates of the physical scene illuminant, color

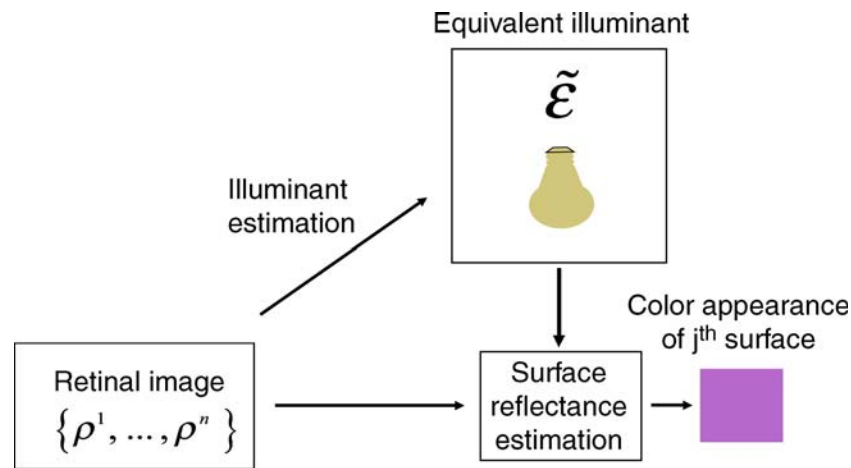


Figure 3. Equivalent illumination models. The information available about the scene (the retinal image $\{\rho^1, \dots, \rho^n\}$) is used to form an estimate $\tilde{\epsilon}$ of the illuminant coordinates, the *equivalent illuminant*. This estimate, in turn, determines how the color signal reflected from each scene surface is transformed to the perceptual representation that is color appearance.

constancy may fail dramatically. The resulting estimates of surface coordinates $\tilde{\sigma}^j$ will also deviate from their physical counterparts.

This type of state space approach is familiar in color science and was perhaps best enunciated by W. S. Stiles in discussing retinal adaptation: "... we anticipate that a small number of variables—adaptation variables—will define the condition of a particular visual area at a given time, instead of the indefinitely many that would be required to specify the conditioning stimuli. The adaptation concept—if it works—divides the original problem into two: what are the values of the adaptational variables corresponding to different stimuli, and how does adaptation, so defined, modify the visual response to given test stimuli (Stiles, 1961, p. 246)." By adopting the equivalent illuminant approach, we, in essence, posit that the adaptation variables may be characterized as an equivalent illuminant and that the effect of adaptation is to stabilize object color appearance given the current equivalent illuminant. If the approach succeeds, then it offers a remarkably parsimonious description of human perception of matte surface color.

Suppose that a two-stage visual system is viewing a flat–matte–diffuse stimulus array with n surface patches whose true surface coordinates are $\sigma^1, \sigma^2, \dots, \sigma^n$. The true illuminant coordinates are ϵ but the visual system's estimate is $\tilde{\epsilon}$. Based on that estimate, the visual system estimates surface coordinates to be $\tilde{\sigma}^1, \tilde{\sigma}^2, \dots, \tilde{\sigma}^n$ and some or all of these estimates may be in error. However, the possible patterns of error that can occur are highly constrained. So although a two-stage algorithm can grossly misestimate surface coordinates, the resulting errors are patterned: knowledge of any one of them determines all the others as well as the misestimate of the illuminant. This constraint allows us to develop experimental

tests of two-stage algorithms as models of human color vision.

Indeed, suppose that we measure the color appearance human vision assigns to surfaces in several illuminated scenes. Discrepancies between the color appearance of the same physical surface across scenes can function as a measure of the type of patterned errors that two-stage algorithms predict. If the pattern of the measured appearance discrepancies is inconsistent with a particular two-stage algorithm, we can reject this algorithm as model of human performance. On the other hand, if the discrepancies conform to the predictions of the algorithm, we can use the data to infer how the visual system's estimate of the equivalent illuminant varies from scene to scene.

The EIM approach in flat–matte–diffuse scenes

The equivalent illuminant approach separates the modeling problem into two parts (Brainard & Wandell, 1992; Krantz, 1968; Maloney & Wandell, 1986; Stiles, 1961). First, what is the parametric form of the visual system's representation of the equivalent illuminant, and how well does this form account for the effect of changing scene context on color appearance? Second, how does the equivalent illuminant depend on the information that arrives at the eyes?

Brainard et al. (1997; see also Speigle & Brainard, 1996) addressed the first question using an asymmetric color matching paradigm (see Arend & Reeves, 1986; Brainard & Wandell, 1992; Burnham, Evans, & Newhall, 1957). Subjects adjusted a test surface so that its appearance

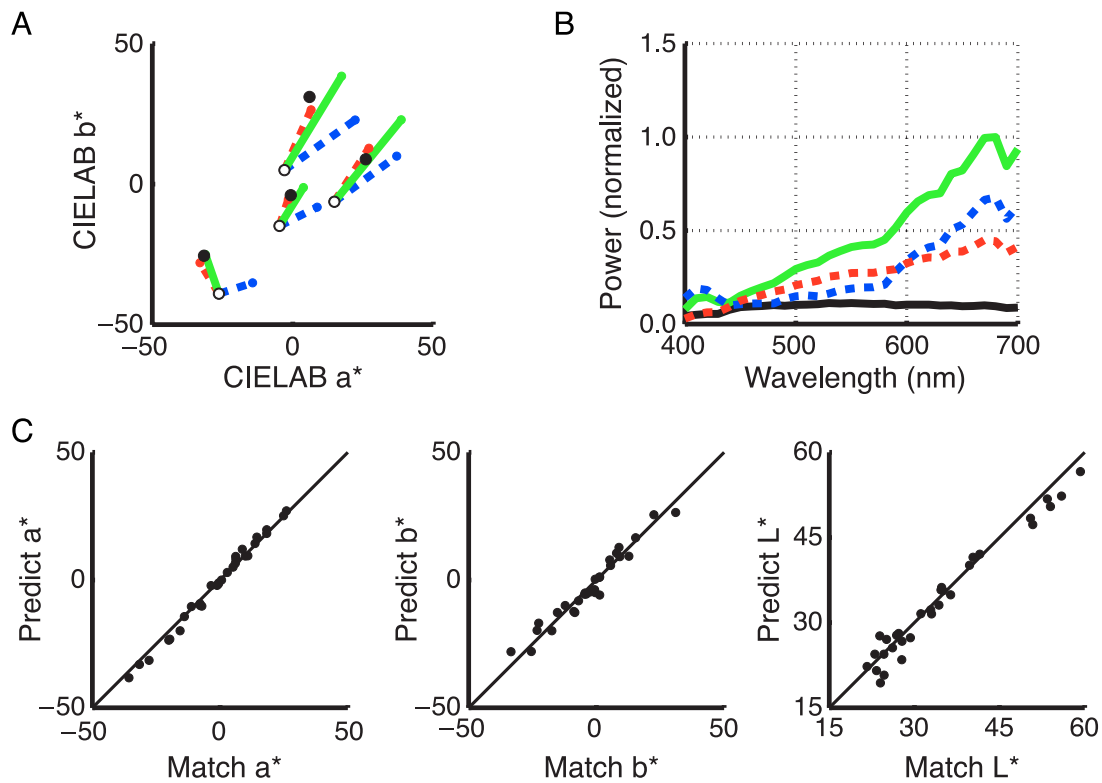


Figure 4. Equivalent illuminants in flat-matte scenes. (A) Asymmetric matching data and predictions. Data are shown as the CIELAB a^* and b^* coordinates of the color signal reaching the eye from the reference and matching test surfaces. Open black circles plot the a^*b^* coordinates of the color signal reflected from a reference surface. Closed black circles are the asymmetric matches, plotted as the coordinates of the light reflected from matching test surfaces. Points indicated by closed green circles (and connected to open black circles by solid green lines) show where the matches would lie for a color constant visual system. Equivalent illumination model predictions for two hypothetical choices of equivalent illuminant are shown by closed red and closed blue circles, connected to black open circles by red and blue dashed lines, respectively. The red closed circles are, in fact, the predictions of the best fitting equivalent illuminant. (B) Reference illuminant spectrum (solid black line), test illuminant spectrum (solid green line), and two equivalent illuminant spectra (red and blue dashed lines). These correspond to the equivalent illuminant predictions shown in (A). The equivalent illuminant shown in red provides the best fit to the data. Spectra shown are all within the parametric model for illuminant spectra and, therefore, differ from the physical spectra used in the experimental apparatus. (C) Quality of equivalent illuminant predictions. The CIELAB a^* , b^* , and L^* components of the predictions are plotted against the corresponding components of the asymmetric matches. All conversions to CIELAB were done using the test illuminant's tristimulus coordinates as the reference white.

matched to that of a reference surface. The test and reference surfaces were seen under separate illuminants. Matches were made for a number of different test surfaces.

Figure 4A plots a subset of Brainard et al.'s (1997) measurements. The data are shown as the CIELAB a^* and b^* coordinates of the color signal reaching the eye from the reference and matching test surfaces.² The open black circles plot the light reflected from the reference surfaces under the reference illuminant. When the illuminant is changed from reference to test, the light reflected from these surfaces changes. If the visual system made no adjustment in response to the change in the illuminant, then matches would have the same cone excitation vectors as the corresponding reference surfaces, and the data would lie near the open black circles. These can be considered the predictions for a visual system with no color constancy.

The closed green circles (connected to the open black circles by solid green lines), on the other hand, plot the light that would be reflected from the reference surfaces under the test illuminant. For a color constant visual system, the subjects' matches would correspond to these plotted points. That is, constancy predicts that the same physical surface has the same color appearance across the change in illuminant. Deviations from this prediction in the asymmetric matches represent discrepancies between the color appearance assigned to the same physical surface across the change in illuminant.

The closed black circles in Figure 4A plot the subject's actual matches. For these particular experimental conditions, the data are intermediate between the predictions of constancy (closed green circles) and the predictions of no constancy (open black circles).

Are the actual matches consistent with an equivalent illumination model? If so, then we should be able to find an equivalent illuminant such that the coordinates of the reference surfaces, rendered under the equivalent illuminant, lie near the data. Such predictions for two possible choices of equivalent illuminant are shown as red and blue closed circles in [Figure 4A](#). There are, of course, infinitely many possible equivalent illuminants, but we show only these two. The closed circles are connected to the open black circles by dashed lines of the same color.

In essence, the equivalent illuminant predictions represent the performance of a visual system that is color constant, up to a mismatch between the equivalent illuminant and the test illuminant. Each choice of possible equivalent illuminant predicts a precise pattern for the asymmetric matches across conditions. The actual data lie near the equivalent illuminant predictions shown in red for the sample matches shown, and indeed the predictions shown in red are those for the equivalent illuminant that best fits the data set. [Figure 4B](#) plots the spectra of the reference illuminant (solid black line), test illuminant (solid green line), and the two equivalent illuminants (dashed red and blue lines). [Figure 4C](#) summarizes the quality of the equivalent illuminant predictions for the full data set, for one observer and test illuminant. Brainard et al. (1997) present data for more subjects and a second pair of reference and test illuminants and show that in each case the equivalent illumination model provides a good fit.

The fact that a single equivalent illuminant can simultaneously predict all of the matches provides evidence that the equivalent illumination model provides a good account of human surface color perception for these experimental conditions. It also allows a compact summary of the entire data set, in the sense that knowing just the equivalent illuminant coordinates $\hat{\epsilon}$ allows accurate prediction of an entire matching data set. This description characterizes human performance in the currency of computational illuminant estimation algorithms and enables subsequent research to ask whether such algorithms can predict how the equivalent illuminant itself depends on the contextual image (see immediately below).

Predicting equivalent illuminants

There is an extensive literature on how a computational system can estimate the illuminant from cone excitation vectors for flat–matte–diffuse conditions (Brainard & Freeman, 1997; Buchsbaum, 1980; D’Zmura & Iverson, 1994; Finlayson, 1995; Finlayson, Hordley, & Hubel, 2001; Forsyth, 1990; Funt & Drew, 1988; Maloney & Wandell, 1986; for reviews, see Hurlbert, 1998; Maloney, 1992, 1999). Building on this computational foundation, Brainard et al. (2006; see Brainard, 2009) asked the second question central to the EIM approach: can the illuminant estimates provided by a computational algorithm predict the visual system’s equivalent illuminants?

Brainard et al. (2006) took advantage of the experimental work described above and employed a simplified experimental procedure, measurement of achromatic loci, that allows estimation of the observers’ equivalent illuminant for any contextual scene but that does not explicitly test the quality of these estimates against a full set of asymmetric matches. The interested reader is referred to papers that elaborate the link between achromatic loci and equivalent illuminants (Brainard, 1998) and the link between achromatic loci and asymmetric matches (Speigle & Brainard, 1999).

Brainard et al. (2006) applied a Bayesian color constancy algorithm (Brainard & Freeman, 1997) to the set of contextual images studied by Delahunt and Brainard (2004). This provided a computational estimate of the illuminant for each contextual image. Importantly, the computational estimates often deviated from the actual scene illuminants, depending on exactly what surfaces were in the scene. Brainard et al. (2006) then compared the computational estimates of the illuminant to the equivalent illuminants for the same images, which were obtained from analysis of the measured achromatic loci. For the 17 scenes studied, the agreement between computed illuminant estimates and measured equivalent illuminants was very good. [Figure 5](#) shows the comparison for 3 of the 17 conditions. Brainard et al. (2006) provide a more extensive analysis, as well as a technical discussion of the illuminant estimation algorithm and how its parameters were adjusted to fit the data.

The key point here is that the Bayesian algorithm provides a method to predict the equivalent illuminants from the contextual images, both for cases where constancy is good (comparison between left and center panels of [Figure 5](#)) and for cases where constancy is poor (comparison between left and right panels of [Figure 5](#)). Together with the fact that the equivalent illuminants can predict asymmetric matches, this means that the general equivalent illuminant approach yields a complete theory of how context affects color appearance for flat–matte–diffuse conditions. Although this theory still needs to be tested against more extensive data sets, the results to date suggest that it holds considerable promise.

Intermediate discussion

So far, we have described the EIM approach and shown how it can account for color constancy and failures thereof under flat–matte–diffuse conditions. Although this success is exciting in its own right, equally important is that the approach provides a recipe for generalization to more complex viewing conditions. The EIM recipe may be summarized as follows:

- I. Define the domain of scenes under study. Above, the scenes were defined by the flat–matte–diffuse assumptions.

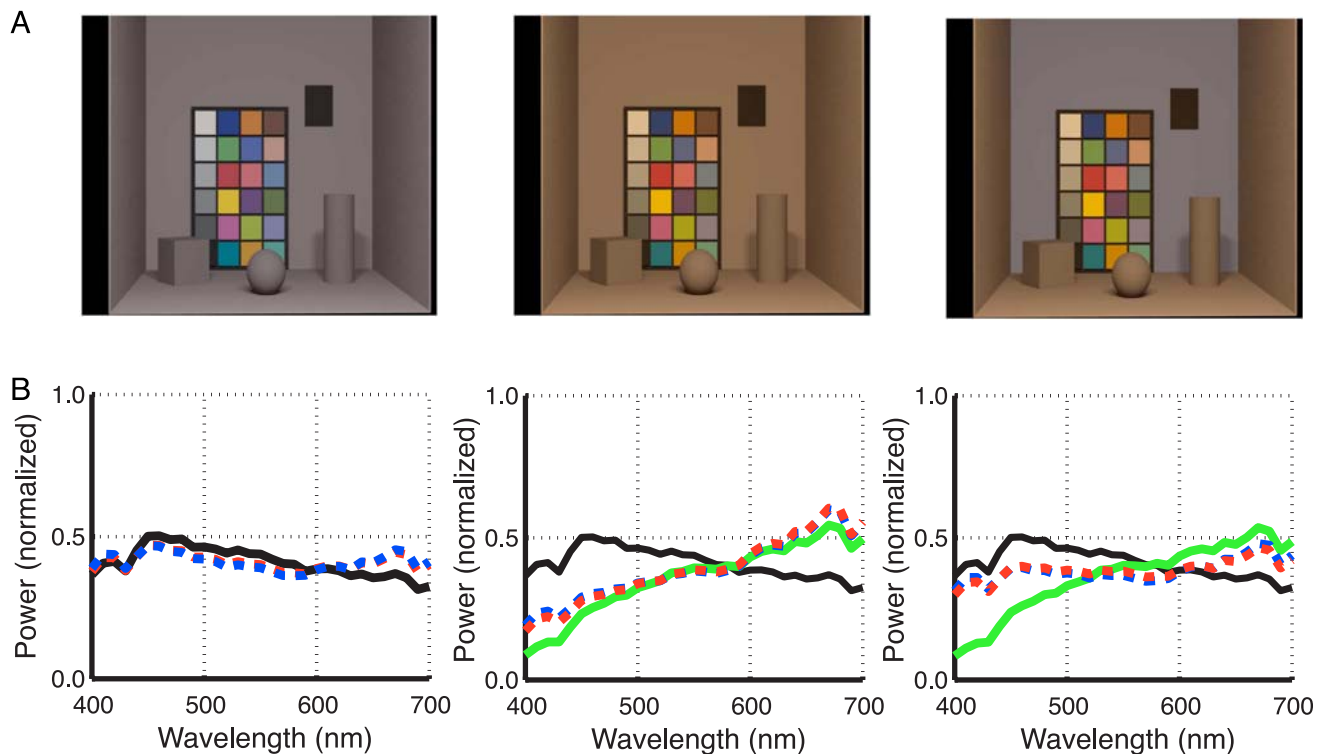


Figure 5. (A) Three scenes studied by Delahunt and Brainard (2004). Subjects viewed renderings of the scenes in stereo and set the chromaticity of a test patch (location is indicated by black rectangle) to appear achromatic. Measurements were made for 17 scenes. Across some scenes, only the illuminant changed (e.g., left to center). Across other scenes, both the illuminant and the reflectance of the back surface changed (e.g., left to right). This latter manipulation eliminated local contrast as a cue to the illuminant change. (B) The measured achromatic locus may be interpreted as the visual system's estimate of the illuminant (Brainard et al., 2006). This is shown as a red dashed line in all three panels. The corresponding illuminant estimated by the Bayesian algorithm is shown as a blue dashed line in each panel. In the left panel, the scene illuminant is shown as a solid black line. This line is replotted in the middle and right panels for comparison. In those panels, the physical illuminant is shown as a solid green line. The Bayesian algorithm predicts the human equivalent illuminants well, both for cases of good constancy (e.g., left to center panel scene change) and for cases of poor constancy (e.g., left to right panel scene change).

II. Choose illuminant and surface coordinates appropriate for the scenes defined in Step I and develop the rendering equation. Ask whether, for each scene studied, there is an equivalent illuminant that characterizes performance for all test surfaces in that scene. Above, the choice of three-parameter illuminant coordinates allowed derivation of equivalent illuminants that successfully predicted asymmetric matches for any reference surface across changes of scene illumination.

III. Develop computational techniques for estimating the illuminant from the information in the retinal image, for the class of scenes studied. Ask whether these techniques successfully predict the visual system's equivalent illuminants for any scene. Above, a Bayesian illuminant estimation algorithm successfully predicted equivalent illuminants derived from achromatic loci for scenes that conformed closely to flat–matte–diffuse conditions.

Of course, there is no *a priori* guarantee that this EIM recipe will prove successful for conditions other than flat–matte–diffuse. We may discover that human performance is not consistent with any choice of equivalent illuminant, correct or misestimated. Nonetheless, the fact that it has

been possible to elaborate Steps I, II, and III above into a promising model for flat–matte–diffuse conditions motivates asking whether the same recipe can be extended to richer viewing conditions. We turn to this question below and show that an affirmative answer is possible with respect to Steps I and II. Whether similar success will be possible for Step III awaits future research.

The EIM approach in more complex scenes

Figure 6 shows a photograph of the same matte surface at two different orientations and illustrates a physical effect that is introduced when we relax the restriction that the illuminant be diffuse and the assumptions that all surfaces lie in a single plane. The surface was embedded in the same scene for each of the two photographs, so that the location and intensity of the light source was the same. The photographs show that there is large effect of orientation




Figure 6. The effect of surface orientation. Photograph of the same flat-matte surface at two different slants relative to a directional light source. The figure illustrates that even when the configuration of light sources in a three-dimensional scene is held fixed, the three-dimensional pose of an object affects the amount of light reflected to the observer. Reprinted from Figure 1 of Ripamonti et al. (2004).

on the intensity of the light reflected from the surface. This occurs because the light source in the scene is largely directional, and for a directional light source, the intensity of the light reflected from even a matte surface varies the angle between the incident light and the surface normal (Pharr & Humphreys, 2004).

It is well established that the visual system can, under some circumstances, partially compensate for variation in surface orientation under directional illumination (Bloj & Hurlbert, 2002; Gilchrist, 1980; Hochberg & Beck, 1954; but see Epstein, 1961). There are no current models that predict the effect of surface azimuth on perceived lightness under varying illumination geometry. Here, we review our work that asks whether the EIM approach can lead us toward such a model.

The first step is to specify the class of scenes of interest and develop the rendering equation. To focus on geometric effects, we begin by restricting attention to achromatic scenes. We consider scenes illuminated by two light sources, one collimated and the other diffuse. A collimated light source can be thought of as a distant point light source with intensity ε_p . The diffuse source is a non-directional ambient illuminant with intensity ε_d . Because we are working only with achromatic lights, the illuminant coordinates are scalars. The theoretical framework developed here is easily extended to scenes in which illuminants are colored and illuminant coordinates are three-dimensional (Boyaci, Doerschner, & Maloney, 2004).

Let the angle between a normal to a surface and the direction to the collimated light source be denoted by θ . The light reflected from an achromatic matte surface is then proportional to

$$\rho = \begin{cases} [\varepsilon_p \cos\theta + \varepsilon_d]\sigma_a, & 0^\circ \leq |\theta| \leq 90^\circ \\ 0, & \text{otherwise,} \end{cases} \quad (4)$$

where σ_a denotes the fraction of light flux reflected from the surface. Since we are working with achromatic surfaces, σ_a (the surface *albedo*) and ρ are scalars. The restriction on θ simply guarantees that the light source is on the side of the surface being viewed.³

The value of θ is determined by the orientation of the surface and the direction to the collimated source. We can specify the orientation of the surface at any point by the azimuth and elevation⁴ $D_S = (\psi_S, \varphi_S)$ of a line perpendicular to the surface at that point, the surface normal. Similarly, we can specify the direction to the collimated light source by azimuth and elevation $D_E = (\psi_E, \varphi_E)$. We can compute θ for any choice of surface orientation D_S and direction to the collimated source D_E by standard trigonometric identities (Gerhard & Maloney, 2010): $\theta = \theta(D_E, D_S)$.

We rewrite Equation 2 (Bloj et al., 2004; Boyaci, Maloney, & Hersh, 2003) as

$$\rho = k[\pi + (1 - \pi)\cos\theta(D_E, D_S)]\sigma_a, \quad (5)$$

where k is a constant that depends on the intensities of the two illumination sources but not on D_E or D_S , and π is a measure of the intensity of the diffuse source relative to the collimated source and is again independent of D_E and D_S . We refer to π as *diffuseness*. We drop the explicit specification of the restriction on the range of θ in Equation 5 and following, but it is still in force.

The interpretation of Equation 5 is straightforward. The total illumination incident on the surface arises as a mixture of light from the collimated and diffuse sources. The amount of the collimated source that is reflected is proportional to the cosine of the difference in angle between the surface normal and direction to the collimated source. When this angular difference is 0, the

proportion from the collimated source reaches a maximum. This occurs when the surface directly faces the collimated source.

The experiments investigating the effect of geometric factors considered below are designed so that we can normalize the data and neglect any effects of this constant. With such normalization, we again simplify the rendering equation:

$$\rho = \langle \varepsilon | \sigma \rangle = [\pi + (1 - \pi) \cos \theta(D_E, D_S)] \sigma_a. \quad (6)$$

Equation 6 is analogous to Equation 3 above. It expresses the information available to the visual system (ρ) as a function of illuminant and surface coordinates. In Equation 6, the illuminant coordinates ε are π , $D_E = (\psi_E, \varphi_E)$ while the surface coordinates σ are σ_a , $D_S = (\psi_S, \varphi_S)$.

Experiments and EIM models

In two sets of psychophysical experiments, we measured how the lightness of a test surface changed as its orientation was varied in azimuth (Boyaci et al., 2003; Ripamonti et al., 2004), by having observers set asymmetric matches between test surfaces presented at different orientations and reference surfaces presented in a separated fixed context and oriented fronto-parallel to the observer. The experiments are complementary in one respect: Boyaci et al.

(2003) used rendered stimuli presented binocularly while Ripamonti et al. (2004) used physical surfaces and lights, also with binocular viewing. In both, the collimated light sources were directed from behind the observer and out of view.

Boyaci et al. considered only one illumination geometry: the parameters π , (ψ_E, φ_E) did not vary across the variations in test surface orientation. Ripamonti et al. examined human performance with the collimated illuminant at either of two azimuths and a single elevation. In both experiments, only the azimuth of the test surface was varied; the parameter φ_S was held constant. These restrictions are to allow us to further simplify the illuminant–surface coordinates by folding the effect of φ_E and φ_S into the normalizing constant k of Equation 5 and subsequently neglecting them. The theoretical goal of the experiments was to determine (a) whether observers' lightness matches were systematically affected by changes in surface azimuth and (b) whether an equivalent illumination model could describe performance, and if so, what were the coordinates $\tilde{\pi}$ and $\tilde{\psi}_E$ of the equivalent illuminant for each set of experimental conditions.

Because of surface–illuminant duality, each choice of equivalent illuminant parameters $\tilde{\pi}$ and $\tilde{\psi}_E$ leads to a prediction of how the (normalized) equivalent surface albedo estimate, $\tilde{\sigma}_a$, should vary as a function of surface azimuth ψ_S . The green solid curve shown in Figure 7A shows this dependence for the actual illumination param-

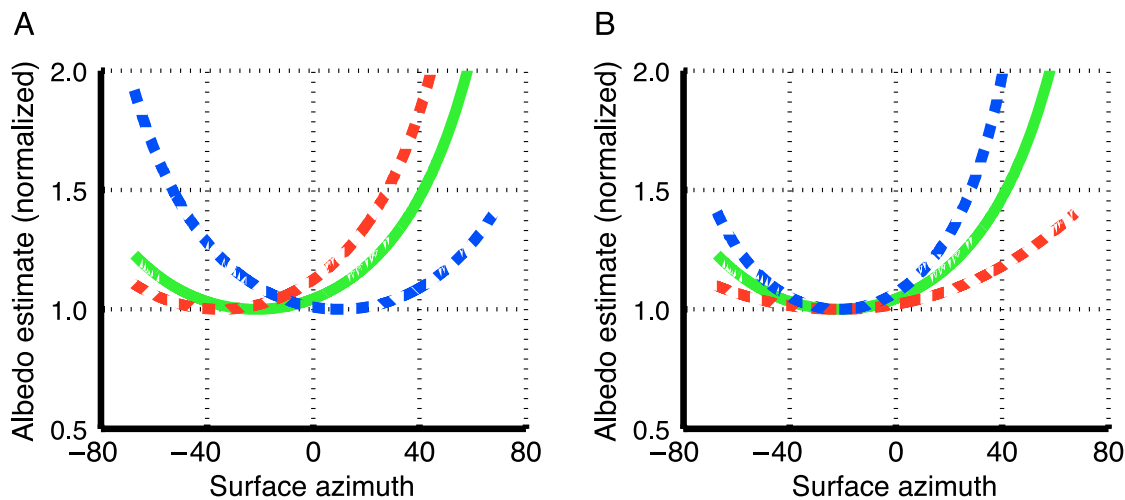


Figure 7. Matching functions. (A) The normalized surface albedo of a surface of constant luminance is plotted versus surface azimuth for a scene illuminated by a combination of collimated and diffuse sources (solid green curve). We refer to such a curve as a *matching function*. Surface albedo is normalized so that its minimum is one. This minimum occurs when the surface azimuth is identical to the azimuth of the collimated source. When the surface azimuth differs from that of the collimated source, the surface must have higher albedo to produce the same luminance. If a visual system misestimates the azimuth of the collimated light source but otherwise computes an estimate of surface albedo correctly, the resulting plot of estimated normalized surface albedo vs. azimuth will be shifted so that its minimum falls at the estimated azimuth but is otherwise unchanged. Two examples are shown as dashed red and blue lines. (B) If a visual system misestimates the relative intensity of the collimated source but otherwise computes an estimate of surface albedo correctly, the resulting plot of relative surface albedo vs. azimuth will be shallower or steeper as shown but otherwise unchanged. Two examples are shown as dashed red and blue lines. The solid green curve is replotted from (A).

eters π , ψ_E used in the experiments of Boyaci et al. This plot represents predictions for the case where the luminance reaching the observer from the test surface is held fixed across the changes of azimuth, again as was done in the experiments of Boyaci et al. In addition, because of the normalization procedure used in the data analysis, the units of albedo are arbitrary and here have been set so that their minimum is one. We refer to the plot as a *matching function*. If the observer's estimates of albedo were based on a correct physical model with correct estimates of the illuminant, then his or her normalized matches would fall along this particular matching function.

However, the observer's equivalent illuminant may differ from the physically correct values. The red and blue dashed curves in Figure 7A show matching functions for two choices of $\tilde{\psi}_E$ that differ from the value of the physical illumination, with $\tilde{\pi}$ set equal to the physically correct value. Each equivalent illuminant matching function reaches its minimum when $\psi_S = \tilde{\psi}_E$, that is when the surface azimuth agrees with the azimuth of the direction estimated for the collimated light source. This change is readily interpretable: a higher albedo is required to predict a constant reflected luminance as the surface normal is rotated away from the direction of the collimated source.

Figure 7B shows the effect of varying the diffuseness parameter $\tilde{\pi}$. The solid green curve is replotted from Figure 7A and again shows the matching function where the observer's equivalent illuminant coincides with the physically correct illuminant. The matching functions for two other values of $\tilde{\pi}$ are shown as dashed red and blue curves. As $\tilde{\pi}$ increases, the curve becomes shallower, and as $\tilde{\pi}$ decreases, the curve becomes steeper. These changes are also interpretable: in scenes illuminated primarily by diffuse light, the intensity of light reflected from a surface varies little with surface orientation, while the maximum variation occurs when the illumination is entirely collimated. In the limiting case where $\tilde{\pi} = 1$, the illumination is purely diffuse and the estimated albedo does not vary with surface azimuth.

The two parameters $\tilde{\psi}_E$ and $\tilde{\pi}$ independently affect the location and shape of the matching function. Given normalized matches obtained psychophysically, parameter search may be used to find the illuminant coordinates that provide the best fit to the data. Figure 8A shows an example for one observer from Boyaci et al. (2003) obtained in this manner. The closed black circles show normalized observer matches, and the dashed red curve shows the matching function for the best fitting equivalent




Figure 8. Estimates of equivalent illuminants in three-dimensional scenes. (A) Data and fit for one observer redrawn from Boyaci et al. (2003). Observers viewed rendered scenes illuminated by a combination of collimated and diffuse light sources and matched the perceived albedo of a reference surface to a test surface that varied in orientation. The luminance rather than the albedo of the test surface was held constant across orientations. Viewing was binocular. The true azimuth of the collimated source is marked by a green vertical line and the true matching function is shown as a solid green curve. The observer's fit matching function has an azimuth estimate close to the true value but is shallower. The equivalent illuminant has a higher diffuseness value than the true value. The fit matching function (dashed red curve) is in good agreement with the data. In particular, the observer is sensitive to the effect of surface azimuth in scenes illuminated by a combination of punctuate and diffuse sources. (B, C) Data for one observer in two conditions redrawn from Ripamonti et al. (2004). Observers viewed real scenes illuminated by a combination of collimated and diffuse light sources and matched the perceived albedo of a reference surface to a test surface that varied in orientation. Viewing was binocular. In this experiment, the luminance of the test surface was not held constant across orientations, but measurements were made for a number of different surface albedos. The data here are combined across surface albedos and replotted to show the effect inferred for constant luminance. Each panel shows a different illuminant configuration. The two configurations differed primarily in the azimuth of the collimated source. The format is as in (A). Once again, the observer's settings are consistent with an equivalent illuminant whose azimuth parameter is close to that of the actual collimated source but whose diffuseness is higher than that of the physical illuminant.

illuminant coordinates. **Figures 8B** and **8C** shows data and fits in the same format for one observer and two collimated source positions from Ripamonti et al. (2004).

If the observers were lightness constant, the data would fall along the solid green line in each panel—this line shows the EIM prediction for the case where the equivalent illuminant parameters match those of the scene illuminant. If the observers were not sensitive to surface azimuth, on the other hand, the data would lie along a horizontal line at a normalized value of one. The actual data are intermediate between these two possibilities and show partial lightness constancy. Moreover, in each case, the data are described by an equivalent illumination model (dashed red line). The original papers (Bloj et al., 2004; Boyaci et al., 2003; Ripamonti et al., 2004) show fits for more observers, provide a detailed analysis of the quality of the model fits, and conclude that the EIM provides a good account of the data sets from both laboratories.

We emphasize that the observers in both experiments were not color constant or even close. However, their matching data are highly constrained, indicating that we can predict their matches as a function of azimuth based on an equivalent illuminant that we could recover from their data.

As with our EIM analysis of color constancy for flat–matte–diffuse conditions, the EIM developed here translates the psychophysical data into the currency of illuminant estimation computations. For example, comparison of **Figures 8B** and **8C** shows that when the physical illuminant is shifted in azimuth, observer’s estimates approximately track this shift. This pattern was shown by all the observers of Ripamonti et al. (2004). In addition, the azimuth of the recovered equivalent illuminants was fairly accurate in all of the experiments.

On the other hand, observers’ equivalent illuminants generally underestimated the strength of the collimated component of the illumination (compare dashed red and solid green lines in **Figure 8**). Recently, Morgenstern, Murray, and Geisler (2010) reported measurements of the directionality of natural illumination fields and showed that these were considerably more diffuse than the stimuli used in the experiments discussed above—perhaps the underestimation is an appropriate response of an ideal observer armed with good priors. In any case, it will be interesting to see whether future algorithms that estimate geometrical illuminant parameters in a fashion that predicts the empirically derived equivalent illuminant parameters can be developed.

Additional experiments

In subsequent work, Boyaci et al. (2004) were able to use the same EIM principles to account for measurements of color appearance for surfaces viewed at different orientations in binocularly presented rendered scenes with a yellow collimated source and a blue diffuse source. They

varied both the azimuth ψ_S and elevation φ_S of a matte test surface and for two locations of the collimated source. Observers adjusted the test surface at each orientation until it appeared achromatic. Their data allowed recovery of equivalent illuminant parameters $\tilde{\pi}$, $\tilde{\psi}_E$, and $\tilde{\varphi}_E$. Two settings determined the full shape of each EIM prediction, and the fact that all the settings fell near a single matching function provided evidence that observers were discounting illumination for a scene much like the physical one but with respect to estimates of the illuminant coordinates that deviated from those of the scene illuminant.

The EIM approach has also been applied to yet more complex lighting conditions. Doerschner, Boyaci, and Maloney (2007) explored surface color perception in scenes with two yellow collimated light sources and a blue diffuse light source and found that observers’ performance was consistent with an EIM that implied that, while they underestimated the intensity of the collimated sources, they successfully compensated for the presence of two light sources in the scenes they viewed.

Discussion

Generalizability

An enormous challenge for understanding color and lightness constancy is that, as the complexity of the scenes studied increases, there is an explosion of scene variables that could affect performance. Once we leave the limited range of scenes consisting of flat, matte, coplanar objects and venture into the full complexity of 3D scenes, the spatial distribution of the light field and geometric properties of object surface reflectance must be considered, along with how they interact with object shape, roughness, and pose within the scene. Without some principled way to generalize the understanding we gain from simple scenes, the task of measuring and characterizing the interaction of all the relevant scene variables in terms of how they affect surface color perception seems so daunting as to be hopeless.

The success of EIMs in accounting for the parametric structure of surface color perception for both flat–matte–diffuse and simple three-dimensional scenes is promising. It tells us that a single set of principles can be used to provide a unified account of constancy across this range of scenes and suggests that the same ideas may continue to work as the field moves to yet richer viewing conditions (for an overview of some recent work, see Maloney & Brainard, 2010).

The generalization success currently applies to the first step of the EIM program, in which the modeling is used to account for the parametric form of the effect of manipulations of a surface (e.g., reflectance, azimuth) within any given context. This step is silent about how the visual

system processes the image data available at the retinas to set the values of the equivalent illuminant parameters for any scene. That is, formulating and testing an equivalent illumination model is the first step in a two-step modeling process. The second step is to understand how the image data determine the equivalent illuminant parameters. A number of authors have emphasized the usefulness of the general two-step approach in the context of color vision (Brainard & Wandell, 1992; Krantz, 1968; Maloney, 1984; Smithson & Zaidi, 2004; Stiles, 1961, 1967). We emphasize that the first step of the modeling approach leads to testable empirical predictions, independent of how the equivalent illuminant parameters are determined by the image.

Equivalent illumination models also lead to a natural approach for the follow-on second step. An EIM “operationalizes” (Bridgman, 1927) what it means to estimate the illuminant as a step in determining surface color, and for this reason, it provides a direct connection to computational algorithms that take image data and estimate properties of the illuminant. As reviewed above, taking this connection seriously has allowed successful prediction of the equivalent illuminants from image data for flat–matte–diffuse conditions (Brainard, 2009; Brainard et al., 2006). This was possible in large part because there was a considerable literature on illuminant estimation algorithms for these conditions (Brainard & Freeman, 1997; Buchsbaum, 1980; D’Zmura & Iverson, 1994; Finlayson, 1995; Finlayson et al., 2001; Forsyth, 1990; Funt & Drew, 1988; Maloney & Wandell, 1986; for reviews, see Hurlbert, 1998; Maloney, 1992, 1999).

It remains to be seen whether generalization success will be possible for the second step as we consider how image information in three-dimensional scenes sets the parameters of the equivalent illuminants, that is, what are the “cues” to scene illumination (Maloney, 2002). Candidate cues to the illumination can be evaluated experimentally and there is some work along these lines (Boyaci, Doerschner, & Maloney, 2006; Kraft, Maloney, & Brainard, 2002; Yang & Maloney, 2001; Yang & Shevell, 2002). Further research in this direction is needed.

Because EIMs translate human performance into parameters that describe the illuminant, however, any algorithm that estimates these same parameters from the images is immediately available as a candidate for this step in the modeling step. This potential to capitalize naturally on developments in the computer vision literature is an important feature of the general modeling approach.

EIMs and constancy

A common misconception about the equivalent illuminant approach is that it predicts perfect constancy. This is a misconception whose origins are unclear to us, but we

encounter it repeatedly when we discuss our work on EIM models. Perhaps it occurs because the computations required to make predictions with EIMs resemble those required to make the predictions for color constant performance. Indeed, if the equivalent illuminant coordinates match those of the actual illumination within the scene, then the EIM computations do predict constancy. The crucial difference is that EIMs do not require such a match between equivalent illuminant and scene illuminant coordinates. Indeed, we are excited about equivalent illumination models in part because they can account quite precisely and parsimoniously for the way in which observers *fail* to be constant, for the constraints that appear in their data, and for individual differences in the degree of constancy failure (see Gilchrist et al., 1999, for a general discussion of the diagnosticity of accounting for failures of constancy).

EIMs and perceptual judgments of the illuminant

Human observers are capable of making explicit judgments about the scene illuminant. Some authors have thus probed EIM-like models by asking whether the explicitly judged scene illuminant is, in essence, consistent with the equivalent illuminant parameters required to account for surface perception in the same scene (Beck, 1961, 1972; Granzier, Brenner, & Smeets, 2009; Logvinenko & Menshikova, 1994; Rutherford & Brainard, 2002). Although the number of empirical studies along these lines is limited, the general conclusion is negative: establishing the equivalent illuminant parameters via explicit judgments of the illuminant does not appear to be an effective method for understanding surface color appearance. Because perception of illumination per se has not been extensively studied, it may be that there are general principles awaiting discovery that will link it to the perception of surface lightness and color. In the meantime, however, the equivalent illuminant parameters should be understood as an implicit characterization of the illuminant, a state variable for perception of surface color, as discussed above. In this sense, the parameters are an important model construct, but one that need not be represented explicitly in consciousness or in neural responses.

EIMs, evolution, and mechanism

Although few would deny that evolution has shaped biological vision, there is little consensus on how this broad observation should shape experiments and theory. One approach focuses on the fact that evolution generally proceeds in a step-by-step fashion, one mutation at a time

(Darwin, 1859/1979; Leibniz, 1764/1996). The neural mechanisms underlying biological vision should, therefore, consist of a series of special purpose neural computations, each of which provided an incremental advantage to the organism at some point in its evolutionary history. This insight, in turn, leads to the view that the mechanisms of neural processing are best described as a “bag of tricks” (Cornelissen, Brenner, & Smeets, 2003; Ramachandran, 1985, 1990).

An emphasis on this aspect of evolution has motivated a number of authors to eschew the development of overarching functional models and instead seek descriptive models for each specialized mechanism. This is a research gambit, based on the hope that the process of characterizing individual mechanisms and their interactions will converge, so that at some point it will be possible to leverage the understanding gained and develop a model that describes overall performance.

The EIM models we describe, on the other hand, may be understood as motivated by a different principle of evolution. This is the idea that natural selection pushes biological systems toward optimal performance, within some set of environmental constraints (Geisler & Diehl, 2003). This observation suggests that, at a minimum, it would be fruitful to compare the actual performance of biological organisms in visual tasks to *normative models* of performance (Geisler, 1989).⁵ If biological performance approximates normative in a particular task, the experimentalist can attempt to develop a *descriptive model* based on the normative model. There is no *a priori* guarantee that a normative model will lead to an accurate descriptive model. The premise that it might do so is also a research gambit.

The apparent conflict between the two approaches is not as severe as it may seem. In terms of Marr’s (1982) influential taxonomy, EIMs are formulated at the computational level—they implicitly posit that an important function of color appearance is to provide information about object surface properties and they seek to describe the functional relation between the visual stimulus and this perceptual attribute. Mechanistic models are formulated at the algorithmic and hardware levels.

Although we agree that visual processing indeed results through the action and coordination of many separate and specialized neural mechanisms, we maintain that it is critical to understand these mechanisms as acting to implement computations that reveal broad regularities. To analogize: the logic circuits of classical digital computers can be expressed in terms of a single “trick,” a NAND logic gate that takes in two Boolean inputs (A, B) and computes the negation of the “and” of the two inputs. Yet to characterize a digital computer as a “bag with one trick” would be incomplete. This characterization fails to specify how the circuits of the computer are coordinated to carry out arbitrarily complex algorithms.

The computational perspective makes clear that mechanistic understanding must go beyond simply enumerating which “tricks” are in the “bag.” It is also critical to understand how the action of each “trick” is coordinated with the others, so that the appropriate processing is accomplished in any particular situation. As with our discussion of neural mechanism above, we think that successful equivalent illuminant models will provide critical functional constraints for organizing a theory of how individual “tricks” contribute to performance, how these “tricks” interact with each other, and how selection may have acted on individual “tricks” to lead to overall performance that is coherent and evolutionarily adaptive.

Conclusion

In this article, we reviewed progress toward understanding human performance for a particular visual task, the perceptual representation of surface properties. In particular, we considered work that takes as its point of departure normative models of object color perception and argue that this approach holds much promise. Although we restrict attention to a particular visual task, the general principles we outline apply generally to essentially any aspect of biological information processing. We close with a prescient remark by Barlow, characterizing the interplay of model and mechanism in the study of biological function: “A wing would be a most mystifying structure if one did not know that birds flew. [W]ithout understanding something of the principles of flight, a more detailed examination of the wing itself would probably be unrewarding. I think we may be at an analogous point in our understanding of the sensory side of the central nervous system. It seems to me vitally important to have in mind possible answers ... when investigating these structures, for if one does not one will get lost in a mass of irrelevant detail and fail to make the crucial observations (Barlow, 1961).”

Acknowledgments

This work was supported by NIH RO1EY10016 (DHB), NIH RO1EY019889 (LTM), NSF1059166 (LTM), and the Alexander von Humboldt Foundation (LTM).

Commercial relationships: none.

Corresponding author: David H. Brainard.

Email: brainard@psych.upenn.edu.

Address: Department of Psychology, University of Pennsylvania, 3720 Walnut Street, Philadelphia, PA, USA.

Footnotes

¹For purposes of this overview, we do not need to develop in detail the computational details implicit in Equation 3 (see Maloney, 1999).

²The CIELAB space represents the color signal reaching the eye and is computed from the cone excitation vectors of that color signal. The representation also incorporates a simple model of early visual processing. Roughly speaking, variation in the a^* direction corresponds to reddish-greenish perceptual variation, while variation in the b^* direction corresponds to bluish-yellowish variation. There is also an L^* coordinate that roughly corresponds to perceptual variation in lightness. In computing CIELAB coordinates, we omitted its model of retinal adaptation by fixing the reference white used in the transformation. Thus, in our use of the space here, the CIELAB coordinates are in a one-to-one invertible relation to cone excitation vectors.

³For matte surfaces, the luminance of the reflected light described by Equation 4 holds independent of observer viewpoint, as long as the observer can see the surface. For this reason, we do not explicitly specify observer position in the formulation.

⁴Azimuth and elevation can be thought of as latitude and longitude on a terrestrial sphere.

⁵We have adopted the terminology of normative and descriptive models from the decision-making literature, where it is widely used (see Bell, Raiffa, & Tversky, 1988). In the vision literature, normative models are often referred to as ideal observer models (Geisler, 1989; Green & Swets, 1966).

References

- Arend, L. E., & Reeves, A. (1986). Simultaneous color constancy. *Journal of the Optical Society of America A*, 3, 1743–1751.
- Barlow, H. B. (1961). Possible principles underlying the transformations of sensory messages. In W. A. Rosenblith (Ed.), *Sensory communication* (pp. 217–234). Cambridge, MA: MIT Press.
- 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.
- Bell, D. E., Raiffa, H., & Tversky, A. (1988). *Decision making: Descriptive, normative, and prescriptive interactions*. Cambridge, UK: Cambridge University Press.
- Bloj, M., Ripamonti, C., Mitha, K., Greenwald, S., Hauck, R., & Brainard, D. H. (2004). An equivalent illuminant model for the effect of surface slant on perceived lightness. *Journal of Vision*, 4(9):6, 735–746, <http://www.journalofvision.org/content/4/9/6>, doi:10.1167/4.9.6. [PubMed] [Article]
- Bloj, M. G., & Hurlbert, A. C. (2002). An empirical study of the traditional Mach card effect. *Perception*, 31, 233–246.
- Boyaci, H., Doerschner, K., & Maloney, L. T. (2004). Perceived surface color in binocularly viewed scenes with two light sources differing in chromaticity. *Journal of Vision*, 4(9):1, 664–679, <http://www.journalofvision.org/content/4/9/1>, doi:10.1167/4.9.1. [PubMed] [Article]
- Boyaci, H., Doerschner, K., & Maloney, L. T. (2006). Cues to an equivalent lighting model. *Journal of Vision*, 6(2):2, 106–118, <http://www.journalofvision.org/content/6/2/2>, doi:10.1167/6.2.2. [PubMed] [Article]
- Boyaci, H., Maloney, L. T., & Hersh, S. (2003). The effect of perceived surface orientation on perceived surface albedo in binocularly viewed scenes. *Journal of Vision*, 3(8):2, 541–553, <http://www.journalofvision.org/content/3/8/2>, doi:10.1167/3.8.2. [PubMed] [Article]
- Brainard, D. H. (1998). Color constancy in the nearly natural image: 2. Achromatic loci. *Journal of the Optical Society of America A*, 15, 307–325.
- Brainard, D. H. (2004). Color constancy. In L. M. Chalupa & J. S. Werner (Eds.), *The visual neurosciences* (pp. 948–961). Cambridge, MA: MIT Press.
- Brainard, D. H. (2009). Bayesian approaches to color vision. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (4th ed., pp. 395–408). Cambridge, MA: MIT Press.
- 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., Longere, P., Delahunt, P. B., Freeman, W. T., Kraft, J. M., & Xiao, B. (2006). Bayesian model of human color constancy. *Journal of Vision*, 6(11):10, 1267–1281, <http://www.journalofvision.org/content/6/11/10>, doi:10.1167/6.11.10. [PubMed] [Article]
- Brainard, D. H., & Wandell, B. A. (1986). Analysis of the retinex theory of color vision. *Journal of the Optical Society of America A*, 3, 1651–1661.
- Brainard, D. H., & Wandell, B. A. (1991). A bilinear model of the illuminant's effect on color appearance. In M. S. Landy & J. A. Movshon (Eds.), *Computational models of visual processing* (pp. 171–186). Cambridge, MA: MIT Press.

- Brainard, D. H., & 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.
- Brainard, D. H., Wandell, B. A., & Chichilnisky, E. J. (1993). Color constancy: From physics to appearance. *Current Directions in Psychological Science*, 2, 165–170.
- Bridgman, P. (1927). *The logic of modern physics*. New York: MacMillan.
- Buchsbaum, G. (1980). A spatial processor model for object colour perception. *Journal of the Franklin Institute*, 310, 1–26.
- Burnham, R. W., Evans, R. M., & Newhall, S. M. (1957). Prediction of color appearance with different adaptation illuminations. *Journal of the Optical Society of America*, 47, 35–42.
- Cohen, J. (1964). Dependency of the spectral reflectance curves of the Munsell color chips. *Psychonomic Science*, 1, 369–370.
- Cornelissen, F. W., Brenner, E., & Smeets, J. (2003). True color only exists in the eye of the observer. *Behavioral and Brain Sciences*, 26, 26–27.
- Darwin, C. R. (1859/1979). *On the origin of species by means of natural selection*. New York: Random House.
- Delahunt, P. B., & Brainard, D. H. (2004). Does human color constancy incorporate the statistical regularity of natural daylight? *Journal of Vision*, 4(2):1, 57–81, <http://www.journalofvision.org/content/4/2/1>, doi:10.1167/4.2.1. [PubMed] [Article]
- DiCarlo, J. J., & Cox, D. D. (2007). Untangling invariant object recognition. *Trends in Cognitive Science*, 11, 333–341.
- DiCarlo, J. M., & Wandell, B. A. (2000). *Illuminant estimation: Beyond the bases*. Paper presented at IS&T/SID Eighth Color Imaging Conference, Scottsdale, AZ, 91–96.
- Doerschner, K., Boyaci, H., & Maloney, L. T. (2007). Testing limits on matte surface color perception in three-dimensional scenes with complex light fields. *Vision Research*, 47, 3409–3423.
- D’Zmura, M., & Iverson, G. (1993a). Color constancy: I. Basic theory of two-stage linear recovery of spectral descriptions for lights and surfaces. *Journal of the Optical Society of America A*, 10, 2148–2165.
- D’Zmura, M., & Iverson, G. (1993b). Color constancy: II. Results for two-stage linear recovery of spectral descriptions for lights and surfaces. *Journal of the Optical Society of America A*, 10, 2166–2180.
- D’Zmura, M., & Iverson, G. (1994). Color constancy: III. General linear recovery of spectral descriptions for lights and surfaces. *Journal of the Optical Society of America A*, 11, 2389–2400.
- Epstein, W. (1961). Phenomenal orientation and perceived achromatic color. *Journal of Psychology*, 52, 51–53.
- Finlayson, G. D. (1995). *Color constancy in diagonal chromaticity space*. Paper presented at Proceedings of the 5th International Conference on Computer Vision, Cambridge, MA, 218–223.
- Finlayson, G. D., Hordley, S. D., & Hubel, P. M. (2001). Color by correlation: A simple, unifying framework for color constancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23, 1209–1221.
- Forsyth, D. A. (1990). A novel algorithm for color constancy. *International Journal of Computer Vision*, 5, 5–36.
- Foster, D. H., & Nascimento, S. M. C. (1994). Relational colour constancy from invariant cone-excitation ratios. *Proceedings of the Royal Society of London B: Biological Sciences*, 257, 115–121.
- Funt, B. V., & Drew, M. S. (1988). *Color constancy computation in near-Mondrian scenes using a finite dimensional linear model*. Proceedings Computer Vision and Pattern Recognition, IEEE Computer Society, pp. 544–549, Ann Arbor, MI, June 1988.
- Geisler, W. S. (1989). Sequential ideal-observer analysis of visual discriminations. *Psychological Review*, 96, 267–314.
- Geisler, W. S., & Diehl, R. L. (2003). A Bayesian approach to the evolution of perceptual and cognitive systems. *Cognitive Science*, 27, 379–402.
- Gerhard, H. E., & Maloney, L. T. (2010). Detection of light transformations and concomitant changes in surface albedo. *Journal of Vision*, 10(9):1, 1–14, <http://www.journalofvision.org/content/10/9/1>, doi:10.1167/10.9.1. [PubMed] [Article]
- Gilchrist, A. (2006). *Seeing black and white*. Oxford, UK: Oxford University Press.
- Gilchrist, A., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., et al. (1999). An anchoring theory of lightness perception. *Psychological Review*, 106, 795–834.
- Gilchrist, A. L. (1980). When does perceived lightness depend on perceived spatial arrangement? *Perception & Psychophysics*, 28, 527–538.
- 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(1):2, 20–27, <http://www.journalofvision.org/content/5/1/2>, doi:10.1167/5.1.2. [PubMed] [Article]
- Granzier, J. J. M., Brenner, E., & Smeets, J. B. J. (2009). Can illumination estimates provide the basis for color constancy? *Journal of Vision*, 9(3):18, 1–11,

- <http://www.journalofvision.org/content/9/3/18>, doi:10.1167/9.3.18. [PubMed] [Article]
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: John Wiley & Sons.
- Helmholtz, H. (1896). *Handbuch der Physiologischen Optik*: Translated into English by J.P.C Southall in 1924. Translation reprinted in 2000 by Thoemmes Press.
- Helson, H., & Jeffers, V. B. (1940). Fundamental problems in color vision: II. Hue, lightness, and saturation of selective samples in chromatic illumination. *Journal of Experimental Psychology*, 26, 1–27.
- Hochberg, J. E., & Beck, J. (1954). Apparent spatial arrangement and perceived brightness. *Journal of Experimental Psychology*, 47, 263–266.
- Hurlbert, A. C. (1998). Computational models of color constancy. In *Perceptual constancy: Why things look as they do* (pp. 283–322). Cambridge, UK: Cambridge University Press.
- Jaaskelainen, T., Parkkinen, J., & Toyooka, S. (1990). A vector-subspace model for color representation. *Journal of the Optical Society of America A*, 7, 725–730.
- Judd, D. B., MacAdam, D. L., & Wyszecki, G. W. (1964). Spectral distribution of typical daylight as a function of correlated color temperature. *Journal of the Optical Society of America*, 54, 1031–1040.
- Katz, D. (1935). *The world of colour*. London: Kegan, Paul, Trench, Trübner.
- Kraft, J. M., & Brainard, D. H. (1999). Mechanisms of color constancy under nearly natural viewing. *Proceedings of the National Academy of Sciences of the United States of America*, 96, 307–312.
- Kraft, J. M., Maloney, S. I., & Brainard, D. H. (2002). Surface–illuminant ambiguity and color constancy: Effects of scene complexity and depth cues. *Perception*, 31, 247–263.
- Krantz, D. (1968). A theory of context effects based on cross-context matching. *Journal of Mathematical Psychology*, 5, 1–48.
- Land, E. H. (1986). Recent advances in retinex theory. *Vision Research*, 26, 7–21.
- Land, E. H., & McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America*, 61, 1–11.
- Leibniz, G. W. (1764/1996). *New essays on human understanding*. Cambridge, UK: Cambridge University Press.
- 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. (1984). *Computational approaches to color constancy*. Unpublished Technical Report, Stanford University.
- Maloney, L. T. (1986). Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *Journal of the Optical Society of America A*, 3, 1673–1683.
- Maloney, L. T. (1992). Color constancy and color perception: The linear models framework. In D. E. Meyer & S. E. Kornblum (Eds.), *Attention and performance XIV: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience*. Cambridge, MA: MIT Press.
- Maloney, L. T. (1999). Physics-based approaches to modeling surface color perception. In *Color vision: From genes to perception* (pp. 387–416). Cambridge, UK: Cambridge University Press.
- Maloney, L. T. (2002). Illuminant estimation as cue combination. *Journal of Vision*, 2(6):6, 493–504, <http://www.journalofvision.org/content/2/6/6>, doi:10.1167/2.6.6. [PubMed] [Article]
- Maloney, L. T., & Brainard, D. H. (2010). Color and material perception: Achievements and challenges. *Journal of Vision*, 10(9):19, 1–6, <http://www.journalofvision.org/content/10/9/19>, doi:10.1167/10.9.19. [PubMed] [Article]
- Maloney, L. T., & Wandell, B. A. (1986). Color constancy: A method for recovering surface spectral reflectances. *Journal of the Optical Society of America A*, 3, 29–33.
- Marr, D. (1982). *Vision*. San Francisco: W. H. Freeman.
- 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.
- Morgenstern, Y., Murray, R. F., & Geisler, W. S. (2010). Real-world illumination measurements with a multi-directional photometer [Abstract]. *Journal of Vision*, 10(7):444, 444a, <http://www.journalofvision.org/content/10/7/444>, doi:10.1167/10.7.444.
- Olkkonen, M., Hansen, T., & Gegenfurtner, K. R. (2009). Categorical color constancy for simulated surfaces. *Journal of Vision*, 9(12):6, 1–18, <http://www.journalofvision.org/content/9/12/6>, doi:10.1167/9.12.6. [PubMed] [Article]
- Pharr, M., & Humphreys, G. (2004). *Physically based rendering: From theory to implementation*. San Francisco: Morgan Kaufmann.
- Ramachandran, V. S. (1985). The neurobiology of perception. *Perception*, 14, 97–105.

- Ramachandran, V. S. (1990). Visual perception in people and machines. In A. Blake & T. Troscianko (Eds.), *AI and the eye* (pp. 21–77). New York: Wiley.
- Ripamonti, C., Bloj, M., Hauck, R., Mitha, K., Greenwald, S., Maloney, S. I., et al. (2004). Measurements of the effect of surface slant on perceived lightness. *Journal of Vision*, 4(9):7, 747–763, <http://www.journalofvision.org/content/4/9/7>, doi:10.1167/4.9.7. [[PubMed](#)] [[Article](#)]
- Rust, N. C., & Stocker, A. A. (2010). Ambiguity and invariance: Two fundamental challenges for visual processing. *Current Opinion in Neurobiology*, 20, 382–388.
- Rutherford, M. D., & Brainard, D. H. (2002). Lightness constancy: A direct test of the illumination estimation hypothesis. *Psychological Science*, 13, 142–149.
- Smithson, H., & Zaidi, Q. (2004). Colour constancy in context: Roles for local adaptation and levels of reference. *Journal of Vision*, 4(9):3, 693–710, <http://www.journalofvision.org/content/4/9/3>, doi:10.1167/4.9.3. [[PubMed](#)] [[Article](#)]
- Speigle, J. M., & Brainard, D. H. (1996). 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.
- Stiles, W. S. (1961). Adaptation, chromatic adaptation, colour transformation. *Anales De La Real Sociedad Espanola De Fisica Y Quimica*, 57, 149–175.
- Stiles, W. S. (1967). Mechanism concepts in colour theory. *Journal of the Colour Group*, 11, 106–123.
- Ullman, S. (1989). Aligning pictorial descriptions: An approach to object recognition. *Cognition*, 32, 193–254.
- Xiao, B., & Brainard, D. H. (2008). Surface gloss and color perception of 3D objects. *Visual Neuroscience*, 25, 371–385.
- Yang, J. N., & Maloney, L. T. (2001). Illuminant cues in surface color perception: Tests of three candidate cues. *Vision Research*, 41, 2581–2600.
- Yang, J. N., & Shevell, S. K. (2002). Stereo disparity improves color constancy. *Vision Research*, 42, 1979–1989.