Bayesian Model Of Human Color Constancy

David H. Brainard
Department of Psychology, University of Pennsylvania, Philadelphia, PA
brainard@psych.upenn.edu
To Whom Correspondence Should Be Addressed

Philippe Longère
Neion Graphics, 2229 route des Crêtes, Valbonne, France
philippe.longere@neiongfx.com

Peter B. Delahunt
Posit Science, 114 Sansome Street, San Francisco, CA
peter.delahunt@positscience.com

William T. Freeman
Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA
billf@ai.mit.edu

James M. Kraft
Faculty of Life Sciences, The University of Manchester, Manchester, UK
j.kraft@manchester.ac.uk

Bei Xiao
Department of Neuroscience, University of Pennsylvania, Philadelphia, PA
beixiao@pobox.upenn.edu

Submitted to Journal of Vision. Comments Welcome.

Please do not quote without permission.

Abstract

Vision is difficult because images are ambiguous. For object color, the ambiguity arises because the same object reflects a different spectrum to the eye under different illuminations. Human vision typically does a good job of resolving this ambiguity – an ability known as color constancy. The past twenty years have seen an explosion of work on color constancy, with advances in both experimental methods and computational algorithms. Here we connect these two lines of research by developing a quantitative model of human color constancy. The model includes an explicit link between psychophysical data and illuminant estimates obtained via a Bayesian algorithm. The model is fit to the data through a parameterization of the prior distribution of illuminant spectral properties. The fit to the data is good, and the derived prior provides a succinct description of human performance.

Vision is useful because it informs an organism about the physical environment. Vision is difficult because the retinal stimulus is ambiguous about the state of the world. Such ambiguity is illustrated in Figure 1 for the case of object color: the same object reflects a different spectrum to the eye when it is viewed under different illuminations. The remarkable feature of human vision is that despite ambiguity in the sense data, our perceptions generally provide an accurate picture of what surrounds us. The core of any general theory of perception must include an understanding of how the ambiguity is so effectively resolved.

Figure 1 about here.

An attractive general principle is that vision resolves ambiguity by taking advantage of the statistical structure of natural scenes: given several physical interpretations that are consistent with the sense data, the visual system chooses the one that is most likely *a priori* (Helmholtz, 1866; Adelson & Pentland, 1996; Purves & Lotto, 2003). This principle can be instantiated quantitatively using Bayesian decision theory, and there is great interest in linking perceptual performance explicitly to Bayesian models (Knill & Richards, 1996; Rao, Olshausen, & Lewicki, 2002; Geisler & Kersten, 2002; Kersten & Yuille, 2003). This paper develops a quantitative Bayesian model of human color constancy. The model is fit to the data by adjusting parameters that describe the visual system's prior knowledge (Mamassian, Landy, & Maloney, 2002; Weiss, Simoncelli, & Adelson, 2002).

First, we review published measurements of human color constancy made by asking observers to adjust the appearance of test objects, embedded in images of three-dimensional scenes, until they appear achromatic (Delahunt & Brainard, 2004). The degree of constancy exhibited by the visual system varies greatly with the structure of the scenes used to assess it. Thus a general understanding requires a model that predicts both successes and failures of constancy (Gilchrist *et al.*, 1999). The core idea of our model is to use an explicit algorithm to estimate the scene illuminant from the image data, and to use this estimate to predict the color appearance data. When the algorithm provides correct estimates of the illuminant, the model predicts good color constancy. When the algorithm's estimates are in error, the model predicts failures of constancy.

We apply this modeling principle using a Bayesian illuminant estimation algorithm (Brainard & Freeman, 1997). The algorithm resolves ambiguity in the image through explicit priors that represent the statistical structure of spectra in natural images. These priors provide a parameterization that allows us to control the algorithm's performance. The model provides a good account of the extant data. The illuminant prior that provides the best fit to the data may be compared to the statistics of natural illuminants.

Methods

Measurements of Human Color Constancy

The experimental methods used in the psychophysical experiments are described in detail by Delahunt and Brainard (2004). The published report also describes a number of control experiments that verify the robustness of the main results.

Bayesian Illuminant Estimation

The principles underlying the Bayesian Illuminant Estimation algorithm are described by Brainard and Freeman (1997). The published algorithm was modified to estimate illuminant chromaticity rather than illuminant spectrum, as described below.

To predict the achromatic loci using the equivalent illuminant principle, we wish to estimate the CIE u'v' chromaticity coordinates of the illuminant. We denote these by the two-dimensional column vector

$$\mathbf{x} = \begin{bmatrix} u' \\ v' \end{bmatrix}.$$

Our goal is to estimate $\mathbf{x}$ from the data available in the retinal image. We denote this data by the vector $\mathbf{y}$. The entries of $\mathbf{y}$ consist of the L-, M-, and S-cone quantal absorption rates at a series of N_y image locations

$$\mathbf{y} = \begin{bmatrix} L_1 & M_1 & S_1 & \cdots & L_{N_y} & M_{N_y} & S_{N_y} \end{bmatrix}^{\mathrm{T}}.$$

The Bayesian approach (Berger, 1985; Lee, 1989) tells us that to estimate $\mathbf{x}$ from $\mathbf{y}$, we should calculate the *posterior probability*

$$p(\mathbf{x} \mid \mathbf{y}) = c\, p(\mathbf{y} \mid \mathbf{x})\, p(\mathbf{x}), \qquad (1)$$

where c is a normalizing constant that depends on $\mathbf{y}$ but not on $\mathbf{x}$, $p(\mathbf{y} \mid \mathbf{x})$ is the *likelihood* that data $\mathbf{y}$ will be observed when the illuminant is in fact $\mathbf{x}$, and $p(\mathbf{x})$ is the *prior probability* of an illuminant with chromaticity $\mathbf{x}$ occurring.

To choose an estimate $\tilde{\mathbf{x}}$ of the illuminant chromaticity, we maximized the posterior probability

$$\tilde{\mathbf{x}} = argmax[p(\mathbf{x} \mid \mathbf{y})].$$

To find the maximum, we need to evaluate both the likelihood and the prior. Since the constant c is independent of $\mathbf{x}$, it may be ignored in the maximization.

The likelihood of observing a data vector $\mathbf{y}$ given an illuminant $\mathbf{x}$ depends on the imaging model, the spectrum of the illuminant, and the surfaces that the illuminant reflects from. We assume that a single diffuse illuminant reflects from N_y distinct matte surfaces. We represent the illuminant spectrum by a column vector $\mathbf{e}$ whose entries denote the illuminant power for a set of wavelength bands that sample the visible spectrum (Brainard, 1995). We represent each individual surface by a column vector $\mathbf{s}_i$ ($1 \le i \le N_y$) whose entries represent the reflectance of the surface for the same set of wavelength bands. The reflected light is absorbed by L-, M-, and S-cone photopigment to produce the mean isomerization rates corresponding to each surface. These mean rates may be computed from $\mathbf{e}$,

$\mathbf{s}_i$, and estimates of the photopigment spectral absorption curves using standard methods (Brainard, 1995). We assume that the observed data specified by $\mathbf{y}$ are the mean absorption rates corresponding to each surface and perturbed by independent zero mean Gaussian noise. The standard deviation of this noise for each cone class was taken to be 1% of the mean value of $\mathbf{y}$ for that cone class. These assumptions allow us to compute explicitly the likelihood of the data $p(\mathbf{y} \mid \mathbf{e}, \mathbf{s}_1, \cdots, \mathbf{s}_{N_y})$ given $\mathbf{e}$, and the vectors $\mathbf{s}_i$, $1 \le i \le N_y$. To find $p(\mathbf{y} \mid \mathbf{x})$ we compute

$$
\begin{aligned}
p(\mathbf{y} \mid \mathbf{x}) &= \int_{\mathbf{e}} \int_{\mathbf{s}_1} \cdots \int_{\mathbf{s}_{N_y}} p(\mathbf{y} \mid \mathbf{e}, \mathbf{s}_1, \cdots, \mathbf{s}_{N_y}) p(\mathbf{e}, \mathbf{s}_1, \cdots, \mathbf{s}_{N_y} \mid \mathbf{x}) d\mathbf{s}_{N_y} \cdots d\mathbf{s}_1 d\mathbf{e} \\
&= \int_{\mathbf{e}} \int_{\mathbf{s}_1} \cdots \int_{\mathbf{s}_{N_y}} p(\mathbf{y} \mid \mathbf{e}, \mathbf{s}_1, \cdots, \mathbf{s}_{N_y}) p(\mathbf{s}_{N_y} \mid \mathbf{x}) \cdots p(\mathbf{s}_1 \mid \mathbf{x}) p(\mathbf{e} \mid \mathbf{x}) d\mathbf{s}_{N_y} \cdots d\mathbf{s}_1 d\mathbf{e}
\end{aligned}.
$$

We assume that the probability of any particular surface occurring in a scene is independent of the other surfaces in the scene, and that the probability of a surface occurring in a scene is independent of the illuminant. This lets us simplify to

$$
p(\mathbf{y} \mid \mathbf{x}) = \int_{\mathbf{e}} p(\mathbf{e} \mid \mathbf{x}) \left[\int_{\mathbf{s}_1} p(\mathbf{y}_1 \mid \mathbf{e}, \mathbf{s}_1) p(\mathbf{s}_1) d\mathbf{s}_1 \right] \cdots \left[\int_{\mathbf{s}_1} p(\mathbf{y}_{N_y} \mid \mathbf{e}, \mathbf{s}_{N_y}) p(\mathbf{s}_{N_y}) d\mathbf{s}_{N_y} \right] d\mathbf{e} \quad (2)
$$

where the notation $\mathbf{y}_i$ refers to the three entries of $\mathbf{y}$ that depend on the light reflected from the i^{th} surface.

We assume that the prior probability of a surface $\mathbf{s}$ occurring in a scene is described by a multivariate Normal distribution over the weights of a linear model (Brainard & Freeman, 1997; Zhang & Brainard, 2004). That is

$$\begin{aligned}\mathbf{s} &= \mathbf{B}_s \mathbf{w}_s \\ \mathbf{w}_s &= N(\mathbf{u}_s, \mathbf{K}_s)\end{aligned} \tag{3}$$

where $\mathbf{B}_s$ is a matrix with N_s columns each of which describes a spectral basis function, $\mathbf{w}_s$ is an N_s dimensional column vector, $\mathbf{u}_s$ is mean of $\mathbf{w}_s$, and $\mathbf{K}_s$ is the covariance of $\mathbf{w}_s$. We used $N_s = 3$ and for the main calculations we chose $\mathbf{B}_s$, $\mathbf{u}_s$, and $\mathbf{w}_s$ as described by Brainard and Freeman (1997). Given this surface prior, each of the integrals in brackets in (2) may be approximated analytically using the methods described by Freeman and Brainard (Freeman, 1993; Brainard & Freeman, 1994).

We also assume that illuminant spectra are characterized by a linear model, so that

$$\mathbf{e} = \mathbf{B}_e \mathbf{w}_e \tag{4}$$

where we chose $\mathbf{B}_e$ to specify that CIE three-dimensional linear model for daylight. Given this linear model constraint, the chromatiticy of an illuminant determines the linear model weights up to a free multiplicative scale factor through a straightforward colorimetric calculation: $\mathbf{w}_e = a\, f(\mathbf{x})$. We assumed that the distribution over the scale factor was uniform, so that

$$p(\mathbf{e} \mid \mathbf{x}) = \begin{cases} c_e, & \mathbf{e} = \mathbf{B}_e \mathbf{w}_e \textit{ and } \mathbf{w}_e = a\, f(\mathbf{x}) \textit{ for some } a \\ 0, & \textit{otherwise} \end{cases}.$$

We also set $p(\mathbf{e} \mid \mathbf{x}) = 0$ for any illuminant that had negative power in any wavelength band. Our assumption about illuminant spectra allows us to evaluate the outer integral in (2) numerically.

Given that we can evaluate the likelihood, computation of the posterior then requires only that we define a prior probability over illuminant chromaticities. We assumed that

$\mathbf{x} \sim N(\mathbf{u}_x, \mathbf{K}_x)$.

For the main calculations, we fixed $\mathbf{u}_x$ at the chromaticity of CIE illuminant D65. Thus the illuminant prior has three free parameters, the independent entries of covariance matrix $\mathbf{K}_x$.

To find $\tilde{\mathbf{x}}$, we used numerical search and found the value of $\mathbf{x}$ that maximized the posterior (Eq. 1).

To run the algorithm on images, 24 points were selected at random from the image data. In some conditions, points were selected uniformly from the image. In other conditions, points were drawn using a Gaussian weighting function of specified standard deviation and centered on the test patch. In the experiments, the test patch chromaticity was randomized at the start of every achromatic adjustment. To simulate this, the same randomization rule was used when drawing points that fell within the test patch. For each condition, the 24 points were drawn 10 separate times and the algorithm was run for each of the 10 sets. The resulting illuminant estimates were then averaged to produce the final estimate.

Calculation of Inferred Achromatic Surface

To predict achromatic chromaticities from illuminant chromaticities, we find an achromatic surface such that the chromaticity of the light reflected from this surface under each illuminant best predicts the achromatic chromaticities.

Let $\mathbf{x}_i^a$ represent the achromatic chromaticity for the i^{th} scene, and let $\tilde{\mathbf{x}}_i$ represent the illuminant chromaticity estimated for the i^{th} scene. We assume that illuminant spectra are described by the CIE three-dimensional linear model for daylight, as in Eq. 4 above, and that the achromatic surface's reflectance function is described by the three-dimensional linear model for surfaces as in Eq. 3. Given the linear model constraint on illuminants, we can compute the spectrum of each scene's estimated illuminant ($\tilde{\mathbf{e}}_i$) the from its chromaticity $\tilde{\mathbf{x}}_i$, up to an undetermined scale factor. Given $\tilde{\mathbf{e}}_i$ and any choice of the linear model weights for the achromatic surface ($\mathbf{w}_s^a$), we can compute the spectrum of the light reflected from the achromatic surface, again up to a undetermined factor. This then yields the chromaticity of the reflected light $\tilde{\mathbf{x}}_i^a$ under the estimated illuminant. Note that the computed chromaticity does not depend on the undetermined scale factor.

The chromaticities $\tilde{\mathbf{x}}_i^a$ serve as our prediction of the achromatic chromaticities $\mathbf{x}_i^a$. We used numerical search over the achromatic surface weights $\mathbf{w}_s^a$ to find the weights that minimized the fit error $\varepsilon_a = \sqrt{\sum_{i=1}^{N_{scenes}} \left\| \mathbf{x}_i^a - \tilde{\mathbf{x}}_i^a \right\|^2 \Big/ N_{scenes}}$. Because the predictions $\tilde{\mathbf{x}}_i^a$ depend only on the relative surface reflectance function, we fixed the first entry of $\mathbf{w}_s^a$ to be 1 and searched over the remaining two entries.

The same procedure was used to predict achromatic chromaticities from actual illuminant chromaticities. In the example calculations involving a subset of images (chromaticity plots in Figures 2, 6, and 7), the weights $\mathbf{w}_s^a$ were chosen so that the prediction error for the standard context was zero, rather than to minimize the prediction error across all seventeen contexts.

Results

Measurements of Human Color Constancy

The data consist of achromatic settings made in the context of rich, semi-naturalistic images (Delahunt & Brainard, 2004). The images at the top right of Figure 2 show four of the seventeen *contextual images* used in the experiments. The first three images are graphics simulations of the same collection of objects rendered under three different illuminants. The lefthand image was rendered under a typical daylight (CIE D65, chromaticity shown as open black circle), while the second image was rendered under a more yellowish daylight (spectrum constructed from CIE linear model for daylights, chromaticity shown as open blue circle). The third image was also a rendering of the same set of surfaces, but the illuminant had a chromaticity atypical of daylight (chromaticity shown as open red circle.) The righthand image was rendered under essentially the same illuminant as the second image (open green circle), but a different background surface was simulated. This background was chosen so that the light reflected from it matched that reflected from the background in the lefthand image.

Figure 2 about here.

Stereo pairs corresponding to each image were displayed on a computer-controlled haploscope. Left- and right-eye images for each pair were obtained by rendering the scene from two viewpoints. For each pair, observers adjusted a test patch (location shown as a black rectangle in each image in Figure 2) so that it appeared achromatic (i.e. gray). Achromatic adjustment has been used extensively to characterize color appearance (Helson & Michels, 1948; Brainard, 1998; Kraft & Brainard, 1999; Chichilnisky & Wandell, 1996; Bauml, 2001; Yang & Maloney, 2001). It provides an excellent first order characterization of how the visual system has adapted to the context provided by the image. Speigle and Brainard showed that achromatic adjustments made in two separate contexts may be used to predict asymmetric color matches (Speigle & Brainard, 1999).

The results of the experiment may be summarized by the chromaticity of the stimulus that appeared achromatic in each contextual image. The solid black, blue, red and green circles in the CIE u'v' chromaticity diagram in Figure 2 indicate the achromatic chromaticities corresponding to each image. This representation characterizes the spectrum reaching the observer from the test patch (i.e. the proximal stimulus), not the reflectance properties of the simulated surface. It is clear that the achromatic chromaticity varies with context.

To relate these data to constancy, first consider the context defined by the lefthand image in Figure 2. We will call this the *standard context*. Constancy imposes no necessary relation between the chromaticity of the simulated illuminant (in this case D65) and the corresponding achromatic chromaticity. We can, however, use the achromatic and illuminant chromaticities to determine the spectral reflectance function of an *inferred achromatic surface* (see METHODS). Such a reflectance function is shown in Figure 2. Here this function is chosen so that when the simulated illuminant reflects from it, the chromaticity of the reflected light matches the measured achromatic chromaticity. This equality is indicated by the overlay of the solid black circle and the small open black circle.

For a color constant observer, a surface that appears achromatic must continue to appear achromatic in any other context. We can compute the chromaticity of the light that would be reflected from the inferred achromatic surface in the contexts defined by the other images (see METHODS). These are shown as the small blue, red, and green open circles in Figure 2. Given the data from the lefthand image, a perfectly color constant observer must judge these chromaticities to be achromatic in the context their respective images.

The constancy predictions for the middle two contextual images (shown as small open blue and red circles) lie in the general vicinity of the corresponding achromatic data (close blue and red circles), but there are clear deviations in each case. To evaluate the magnitude of the deviation, bear in mind that for an observer with no constancy the achromatic chromaticity would not vary with the illuminant and would thus overlay the data for the standard context (solid black circle.) In each of these two cases, the data are closer to the constancy predictions than to the prediction for no constancy. This result is typical for studies where *only* the illuminant is varied, and is the basis for assertions in the literature that the human visual system is approximately color constant (Brainard, 2004).

The result for the righthand contextual image is different. The simulated scene that produced this image has essentially the same illuminant as the second image, but a different background surface. Here the achromatic chromaticity (solid green circle) falls closer to the no constancy prediction (solid black circle) than to the constancy prediction (small open green circle). We have shown this basic result previously using stimuli that consist of real illuminated surfaces (Kraft & Brainard, 1999). Intuitively, less constancy is shown because the change in background surface silences the cue provided by local contrast. The fact that the achromatic chromaticity is not exactly the same as in the standard context (solid green and solid black circles differ) indicates that local contrast is not the only cue mediating constancy, again replicating our earlier results obtained with real surfaces (Kraft & Brainard, 1999).

We can generalize the logic of Figure 2 to make constancy predictions for the achromatic chromaticities measured in all seventeen contexts used by Delahunt and Brainard. We used numerical parameter search to find a single inferred achromatic surface, common to all seventeen contexts (see METHODS). This surface is chosen to minimize the sum-of-squared errors between the u'v' chromaticities of the light reflected from it in each context and the corresponding achromatic chromaticities. The best fit to the data is summarized in Figure 3. As one might expect from the example data shown in Figure 2, the data set is poorly fit by a model that embodies the assumption that the visual system is perfectly color constant.

Figure 3 about here.

Figures 2 and 3 show that any theory of surface color appearance that simply characterizes the overall degree of constancy (e.g. "humans are approximately color constant", "human color constancy is poor", "humans are 83% color constant") is doomed: the degree of constancy depends critically on how the context is manipulated. Since we cannot describe human performance as simply color constant, or as not color constant, our goal is to develop a method for predicting the achromatic setting given the contextual image. Here we pursue a model based on an analysis of the computations required to achieve color constancy.

Equivalent Illuminant Model

Figure 4 illustrates the concept underlying our modeling approach. The observer views an image formed when illumination reflects off objects. The illuminant is characterized by its spectral power distribution $E(\lambda)$, and each surface is characterized by its reflectance function $S(\lambda)$. The spectrum of the light reaching the eye is given by $C(\lambda) = E(\lambda)S(\lambda)$. If the visual system has access to an estimate of the illuminant, $\tilde{E}(\lambda)$, it is straightforward for it to produce a representation of object color based on an estimate of surface reflectance $\tilde{S}(\lambda) = C(\lambda) / \tilde{E}(\lambda)$.* As long as $\tilde{E}(\lambda) \approx E(\lambda)$, this representation will be stable across changes in context. To put it another way, it is easy for a visual system with access to the physical illuminant $E(\lambda)$ to be approximately color constant.• If the visual system applies this strategy with illuminant estimates that deviate substantially from the actual illuminants, however, there will be commensurate deviations from constancy.

Figure 4 about here.

We assume in our modeling that the visual system tries to achieve color constancy, and that it does so through an inverse optics computation that is

---

* The actual estimation is more complicated, as the visual system does not have direct access to $C(\lambda)$ but instead must use the responses of the L-, M-, and S-cones to make the estimate. There are well-established methods for doing so (Wandell, 1987; Brainard, 1995). The tilde in the notation distinguishes perceptual from physical quantities.

• Even for a visual system with access to the physical illuminant, errors in constancy may still occur because of metamerism: two surfaces that produce the same cone responses correct up to possible errors in its estimate of the illuminant $\tilde{E}(\lambda)$. It should not be surprising that illuminant estimates are sometimes in error: the reason that constancy is intriguing rests on the fundamental difficulty of achieving it across all possible scenes. We refer to $\tilde{E}(\lambda)$ as the equivalent illuminant, since in the model it replaces the physical illuminant in the visual system's calculations. More generally, we refer to the class of models developed here as *equivalent illuminant models*.

Equivalent illuminant models have been proposed previously (Beck, 1959; Gilchrist & Jacobsen, 1984; Logvinenko & Menshikova, 1994; Maloney & Yang, 2001; Rutherford & Brainard, 2002). One empirical approach to testing the equivalent illuminant idea in the context of lightness is to ask whether explicit judgments of illuminant intensity covary lawfully with explicit judgments of surface lightness This approach has led to mixed results (Beck, 1959; Logvinenko & Menshikova, 1994; Rutherford & Brainard, 2002), but as a whole cast doubt on the idea that *explicit* illuminant judgments tap an equivalent illuminant that may be used to predict surface lightness.

Another empirical approach is to ask whether the equivalent illuminant idea leads to a parametric model that predicts how surface appearance varies with some ancillary stimulus variable (e.g. chromaticity of reflected

---

under one illuminant may produce different cone responses under a changed illuminant (Brainard, 1995).

light, surface slant). This has proven quite successful in a variety of empirical contexts (Speigle & Brainard, 1996; Brainard, Brunt, & Speigle, 1997; Boyaci, Maloney, & Hersh, 2003; Brainard, Kraft, & Longère, 2003; Bloj *et al.*, 2004; Boyaci, Doerschner, & Maloney, 2004; Doerschner, Boyaci, & Maloney, 2004). This approach, however, does not specify how the equivalent illuminant is determined by the image. Rather, the surface appearance data themselves are fit to identify the parameters of an *implicit* equivalent illuminant.

Here we develop the equivalent illuminant concept further by asking whether we can predict human performance from equivalent illuminants obtained directly through an illuminant-estimation algorithm. We used a Bayesian algorithm that we have described previously (Brainard & Freeman, 1997). The algorithm's performance depends on specifying a prior over illuminant spectral power distributions and surface reflectance functions. The output of the algorithim is an estimate of the chromaticity of scene illuminant.

Figure 5 illustrates how we characterized the prior over illuminants. The left panel (green dots) shows the u'v' chromaticity coordinates of 10760 daylights measured by DiCarlo and Wandell (2000). Consistent with previous reports, these cluster along the CIE daylight locus (black line in both panels). To capture this regularity in broad terms, we modeled the prior probability of illuminant chromaticities as a bivariate Normal distribution (see METHODS.) The blue ellipse in the right panel shows an iso-probabililty contour of the Normal distribution we chose to capture the broad statistics of the daylight chromaticity distribution. We choose the mean of this distribution as the chromaticity of CIE illuminant D65, and the covariance matrix as that of the DiCarlo and Wandell measurements. We refer to this illuminant prior as the *daylight prior*. The prior over scene surfaces and other details of the algorithm are provided in METHODS.

Figure 5 about here.

We applied the Bayesian algorithm with the daylight prior to the seventeen images used by Delahunt and Brainard (2004) to obtain estimates of the illuminant chromaticities. We then repeated the analysis used to make the constancy predictions shown in Figure 3 above, but with the algorithm's estimates in place of the actual illuminant chromaticity for each context.[*•] The resulting predictions are shown in Figure 6. The top panel shows predictions for the same example contexts as in Figures 2. The middle and bottom panels summarize the prediction quality in the same format as Figure 3. The predictions are, on aggregate, improved slightly from those obtained using the actual chromaticities of the scene illuminants. The value of the error measure ε_a is 0.0157 for the predictions based on the

[*•] As described in METHODS, the calculation of equivalent surface and achromatic predictions depends on the illuminant spectrum only through its chromaticity. This is possible because we assume that the CIE linear model for daylights describes the actual illuminant spectra.

actual illumiant chromaticities and 0.0114 for the predictions based on the Bayesian estimates obtained with the daylight prior.

The reason for the net improvement is that for some scenes the algorithm fails to estimate the physical illuminant correctly, and these failures are generally consistent with human performance. Two examples of this may be seen by comparing Figure 2 and the top panel of Figure 6: the predictions shown by the small open blue and green circles are better in the latter. The prediction improvement is largest for the case where the background surface in the image was varied to silence local contrast as a cue to the illuminant (green circles), although the prediction for this case still deviates from the data.

On the other hand, not all of the predictions are improved: the small open red circle lies further from the point it predicts in the top panel of Figure 6 than in Figure 2. Here the Bayesian prediction lies much closer to the daylight locus than the corresponding datum. Intuitively, this occurs because the illuminant prior we used places very little probability weight in the vicinity of the scene illuminant used in the experiment. Overall, the model with the daylight prior does not adequately describe the data.

Figure 6 about here.

The poor prediction performance obtained with the daylight prior for some scenes is consistent with Delahunt and Brainard's (2004) conclusion that human color constancy is robust with respect to atypical illuminant changes. Within a Bayesian framework such robustness can be described through the use of a broader illuminant prior distribution, and we wondered whether allowing this possibility could improve the predictions of the equivalent illuminant model. A second non-realistic aspect of the calculations described above is that the 24 sample points provided to the algorithm were drawn uniformly from the whole image. It seems likely that observers spend more time looking in the vicinity of the test patch when making an achromatic adjustment, and we also explored the effect of drawing samples according to a spatial Gaussian weighting function centered on the test patch.

We parameterized the illuminant prior through the three independent entries of its covariance matrix $\mathbf{K}_x$ and repeated the calculations for many choices of $\mathbf{K}_x$. We identified the choice of $\mathbf{K}_x$ that led to the best prediction results (minimum ε_a). We then varied the spatial sampling, choosing the standard deviation of the Gaussian weighting function that minimized ε_a. The isoprobability contour corresponding to the best $\mathbf{K}_x$ is shown as the red ellipse in Figure 5. This distribution is considerably broader than the daylight prior. The best standard deviation for the Gaussian weighting function was 4° of visual angle. Figure 7 shows that the quality of the resulting fit is very good. The predictions for the example conditions (top panel) all lie close to the measurements, while the points in the summary graphs (middle and bottom panels) are generally near the diagonal. Most of the improvement in ε_a was due to the change in prior, not in spatial sampling.[••]

Figure 7 about here.

Discussion

Insight about perception is often available from consideration of the computational task faced by the perceptual system (Marr, 1982; Barrow & Tenenbaum, 1978). The premise is that if we formulate explicit algorithms that can accomplish this task, then we can use these as the basis of models for human performance. In the case of color constancy, McCann et al. (1976) took this general approach by adopting Land's retinex algorithm (Land, 1964; Land, 1986; see Brainard & Wandell, 1986) as a candidate model. The retinex algorithm does not explicitly estimate the illuminant, and focused tests of its core properties (e.g. Kraft & Brainard, 1999; Delahunt & Brainard, 2000) show clear deviations between its predictions and human performance.

We also use the computational approach to develop a quantitative model of surface appearance. Here the predictions are derived from a principled

[••] The value of ε_a with best illuminant prior and spatial weighting was 0.0044. The value of ε_a obtained with the best illuminant prior when points were drawn from the entire image was 0.0057. We also explored whether any choice of spatial sampling could produce a similar fit with the daylight prior, and whether varying the surface prior parameters allowed a good fit with the daylight prior. Across all variations we explored, the minimum value of ε_a obtained with the daylight prior was 0.0110, over twice that obtained with the broadened prior.

analysis of how the illuminant can be estimated from image data, and this analysis is connected to the data through the equivalent illuminant principle.

In contrast with earlier work with equivalent illuminant models, both from our lab (Brainard, Brunt, & Speigle, 1997; Speigle & Brainard, 1996; Bloj *et al.*, 2004) and by others (Boyaci, Maloney, & Hersh, 2003; Boyaci, Doerschner, & Maloney, 2004; Doerschner, Boyaci, & Maloney, 2004), the equivalent illuminants here were not derived to fit individually to the experimental data from the corresponding condition but were instead obtained through the application of an independent algorithm. We have previously argued that modeling of color appearance may be fruitfully pursued in two steps: first determine the appropriate parametric form to account for context effects, then determine how the parameters are determined by the image context (Brainard & Wandell, 1992; Brainard, 2004). Prior work on equivalent illuminant models address the first step. The present work tackles the question of how the parameters are determined.

The model fitting was restricted to varying a few parameters that govern the overall behavior of the algorithm, and the same parameters were used for all the experimental conditions. The good quantitative fit obtained illustrates the promise of the approach. The introduction of an explicit algorithm represents a novel stage of theoretical development and increases the force of the general modeling approach. Parameterizing the prior distribution to allow fitting of a Bayesian model to human data has also proved effective for modeling of shape (Mamassian, Landy, & Maloney, 2002) and motion (Weiss, Simoncelli, & Adelson, 2002) perception. In that they provide a principled benchmark against which to compare human performance, these Bayesian models of appearance may be understood as close relatives of the ideal observer models that have been highly successful in clarifying how vision detects and discriminates signals (Green & Swets, 1966; Geisler, 1989).

In addition to providing a predictive quantitative model, the parametric fit allows a succinct summary of human performance across all the conditions, in the form of the illuminant prior and inferred achromatic surface. The derived equivalent surfaces are close to spectrally flat, consistent with the properties of achromatic reference surfaces used in the photographic industry. On the other hand, the derived illuminant prior is considerably broader than would be indicated by measurements of natural daylight.

The discrepancy is interesting, and the reasons for it are unclear. A different distributional form of the illuminant prior, perhaps one with heavier tails, might allow reconciliation of the daylight measurements and derived prior. Or direct measurements of daylight spectra may not be the appropriate database from which to derive an illuminant prior: in natural scenes light often reaches surfaces after reflecting from other surfaces in the scene. Extensive measurements of the spectra of light actually reaching objects in typical scenes are not yet available. A final possibility is that formulating the illuminant estimation problem with a different loss function (Brainard & Freeman, 1997; Mamassian, Landy, & Maloney, 2002) would lead to better performance with the daylight prior.

Although the fits obtained here are good, the experimental conditions were not chosen to test the present model. Future tests will be most forceful if they contain conditions designed to probe it sharply. Kraft and Brainard (Kraft & Brainard, 1999) emphasized that a good way to do this is to choose contextual images where the scene illuminant differs while the model at hand predicts no difference in performance.

The general analysis presented here is not specific to the Brainard/Freeman algorithm, in that the same logic may be used to link any such algorithm (Buchsbaum, 1980; Maloney & Wandell, 1986; Funt & Drew, 1988; D'Zmura & Iverson, 1993a; D'Zmura & Iverson, 1993b; Finlayson, Hubel, & Hordley, 1997) to the data. Within the general framework presented here, differentiating between candidate algorithms again requires careful choice of experimental conditions, so that the data set includes conditions where different algorithms make substantially different predictions.

The algorithm we used does not take advantage of information carried by the geometric structure of the scene, and the experimental manipulations used were spectral rather than geometric. Much current experimental work on human color and lightness constancy focuses on such geometric manipulations (Bloj, Kersten, & Hurlbert, 1999; Ripamonti *et al.*, 2004; Williams, McCoy, & Purves, 1998; Adelson, 1999; Gilchrist *et al.*, 1999), and several authors have modeled this type of data with equivalent illuminant models where the illuminant parameters are fit directly to the data in each condition (Boyaci, Maloney, & Hersh, 2003; Boyaci, Doerschner, & Maloney, 2004; Bloj *et al.*, 2004; Doerschner, Boyaci, & Maloney, 2004). As algorithms for achieving constancy with respect to geometric manipulations become available, the principles underlying our present work should allow these algorithms to be gracefully integrated into the modeling effort.

The model reported here provides a functional description and does not speak directly to the neural mechanisms that mediate color constancy and color appearance. A great deal is known, however, about the early stages of chromatic adaptation (Stiles, 1967; Jameson & Hurvich, 1972; Wyszecki, 1986; Brainard, 2001). An important challenge for the current work remains to understand how known neural mechanisms might implement the sort of calculations that we have performed using a digital computer. We have begun to consider this question by examining how well parametric models of adaptation can adjust to changes in illumination and the composition of surfaces in a scene (Björnsdotter Abrams, Hillis, and Brainard, submitted for review).

Gilchrist and colleagues have argued (Gilchrist *et al.*, 1999) that the key to understanding the perception of color and lightness is to model deviations from constancy (in their terms, to model 'errors' in perception). They conclude that a computational approach is unlikely to provide a satisfactory account, because of the many such deviations observed in the literature. We close by noting that the model presented here, which is heavily motivated by computational considerations, successfully accounts for both conditions where the visual system is color constancy and where it is not. Our experimental conditions differ substantially from those considered by Gilchrist et al., but it is clear that there is no fundamental constraint that prevents a computationally motivated model from accounting for both successes and failures of constancy.

Acknowledgements

This work was supported by NIH EY 10016 and a gift from Agilent, Inc. An early version of the theory developed here was outlined in Brainard et al. (2003) and the work in close to final form was presented at the 2003 COSYNE meeting. We thank J. DiCarlo and B. Wandell for providing us with their spectral measurements of daylight, and M. Kahana and the Penn Psychology Department for providing access to a computer cluster.

References

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

Adelson, E. H., & Pentland, A. P. (1996). The perception of shading and reflectance. In Knill, D. & Richards, W. (Eds.), *Visual perception: Computation and psychophysics* (pp. 409-423). New York: Cambridge University Press.

Barrow, H. G., & Tenenbaum, J. M. (1978). Recovering intrinsic scene characteristics from images. In Hanson, A. R. & Riseman, E. M. (Eds.), *Computer vision systems*. New York: Academic Press.

Bauml, K. H. (2001). Increments and decrements in color constancy. *J Opt Soc Am A Opt Image Sci Vis, 18*(10), 2419-2429.

Beck, J. (1959). Stimulus correlates for the judged illumination of a surface. *Journal of Experimental Psychology, 58*(4), 267-274.

Berger, T. O. (1985). *Statistical decision theory and bayesian analysis*. New York: Springer-Verlag.

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

Bloj, M., Ripamonti, C., Mitha, K., Greenwald, S., Hauck, R., & Brainard, D. H. (2004). An equivalent illuminant model for the effect of surface slant on perceived lightness. *Journal of Vision, submitted*.

Boyaci, H., Doerschner, K., & Maloney, L. T. (2004). Perceived surface color in binocularly viewed scenes with two light sources differing in chromaticity. *Journal of Vision, 4*, 664-679.

Boyaci, H., Maloney, L. T., & Hersh, S. (2003). The effect of perceived surface orientation on perceived surface albedo in binocularly viewed scenes. *Journal of Vision, 3*, 541-553.

Brainard, D. H. (1995). Colorimetry. In Bass, M. (Ed.), *Handbook of optics: Volume 1. Fundamentals, techniques, and design* (pp. 26.21-26.54). New York: McGraw-Hill.

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. (2001). Color vision theory. In Smelser, N. J. & P.B., B. (Eds.), *International encyclopedia of the social and beahioral sciences* (Vol. 4, pp. 2256-2263). Oxford, U.K.: Elsevier.

Brainard, D. H. (2004). Color constancy. In Chalupa, L. & Werner, J. (Eds.), *The visual neurosciences* (Vol. 1, pp. 948-961). Cambridge, MA, U.S.A.: 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. (1994). *Bayesian method for recovering surface and illuminant properties from photoreceptor responses.* Paper presented at the SPIE Conference on Human Vision, Visual Processing, and Display, San Jose, CA.

Brainard, D. H., & Freeman, W. T. (1997). Bayesian color constancy. *Journal of the Optical Society of America A, 14*(7), 1393-1411.

Brainard, D. H., Kraft, J. M., & Longère, P. (2003). Color constancy: Developing empirical tests of computational models. In Mausfeld, R. & Heyer, D. (Eds.), *Colour perception: Mind and the physical world* (pp. 307-334). Oxford: Oxford University Press.

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. (1992). Asymmetric color-matching: How color appearance depends on the illuminant. *Journal of the Optical Society of America A, 9*(9), 1433-1448.

Buchsbaum, G. (1980). A spatial processor model for object colour perception. *Journal of the Franklin Institute, 310*, 1-26.

Chichilnisky, E. J., & Wandell, B. A. (1996). Seeing gray through the on and off pathways. *Visual Neuroscience, 13*(3), 591-596.

D'Zmura, M., & Iverson, G. (1993a). Color constancy. I. Basic theory of two-stage linear recovery of spectral descriptions for lights and surfaces. *J. Opt. Soc. Am. 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. *J. Opt. Soc. Am. A., 10*, 2166-2180.

Delahunt, P. B., & Brainard, D. H. (2000). Control of chromatic adaptation: Signals from separate cone classes interact. *Vision Research, 40*, 2885-2903.

Delahunt, P. B., & Brainard, D. H. (2004). Does human color constancy incorporate the statistical regularity of natural daylight? *Journal of Vision, 4*(2), 57-81.

DiCarlo, J. M., & Wandell, B. A. (2000). *Illuminant estimation: Beyond the bases.* Paper presented at the Eigth Color Imaging Conference, Scottsdale, AZ.

Doerschner, K., Boyaci, H., & Maloney, L. T. (2004). Human observers compensate for secondary illumination originating in nearby chromatic surfaces. *Journal of Vision, 4*, 92-105.

Finlayson, G. D., Hubel, P. H., & Hordley, S. (1997). *Color by correlation.* Paper presented at the IS&T/SID Fifth Color Imaging Conference: Color Science, Systems, and Applications, Scottsdale, AZ.

Freeman, W. T. (1993). *Exploiting the generic view assumption to estimate scene parameters.* Paper presented at the Proc. 4th Intl. Conf. Computer Vision, IEEE.

Funt, B. V., & Drew, M. S. (1988). *Color constancy computation in near-mondrian scenes using a finite dimensional linear model.* Paper presented at the IEEE Comp. Vis. & Pat. Recog, Ann Arbor, MI.

Geisler, W. S. (1989). Sequential ideal-observer analysis of visual discriminations. *Psychological Review, 96*(2), 267-314.

Geisler, W. S., & Kersten, D. (2002). Illusions, perception, and bayes. *Nature Neuroscience, 5*, 508-510.

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.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: John Wiley and Sons.

Helmholtz, H. v. (1866). *Treatise on physiological optics* (Southall, J. P. C., Trans. 2nd ed. Vol. 1962). New York: Dover.

Helson, H., & Michels, W. C. (1948). The effect of chromatic adaptation on achromaticity. *Journal of the Optical Society of America, 38*, 1025-1032.

Jameson, D., & Hurvich, L. M. (1972). Color adaptation: Sensitivity, contrast, and afterimages. In Hurvich, L. M. & Jameson, D. (Eds.), *Handbook of sensory physiology* (Vol. VII/4, pp. 568-581). Berlin: Springer.

Kersten, D., & Yuille, A. (2003). Bayesian models of object perception. *Current Opinion in Neurobiology, 13*, 1-9.

Knill, D., & Richards, W. (Eds.). (1996). *Perception as bayesian inference*. Cambridge: Cambridge University Press.

Kraft, J. M., & Brainard, D. H. (1999). Mechanisms of color constancy under nearly natural viewing. *Proc. Nat. Acad. Sci. USA, 96*(1), 307-312.

Land, E. H. (1964). The retinex. In Reuck, A. V. S. d. & Knight, J. (Eds.), *Ciba foundation symposium on color vision* (pp. 217-227). Boston: Little, Brown and Company.

Land, E. H. (1986). Recent advances in retinex theory. *Vision Research, 26*, 7-21.

Lee, P. M. (1989). *Bayesian statistics*. London: Oxford 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*(9), 1007-1023.

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.

Maloney, L. T., & Yang, J. N. (2001). The illuminant estimation hypothesis and surface color perception. In Mausfeld, R. & Heyer, D. (Eds.), *Colour perception: From light to object*. Oxford: Oxford University Press.

Mamassian, P., Landy, M. S., & Maloney, L. T. (2002). Bayesian modelling of visual perception. In Rao, R. P. N., Olshausen, B. A. & Lewicki, M. S. (Eds.), *Probabilistic models of the brain: Perception and neural function* (pp. 13-36). Cambridge, MA: MIT Press.

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.

Purves, D., & Lotto, R. B. (2003). *Why we see what we do: An empirical theory of vision*. Sunderland, MA: Sinauer.

Rao, R. P. N., Olshausen, B. A., & Lewicki, M. S. (Eds.). (2002). *Probabilistic models of the brain: Perception and neural function*. Cambridge, MA: MIT Press.

Ripamonti, C., Bloj, M., Hauck, R., Mitha, K., Greenwald, S., Maloney, S. I., & Brainard, D. H. (2004). Measurements of the effect of surface slant on perceived lightness. *Journal of Vision, 4*, 747-763.

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

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*(3), 436-451.

Speigle, J. M., & Brainard, D. H. (1999). Predicting color from gray: The relationship between achromatic adjustment and asymmetric matching. *J. Opt. Soc. Am. A, 16*(10), 2370-2376.

Stiles, W. S. (1967). Mechanism concepts in colour theory. *Journal of the Colour Group, 11*, 106-123.

Wandell, B. A. (1987). The synthesis and analysis of color images. *IEEE Transactions on Pattern Analysis and Machine Intelligence, PAMI-9*, 2-13.

Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience, 598-604*.

Williams, S. M., McCoy, A. N., & Purves, D. (1998). An empirical explanation of brightness. *Proc. Natl. Acad. Sci. USA, 95*, 13301-13306.

Wyszecki, G. (1986). Color appearance. In Boff, K. R., Kaufman, L. & Thomas, J. P. (Eds.), *Handbook of perception and human performance* (Vol. Vol I: Sensory Processes and Perception, pp. 9.1-9.56). New York: John Wiley & Sons.

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

Zhang, X., & Brainard, D. H. (2004). *Bayesian color-correction method for non-colorimetric digital image sensors.* Paper presented at the 12th IS&T/SID Color Imaging Conference, Scottsdale, AZ.

Figures

Figure 1. Physical interaction of surfaces and illuminants. The images show the same house at two different times. The light reflected to the camera differs greatly because of changes in the illuminant. The squares above each image show the same location in each image and emphasize the physical effect. The images were taken by the first author. Automatic color balancing usually performed by the camera was turned off. Despite the physical change recorded by the camera, the house appeared approximately the same yellow at both times. Figure reproduced from Figure 61.1/Plate 36 of Brainard (2004).

Figure 2. Example achromatic adjustment results, and interpretation in terms of constancy. Average achromatic chromaticities set by seven observers in four contexts are shown as solid black, blue, red, and green circles. Data are shown in the standard CIE u'v' chromaticity diagram. Stimuli that plot to the same point in this diagram produce the same relative L-, M-, and S-cone isomerization rates but may have different intensities, while stimuli that plot to different points produce different relative isomerization rates. The chromaticities of the four scene illuminants are shown as large open circles. The contextual images are shown above the plot, with the symbols under each identifying the corresponding datum and illuminant. Error bars show +/- 1 standard deviation for data from the seven observers. Predictions for a color constant observer are shown as small open circles, as explained in the text. The spectral plot (x-axis: wavelength, y-axis: reflectance) shows the reflectance of the inferred achromatic surface used to make the constancy predictions.

Figure 3. Constancy predictions. Left panel shows predicted u′ chromaticity versus measured achromatic u′ chromaticity. Right panel shows predicted v′ chromaticity versus measured v′ chromaticity. Perfect prediction would be indicated if the data fell along the postive diagonals in both panels. The inset shows the reflectance spectrum of the achromatic surface inferred from the data. Error bars indicate standard deviations of the data across seven observers. The value of the fit error ε_a (see METHODS) is 0.0157.

Figure 4. Equivalent illuminant concept. See description in text.

Figure 5. Illuminant priors. The left panel shows the CIE u′v′ chromaticity of 10760 daylight measurements made by DiCarlo and Wandell (2000). The right panel shows an isoprobability contour of a bivariate Normal distribution chosen to capture the distribution of daylight chromaticities (blue ellipse). This is the daylight prior described in the text. The red ellipse shows an isoprobability contour of a bivariate Normal distribution that in conjunction with our Bayesian algorithm led to equivalent illuminants that provided good predictions of human performance. Both isoprobability contours are scaled so that they contain 90% of draws from their corresponding Normal distributions. For reference, the solid black line in both panels plots the CIE daylight locus for correlated color temperatures between 4000° K and 10000° K, and the filled black circle plots the chromaticity of CIE illuminant D65.

Figure 6. Model predictions with daylight prior. The top panel shows results and predictions for the same four scenes as in Figure 2. Same format as Figure 2. The middle and bottom panels summarize performance

in the same format as Figure 3. The value of the fit error ε_a (see METHODS) is 0.0114.

Figure 7. Model predictions with best prior and spatial weighting. Same format as Figure 6. The value of the fit error ε_a is 0.0044.

Brainard et al. Figure 1

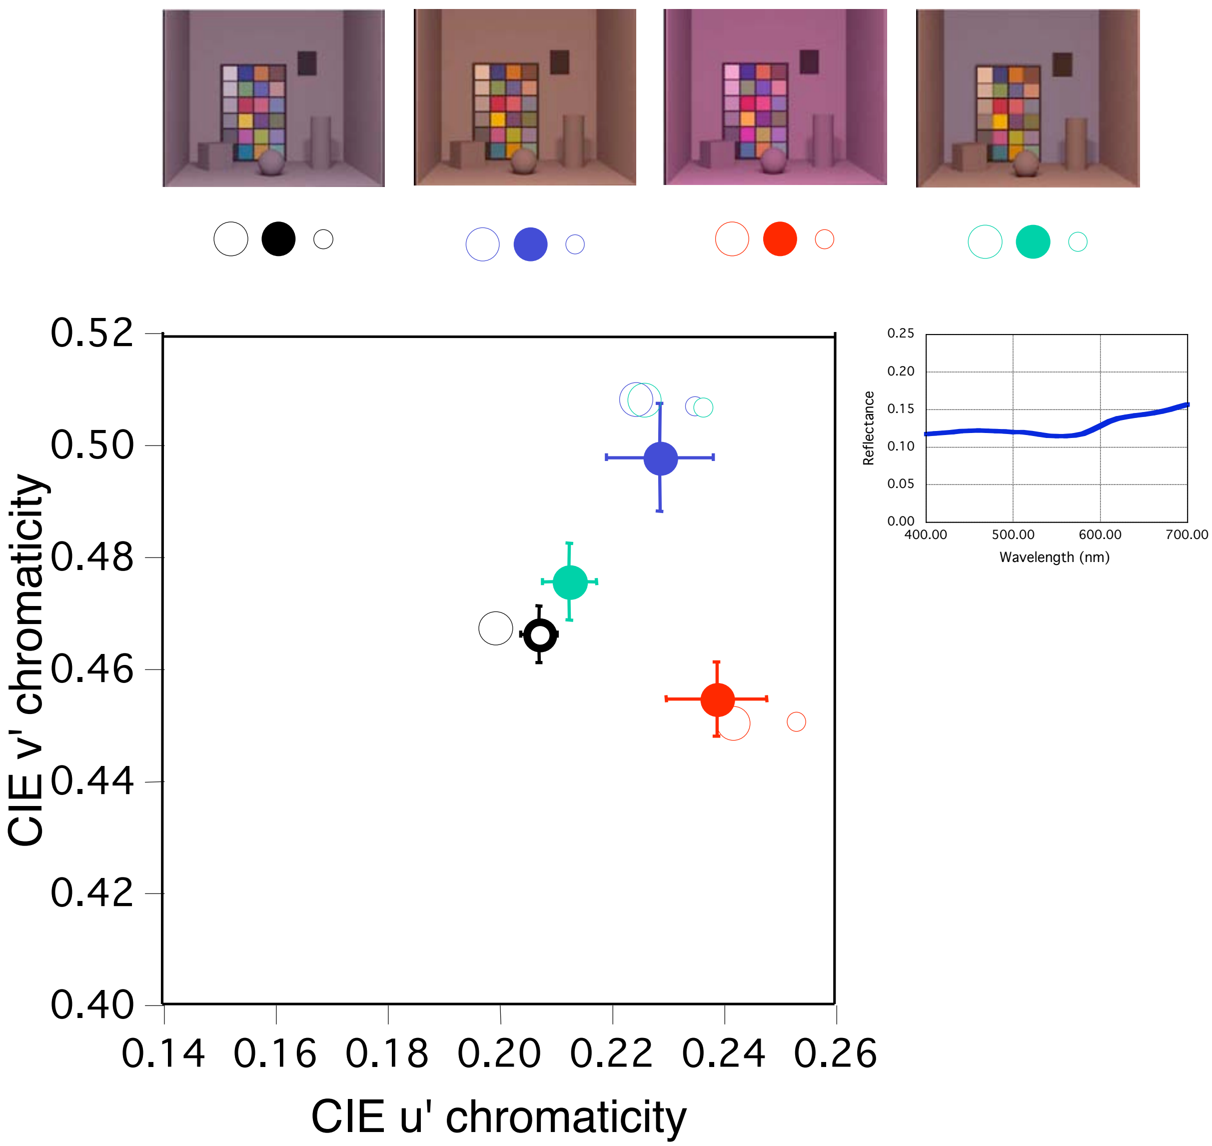

Brainard et al. Figure 2

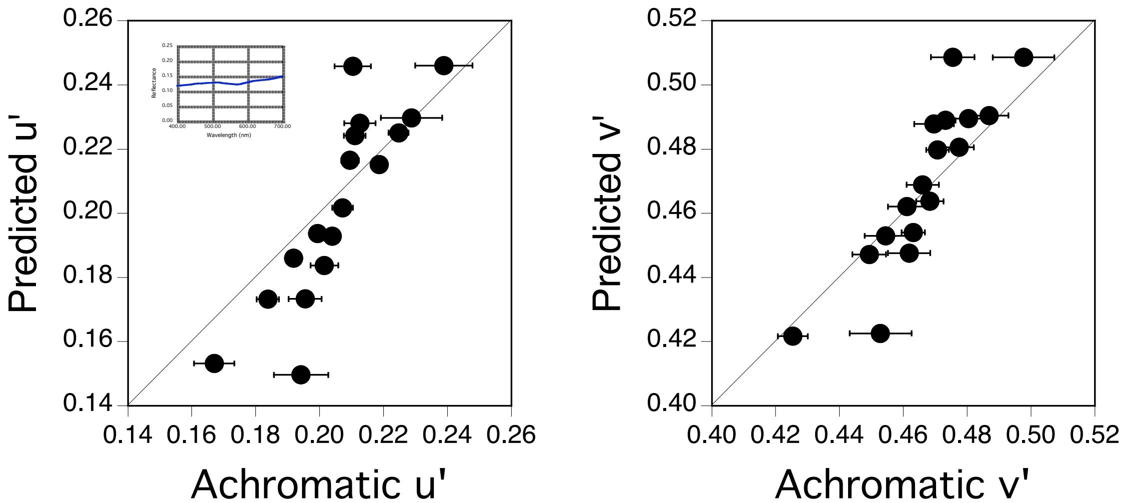

Brainard et al.

Figure 3

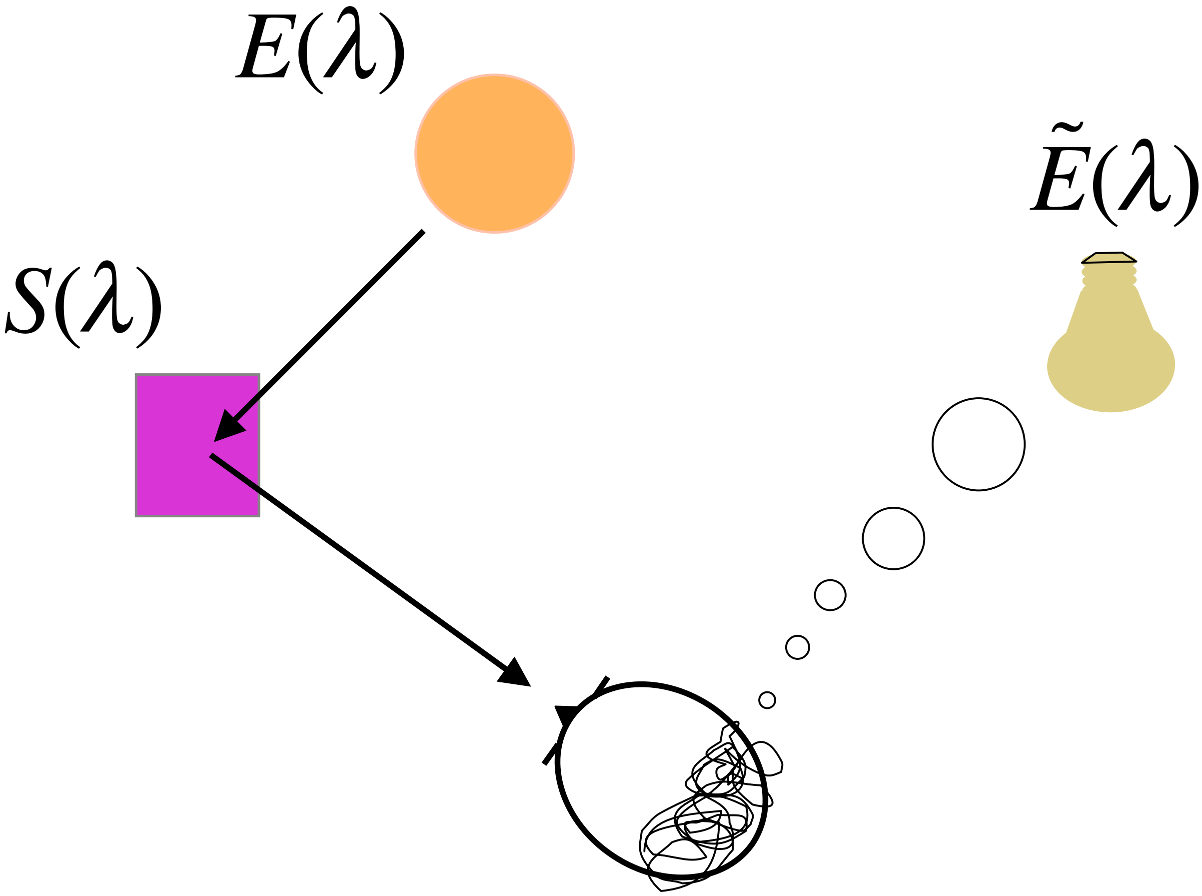
$E(\lambda)$
$S(\lambda)$
$\tilde{E}(\lambda)$

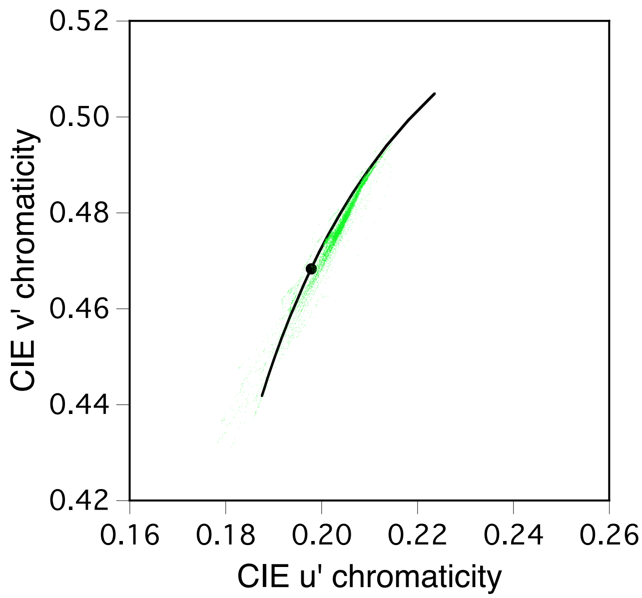

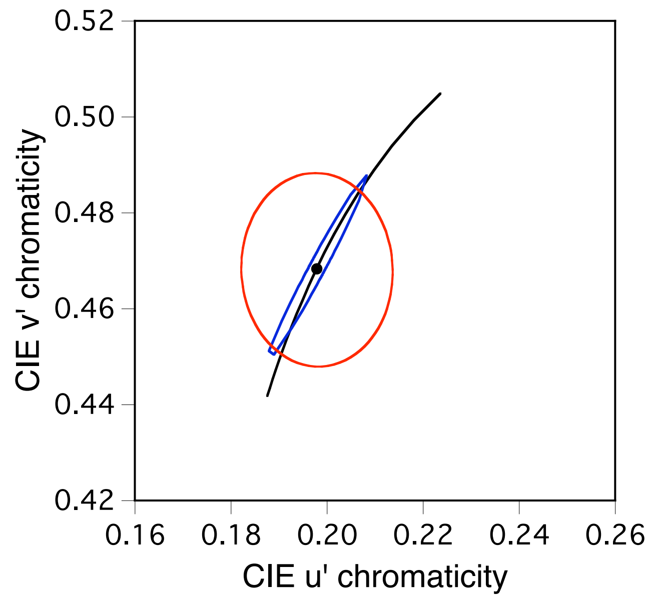

Brainard et al.

Figure 5

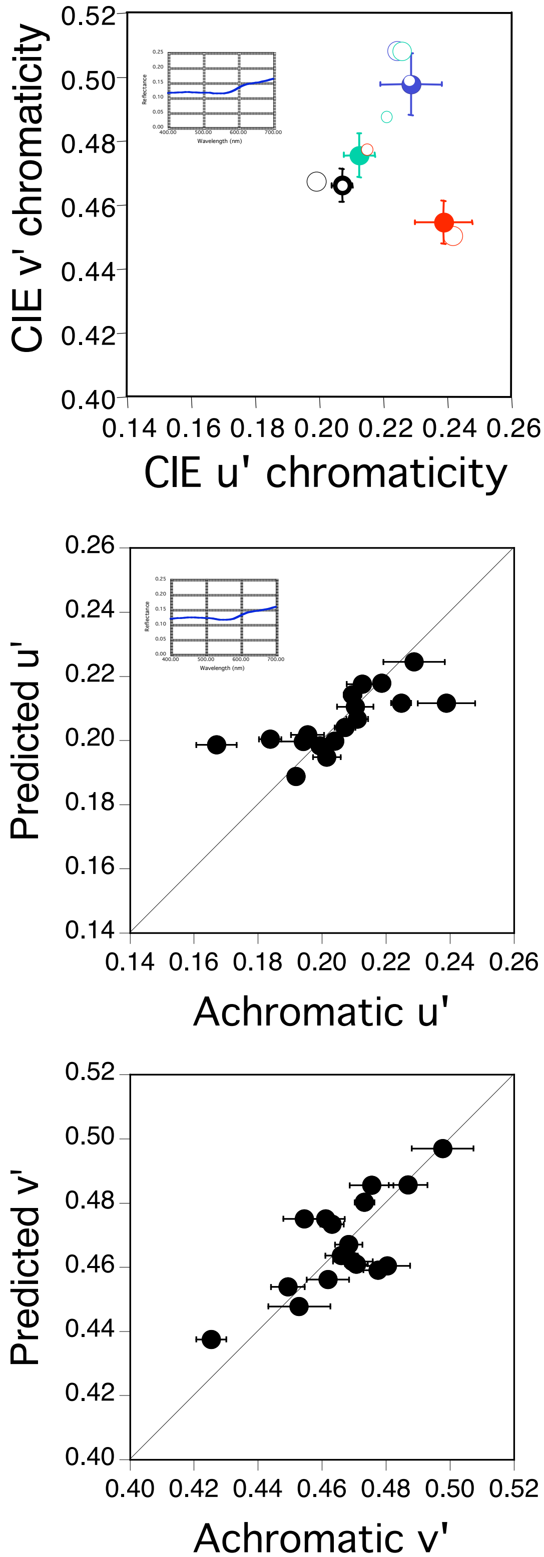

CIE v' chromaticity
CIE u' chromaticity
0.52
0.50
0.48
0.46
0.44
0.42
0.40
0.14 0.16 0.18 0.20 0.22 0.24 0.26
Predicted u'
Achromatic u'
0.26
0.24
0.22
0.20
0.18
0.16
0.14
Predicted v'
Achromatic v'
0.40 0.42 0.44 0.46 0.48 0.50 0.52

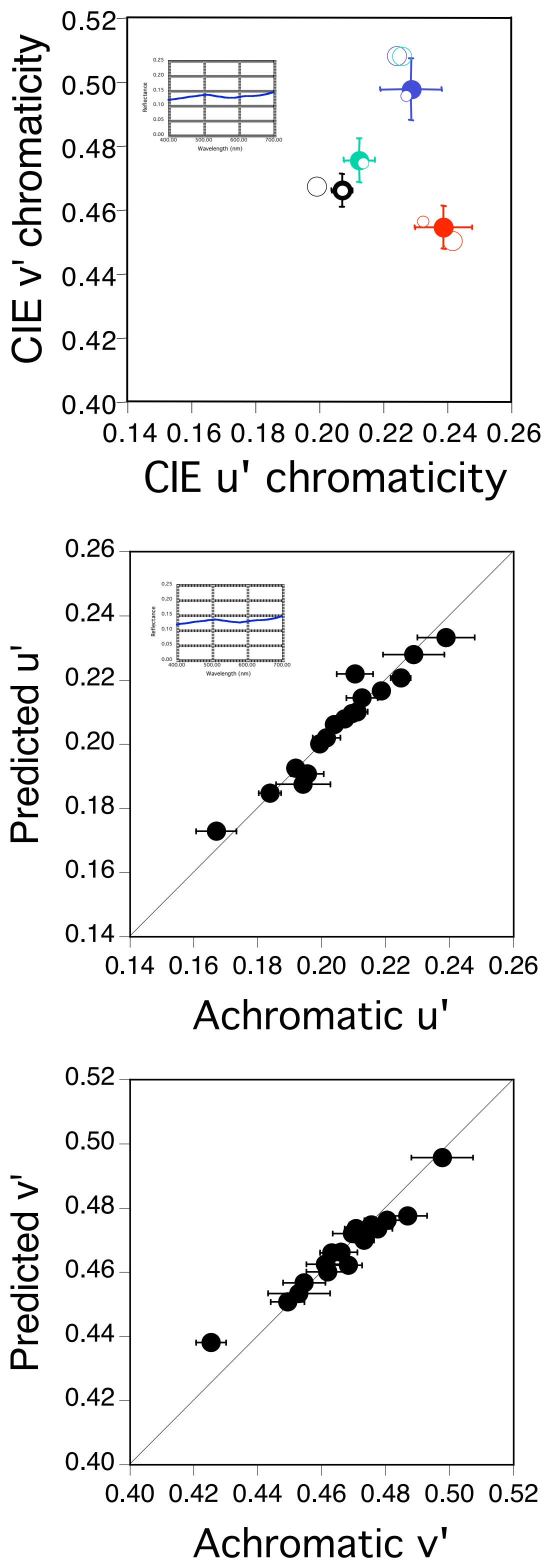

Brainard et al. Figure 7