

An Equivalent Illuminant Model for the Effect of Surface Slant on Perceived Lightness

M. Bloj

Department of Optometry, University of Bradford
Bradford, UK



C. Ripamonti

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

K. Mitha

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

R. Hauck

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

S. Greenwald

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

D. H. Brainard

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



Keywords: computational vision, equivalent illuminant, lightness constancy, scene geometry, surface slant

Abstract

In the companion paper (Ripamonti et al., 2004), we present data that measure the effect of surface slant on perceived lightness. Observers are neither perfectly lightness constant nor luminance matchers, and there is considerable individual variation in performance. This paper develops a parametric model that accounts for how each observer's lightness matches vary as a function of surface slant. The model is derived from consideration of an inverse optics calculation that could achieve constancy. The inverse optics calculation begins with parameters that describe the illumination geometry. If these parameters match those of the physical scene, the calculation achieves constancy. Deviations in the model's parameters from those of the scene predict deviations from constancy. We used numerical search to fit the model to each observers data. The model accounts for the diverse range of results seen in the experimental data in a unified manner, and examination of its parameters allows

interpretation of the data that goes beyond what is possible with the raw data alone. In particular, the model allows calculation of a constancy index that takes on a value of 0 for luminance matching and 1 for perfect constancy. Across our experiments, the average constancy index was 0.57.

Introduction

In the companion paper (Ripamonti et al., 2004), we present psychophysical data that measure how perceived surface lightness varies with scene geometry. In particular, we measured the effect of surface slant. The data indicate that observers take geometry into account when they judge surface lightness, but that there are large individual differences. This paper develops a quantitative model of our data. The model is derived from an analysis of the physics of image formation and of the computations that the visual system would have to perform to achieve lightness constancy. The model allows for failures of lightness constancy by supposing that observers do not perfectly estimate the lighting

geometry. Individual variation is accounted for within the model by parameters that describe each observer's representation of that geometry.

Figure 1 replots experimental data for three observers (HWK, EEP, and FGS) from Ripamonti et al. (2004). Observers matched the lightness of a standard object to a palette of lightness samples, as a function of the slant of the standard object. The data consist of the normalized relative match reflectance at each slant. If the observer had been perfectly lightness constant, the data would fall along a horizontal line, indicated in the plot by the red dashed line. If the observer were making matches by equating the reflected luminance from the standard and palette sample, the data would fall along the blue dashed curves shown in the figure. The complete data set demonstrates reliable individual differences ranging from luminance matches (e.g. HWK) to approximations of constancy (e.g. FGS). Most of the observers, though, showed intermediate performance (e.g. EEP).

Figure 1 about here

Figure 1. Normalized relative matches, replotted from Ripamonti et al. (2004). Data are for Observer HWK (Paint instructions), Observer EEP (Neutral instructions), and Observer FGS (Neutral instructions). See companion paper for experimental details.

Given that observers are neither perfectly lightness constant nor luminance matchers, our goal is to develop a parametric model that can account for how each observer's matches vary as a function of slant. Establishing such a model offers several advantages. First, individual variability may be interpreted in terms of variation in model parameters, rather than in terms of the raw data. Second, once a parametric model is established, one can study how variations in the scene affects the model parameters (c.f. Krantz, 1968; Brainard & Wandell, 1992)). Ultimately, the goal is to develop a theory that allows

prediction of lightness matches across a wide range of scene geometries.

A number of broad approaches have been used to guide the formulation of quantitative models of context effects. Helmholtz (Helmholtz, 1896) suggested that perception should be conceived of as a constructed representation of physical reality, with the goal of the construction being to produce stable representations of object properties. The modern instantiation of this idea is often referred to as the computational approach to understanding vision (Marr, 1982; Landy & Movshon, 1991). Under this view, perception is difficult because multiple scene configurations can lead to the same retinal image. In the case of lightness constancy, the ambiguity arises because illuminant intensity and surface reflectance can trade off to leave the intensity of reflected light unchanged.

Given the ambiguity of the retinal image about the physical configuration of the scene, what we see must depend not only on the image but also on the rules the visual system employs to choose one configuration as that perceived. Although this general notion is well accepted, different authors choose to formulate the nature of these rules in a variety of ways, with some focusing on constraints imposed by known mechanisms (e.g. Stiles, 1967; Cornsweet, 1970) and others on constraints imposed by the statistical structure of the environment (e.g. Gregory, 1968; Marr, 1982; Landy & Movshon, 1991; Wandell, 1995; Geisler & Kersten, 2002; Purves & Lotto, 2003). In our view, the modeling agenda is less to decide amongst these broad conceptions but rather to adopt one of them and use it to provide a quantitative account of a wide range of data. Here we adopt the basic computational approach and use it to develop a model of our data. Our model is essentially identical to that formulated recently by Boyaci et al. (2003) to account for their measurements of the effect of surface slant on perceived lightness.

The starting point is the idea is that the visual system attempts to parse the retinal image and recover the physical characteristics of the scene, so that perceived lightness would be closely correlated with object surface reflectance. In this approach, lightness constancy is found when the visual system is able to do this accurately, while deviations from constancy indicate inaccuracies in the calculation. In previous work, we have elaborated *equivalent illuminant* models of observer performance for tasks where surface mode or surface color was judged (Speigle & Brainard, 1996; Brainard, Brunt, & Speigle, 1997; see also Brainard, Wandell, & Chichilnisky, 1993; Maloney & Yang, 2001; Boyaci et al., 2003). In such models, the observer is assumed to be correctly performing a constancy computation, with the one exception that their estimate of the illuminant deviates from the actual illuminant. Thus in these models, the parameterization of the observer's illuminant estimate determines the range of performance that may be explained, with the detailed calculation then following from an analysis of the physics of image formation.

Equivalent Illuminant Model

Overview

Our model is derived from consideration of an inverse optics calculation that could achieve constancy. The inverse optics calculation begins with parameters that describe the illumination geometry. If these parameters match those of the physical scene, the calculation achieves constancy. Deviations in the model's parameters from those of the scene predict deviations from constancy. In the next sections we describe the physical model of illumination and how this model can be incorporated into an inverse optics calculation to achieve constancy. We then show how the formal development leads to a

parametric model of observer performance.

Physical Model

Consider a Lambertian flat matte standard object¹ that is illuminated by a point² directional light source. The standard object is oriented at a slant θ_N with respect to a reference axis (x-axis in Figure 2). The light source is located at a distance d from the standard surface. The light source azimuth is indicated by θ_D and the light source declination (with respect to the z-axis) by ϕ_D .

The luminance L_{i,θ_N} of the light reflected from the standard surface i depends on its surface reflectance r_i , its slant θ_N , and the intensity of the incident light E :

$$L_{i,\theta_N} = r_i E . \quad (1)$$

When the light arrives only directly from the source, we can write

$$E = E_D \quad (2)$$

where

$$E_D = \frac{I_D \sin \phi_D [\cos(\theta_D - \theta_N)]}{d^2} . \quad (3)$$

Here I_D represents the luminous intensity of the light source.

Equation 3 applies when

$-90^\circ < (\theta_D - \theta_N) < 90^\circ$. For a purely directional source and $(\theta_D - \theta_N)$ outside of this range, $E_D = 0$.

Figure 2 about here

Figure 2. Reference system centered on the standard object. The standard object is oriented so that its surface normal forms an angle θ_N with respect to the x-axis. The light source is located at a distance d from this point, the light source azimuth (with respect to the x-axis) is θ_D , and the light source declination (with respect to the z-axis) is ϕ_D .

In real scenes, light from a source arrives both directly and after reflection off other objects. For this reason, the incident light E can be described more accurately as a compound quantity made of the contribution of directional light E_D and some diffuse light E_A . The term E_A provides an approximate description of the light reflected off other

objects in the scene. We rewrite Equation 2 as

$$E = E_D + E_A \quad (4)$$

and Equation 1 becomes

$$L_{i\theta_N} = r_i \frac{I_D \sin \phi_D [\cos(\theta_D - \theta_N)]}{d^2} + E_A. \quad (5)$$

The luminance of the standard surface $L_{i\theta_N}$ reaches its maximum value when $\theta_D - \theta_N = 0$ and its minimum for -90° ($\theta_D - \theta_N$) 90° . In the latter case only the ambient light E_A illuminates the standard surface.

It is thus useful to simplify Equation 5 by factoring out a multiplicative scale factor α that is independent of θ_N :

$$L_{i\theta_N} = \alpha (\cos(\theta_D - \theta_N) + F_A). \quad (6)$$

In this expression $\alpha = r_i \frac{I_D \sin \phi_D}{d^2}$ and F_A is given by $F_A = \frac{d^2 E_A}{I_D \sin \phi_D}$.

Physical Model Fit

How well does the physical model describe the illumination in our apparatus.? We measured the luminance of our standard objects under all experimental slants, and averaged these over standard object reflectance. Figure 3 (solid circles) shows the resulting luminances from each experiment of the companion paper (Ripamonti et al., 2004) plotted versus the standard object slant. For each experiment, the measurements are normalized to a value of 1 at $\theta_N = 0^\circ$. We denote the normalized luminances by $L_{\theta_N}^{norm}$. The solid curves in Figure 3 denote the best fit of Equation 6 to the measurements, where θ_D , F_A and α were treated as a free parameters and chosen to minimize the mean squared error between model predictions and measured normalized luminances.

Figure 3 about here

Figure 3. The green symbols represent the relative normalised luminance measured for standard objects used in (Ripamonti et al., 2004) and the colored curves illustrate the fit of the model described in the text. Top panel corresponds to the light source set-up used in Experiments 1 and 2, middle panel to Experiment 3 light source on the left, bottom panel for Experiment 3 light source on the right.

The fitting procedure returns two estimated parameters of interest: the azimuth θ_D of the light source and the amount F_A of ambient illumination. (The scalar α simply normalizes the predictions in accordance with the normalization of the measurements.) We can represent these parameters in a polar plot, as shown in Figure 4. The azimuthal position of the plotted points represents θ_D , while the radius v at which the points are plotted is a function of F_A :

$$v = \frac{1}{F_A + 1}. \quad (7)$$

If the light incident on the standard is entirely directional then the radius of the plotted point will be 1. In the case where the light incident is entirely ambient, the radius will be 0.

Figure 4 about here

Figure 4. Light source position estimates of the physical model. Green lines represent the light source azimuth as measured in the apparatus. In Experiments 1, 2 and 3 -light source on the left- the actual azimuth was $\theta_D = -36^\circ$. In Experiment 3 -light source on the right- the actual azimuth was $\theta_D = 23^\circ$. The red symbol represents light source azimuth estimated by the model for Experiments 1 and 2 ($\theta_D = -25^\circ$). For the light source on the left, in Experiment 3, the model estimate is indicated in blue ($\theta_D = -30^\circ$); for the light source on the right in purple ($\theta_D = 25^\circ$). The radius of the plotted points provides information about the relative contributions of directional and ambient illumination to the light incident on the standard object through Equation 7. The radius of the outer circle in the plot is 1. The parameter values obtained for F_A are $F_A = 0.18$ (Experiments 1 and 2), $F_A = 0.43$ (Experiment 3 left), and $F_A = 0.43$ (Experiment 3 right).

The physical model provides a good fit to the dependence of the measured luminances on standard object slant. It should be noted, however, that the recovered azimuth of the directional light source differs from our direct measurement of this azimuth. The most

likely source of this discrepancy is that the ambient light arising from reflections off the chamber walls has some directional dependence. This dependence is absorbed into the model's estimate of θ_D .

Equivalent Illuminant Model

Suppose an observer has full knowledge of the illumination and scene geometry and wishes to estimate the reflectance of the standard surface from its luminance. From Equation 6 we obtain the estimate

$$\tilde{r}_{i,\theta_N} = \frac{L_{i,\theta_N}}{\alpha (\cos(\theta_D - \theta_N) + F_A)}. \quad (8)$$

We use a tilde to denote perceptual analogs of physical quantities.

To the extent that the physical model accurately predicts the luminance of the reflected light, Equation 8 predicts that the observer's estimates of reflectance will be correct and thus Equation 8 predicts lightness constancy. To elaborate Equation 8 into a parametric model that allows failures of constancy, we replace the parameters that describe the illuminant with perceptual estimates of these parameters:

$$\tilde{r}_{i,\theta_N} = \frac{L_{i,\theta_N}}{\alpha (\cos(\tilde{\theta}_D - \theta_N) + \tilde{F}_A)} \quad (9)$$

where $\tilde{\theta}_D$ and \tilde{F}_A are perceptual analogs of θ_D and F_A . Note that the dependence of \tilde{r}_{i,θ_N} on slant in Equation 9 is independent of r_i .

Equation 9 predicts an observer's reflectance estimates as a function of surface slant, given the parameters $\tilde{\theta}_D$ and \tilde{F}_A of the observer's *equivalent illuminant*. These parameters describe the illuminant configuration that the observer uses in his or her inverse optics computation.

Our data analysis procedure aggregates observer matches over standard object reflectance to produce relative normalized matches $\bar{r}_{\theta_N}^{norm}$. The relative normalized matches describe the overall dependence of

observer matches on slant. To link Equation 8 with the data, we assume that the normalized relative matches obtained in our experiment (see Appendix of Ripamonti et al., 2004) are proportional to the computed \tilde{r}_{i,θ_N} , leading to the model prediction

$$\bar{r}_{\theta_N}^{norm} = \beta \frac{L_{\theta_N}^{norm}}{(\cos(\tilde{\theta}_D - \theta_N) + \tilde{F}_A)} \quad (10)$$

where β is a constant of proportionality that is determined as part of the model fitting procedure. In Equation 10 we have substituted $L_{\theta_N}^{norm}$ for L_{i,θ_N} since the contribution of surface reflectance r_i can be absorbed into β .

Equation 10 provides a parametric description of how our measurements of perceived lightness should depend on slant. By fitting the model to the measured data, we can evaluate how well the model is able to describe performance, and whether it can capture the individual differences we observe. In fitting the model, the two parameters of interest are $\tilde{\theta}_D$ and \tilde{F}_A , while the parameter β simply accounts for the normalization of the data.

In generating the model predictions, values for θ_N and $L_{\theta_N}^{norm}$ are taken as veridical physical values. It would be possible to develop a model where these were also treated as perceptual quantities and thus fit to the data. Without constraints on how $\tilde{\theta}_N$ and $\tilde{L}_{\theta_N}^{norm}$ are related to their physical counterparts, however, allowing these as parameters would lead to excessive degrees of freedom in the model. In our slant matching experiment, observer's perception of slant was close to veridical and thus using the physical values of θ_N seems justified. We do not have independent measurements of how the visual system registers luminance.

Model Fit

Fitting the model

For each observer, we used numerical search to fit the model to the data. The search procedure found the equivalent illuminant parameters $\tilde{\theta}_D$ and \tilde{F}_A and the scaling parameter β that provided the best fit to the data. The best fit was determined as follows. For each of the three sessions $k=1,2,3$ we found the normalized relative matches for that session, $\bar{r}_{k,\theta_N}^{norm}$. We then found the parameters that minimized the mean squared error between the model's prediction and these $\bar{r}_{k,\theta_N}^{norm}$. The reason for computing the individual session matches and fitting to these, rather than fitting directly to the aggregate $\bar{r}_{\theta_N}^{norm}$ is that the former procedure allows us to compare the model's fit to that obtained by fitting the session data at each slant to its own mean.

Model Fit

Model fit results are illustrated in the left hand columns of Figures 5 to 10. The dot symbols (green for Neutral instructions and purple for Paint Instructions) are observers' normalized relative matches and the orange curve in each panel shows the best fit of our model. We also show the predictions for luminance and constancy matches as, respectively, a blue or red dashed line. The right hand columns of Figures 5 to 10 show the model's $\tilde{\theta}_D$ and \tilde{F}_A for each observer, using the same polar format introduced in Figure 4.

Figure 5 about here

Figure 5. Model fit to observers' relative normalized matches. In the left column the green dots represent observers' relative normalized matches as a function of slant for Experiment 1. Error bars indicate 90% confidence intervals. The orange curve is the model's best fit for that observer. The blue dashed curve represents predictions for luminance matches and the red dashed line for constancy matches. The right column shows the equivalent illuminant parameters (green symbols) in the same polar format introduced in Figure 4. The polar plot also shows the illuminant parameters obtained by fitting the physical model to the measured luminances (red symbols). The numbers at the top left of each square are the constancy index (CI) for the observer. Observers are listed in increasing order of CI.

Figure 6 about here

Figure 6. Model fit to observers' relative normalized matches for Experiment 2. Same format as Figure 5, except that observer data and parameters are shown in purple.

Figure 7 about here

Figure 7. Model fit to observers' relative normalized matches for Experiment 3 (light on the left, Neutral instructions). Same format as Figure 5.

Figure 8 about here

Figure 8. Model fit to observers' relative normalized matches for Experiment 3 (light on the right, Neutral instructions). Same format as Figure 5. Note that the observer order differs from Figure 7 as the ordering of the observers in terms of CI was not the same for left and right light positions.

Figure 9 about here

Figure 9. Model fit to observers' relative normalized matches for Experiment 3 (light on the left, Paint instructions). Same format as Figure 6.

Figure 10 about here

Figure 10. Model fit to observers' relative normalized matches for Experiment 3 (light on the right, Paint instructions). Same format as Figure 6. Note that the observer order differs from Figure 9 as the ordering of the observers in terms of CI was not the same for left and right light positions.

With only a few exceptions, the equivalent illuminant model captures the wide range of performance exhibited by individual observers in our experiment. To evaluate the quality of the fit, we can compare the mean squared error for the equivalent illuminant model to the variability in the data. To make this comparison, we fit the $\bar{r}_{k,\theta_N}^{norm}$ at each session and

slant by their own means. This provides a lower bound on the squared error that could be obtained by any model. For comparison, we also compute the error obtained by fitting the $\tilde{r}_{k,\theta_N}^{norm}$ with the luminance matching and constancy matching predictions. Figure 11 shows the fit errors (averaged over all experiments and observers) when the data are fit by their own mean, by the equivalent illuminant model, and by the predictions of luminance and constancy matching. We see that the error from the equivalent illuminant model is only slightly greater than the error based on predicting with the mean, and considerably smaller than the error obtained with either luminance matching or constancy predictions.

Figure 11 about here

Figure 11. Mean squared errors obtained when the matching data for each slant, session and observer is fitted by: a) its own mean b) the equivalent illuminant model, c) the luminance match prediction, and d) the lightness constancy prediction.

Using the Model

The equivalent illuminant allows interpretation of the large individual differences observed in our experiments. In the context of the model, these differences are revealed as variation in the equivalent illuminant model parameters $\tilde{\theta}_D$ and \tilde{F}_A , rather than as a qualitative difference in the manner in which observers perform the matching task. In the polar plots we see that for each condition, the equivalent illuminant model parameters are in the vicinity of the corresponding physical model parameters, but with some scatter. Observers whose data resembles luminance matching have parameters that plot close to the origin, while those whose data resemble constancy matching have parameters that plot close to those of the physical model. Given the general difficulty of inverse optics problems, it is perhaps not surprising that individual observers would vary in this regard.

Various patterns in the data shown by many observers, particularly the sharp drop in match for $\theta_N = 60^\circ$ when the light is on the left and the non-monotonic nature of the matches with increasing slant, require no special explanation in the context of the equivalent illuminant model. Both of these patterns are predicted by the model for reasonable values of the parameters. Indeed, striking to us was the richness of the model's predictions for relatively small changes in parameter values.

A question of interest in Experiment 3 was whether observers are sensitive to the actual position of the light source. That they are is indicated by comparing the parameter $\tilde{\theta}_D$ across changes in the physical position of the light source. The average value of $\tilde{\theta}_D$ when the light source was on the left in Experiment 3 was -35° , compared to 16° when it was on the right. The shift in equivalent illuminant azimuth of 51° is comparable to the corresponding shift in the physical model parameter (55°), indicating that observers' demonstrate considerable sensitivity to the light source position.

Constancy Index

In previous studies, it has proved useful to develop a *constancy index* (CI) that provides a rough measure of the degree of constancy indicated by the data (e.g. Brainard & Wandell, 1991; Arend, Reeves, Schirillo, & Goldstein, 1991; Brainard et al., 1997; Brainard, 1998; Kraft & Brainard, 1999; Delahunt & Brainard, 2004). Such an abstraction from the data often enables examination of broad patterns that are difficult to examine in the raw data. We use our model fit to develop such an index for our data. Let the vector

$$\mathbf{v} = \begin{pmatrix} v \sin \theta_D \\ v \cos \theta_D \end{pmatrix} \quad (12)$$

be a function of the physical model's parameters θ_D and F_A , with the scalar v computed from F_A using Equation 7

above. Let the vector $\tilde{\mathbf{v}}$ be the analogous vector computed from the equivalent illuminant model parameters $\tilde{\theta}_D$ and \tilde{F}_A . Then we define our constancy index as

$$CI = 1 - \frac{\|\mathbf{v} - \tilde{\mathbf{v}}\|}{\|\mathbf{v}\|}. \quad (13)$$

Intuitively, this index takes on a value of 1 when the equivalent illuminant model parameters match the physical model parameters and a value near 0 when the equivalent illuminant model parameter \tilde{F}_A is very large. This latter case corresponds to where the model predicts luminance matching.

We have computed this CI for each observer/condition, and the resulting values are indicated on the top left of each polar plot in Figures 5-10. Generally, the indices are in accord with a visual assessment of where the data lie with respect to the luminance matching and constancy predictions, although there are some exceptions. This is not surprising given that a single number cannot describe all aspects of the data.

Figure 12 summarizes the constancy indices across all of our experiments. The mean constancy index was 0.57, with large individual variation. This mean CI quantifies, in a rough manner, the overall degree of lightness constancy shown by our observers. Given the computational difficulty of recovering lighting geometry from images, we regard this degree of constancy as a fairly impressive achievement. Across all experiments, the mean CI for observers given Neutral Instructions (0.56) was essentially identical to that for observers given Paint Instructions (0.58), consistent with our previous conclusion that our instructional manipulation had little effect.

Figure 12 about here.

Figure 12. Constancy indices. Individual observer and mean constancy indices are plotted for each experimental condition. Orange filled circles, individual observer indices. Green filled circles, mean constancy indices.

Figure 13 plots the constancy index for each observer when the light source was on the right against the corresponding index when the light was on the left. These are only moderately correlated ($r = 0.37$). The small positive correlation indicates that whatever factors mediate the degree of constancy shown by individual observers have something in common across the change in light source position, but that this commonality is not dominant.

Figure 13 about here

Figure 13. Comparison of CI for Experiment 3. CI for the illuminant on the right side condition versus illuminant on the left. The correlation between CIs on the left and right was 0.37.

Discussion

Equivalent Illuminant Models

The equivalent illuminant model presented here accounts well for the data reported in the companion paper. The model was derived from an analysis of the computational problem of lightness constancy. The model has two equivalent illuminant parameters, $\tilde{\theta}_D$ and \tilde{F}_A , that describe the lighting geometry. These parameters are not, however, set by measurements of the physical lighting geometry but are fit to each observer's data. Given the equivalent illuminant parameters, the model predicts the lightness matches through an inverse optics calculation.

The equivalent illuminant model developed here is an instance of a general class of models that connect perception to computations that achieve constancy. The general equivalent illuminant model has the following form. First one writes an explicit imaging model that describes how the image data available to the visual system about an object's surface depends on scene parameters that describe physical properties of the illumination and object surface (e.g. Equation 6). The parameters that describe the scene illumination are taken the parameters of the observer model (e.g. $\tilde{\theta}_D$ and \tilde{F}_A).

Given these parameters, the image model is inverted to recover the parameters that describe the object's surface (e.g. Equation 10, which obtains an estimate of object reflectance from the reflected luminance). Perception of the object's surface properties (e.g. lightness, color, glossiness) is then taken to be a function of the estimated object surface properties.

Through different choices of parameterization of illuminant and surface properties, equivalent illuminant models can be developed for a wide range of psychophysical experiments. The most closely related model in the literature is that of Boyaci et al. (Boyaci et al., 2003), which was formulated to account for lightness matching data as a function of standard object slant. This model is substantively identical to ours. Brainard and colleagues have used equivalent illuminant models to account for data on color constancy (Brainard et al., 1997) and the perception of luminosity for chromatic stimuli (Speigle & Brainard, 1996). In these applications, the illuminant parameters describe illuminant spectral power distributions, the surface parameters describe surface spectral reflectance functions, and the imaging model relates reflected spectra to cone photoreceptor isomerization rates.

It is tempting to associate the parameters $\tilde{\theta}_D$ and \tilde{F}_A with the observers perceptual estimate of the illumination geometry. Since our experiments do not explicitly measure this aspect of perception, we have no empirical basis for making the association. In interpreting the parameters as observer estimates of the illuminant, it is important to bear in mind that they are derived from surface lightness matching data and thus, at present, should be treated as illuminant estimates only in the context of our model of surface lightness. It is possible that a future explicit comparison could tighten the link between the derived parameters and conscious perception of

the illuminant. Prior attempts to make such links between implicit and explicit illumination perception, however, have not led to positive results (see Rutherford, 2000).

Comparison with Other Models

A useful approach to modeling context effects is to split the problem into two parts (Krantz, 1968; Brainard & Wandell, 1992). The first part is to determine the parameters that vary across changes in context (or across observers). The second is to determine what features of the context (or observers) set the parameter values and how they do so. The equivalent illuminant model presented here addresses only the first part of the general modeling problem. We have identified two parameters $\tilde{\theta}_D$ and \tilde{F}_A that vary with context (e.g. light source position) and across observer, and connected these parameters to the experimental measurements. Our experiments and analysis are mainly silent about the second part of the general problem: we do not address the question of what about the experimental stimuli causes the parameters to vary. (The one conclusion we draw in this regard is that our instructional manipulation has at most a small effect on the parameters.) In thinking about other approaches to modeling lightness constancy, it is worth keeping the two part distinction in mind.

The retinex theory of Land and McCann (Land & McCann, 1971) addresses both parts of the modeling problem: the modeling computations operate on the directly on the retinal image and produce predictions of lightness at each image location. Thus the retinex model does not explicitly parameterize the effect of context and separate the parameterization from the effect of context on the parameters.

In any case, we do not see how the retinex model can account for our data. The manipulation of object slant produces only small changes in the projected retinal image but large changes in perceived lightness. Since retinex does not account for the

three-dimensional scene geometry, it will have difficulty predicting our effects (see also Gilchrist, 1980; Gilchrist, Delman, & Jacobsen, 1983).

Indeed, any model that accounts for lightness in terms of the contrast between the standard and its surround (either global or local, computed from the two-dimensional image) will have difficulty accounting for our results, since this surround is essentially constant across our slant manipulation. Others have emphasized the difficulties that contrast-based theories of lightness encounter for spatially-rich stimuli (e.g. Gilchrist et al., 1999; Adelson, 1999).

Purves and Lotto (2003) formulate an "empirical" theory of perception and in particular of lightness perception. Like us, they adopt the viewpoint that the visual system attempts to produce percepts that represent object properties. Rather than proceeding via an analytic treatment of the computational problem of constancy, however, these authors suggest that insight about performance is best gained through an empirical database that relates image properties (e.g. luminance) to scene properties (e.g. reflectance). Although we have no objection to this approach in principle, we remain unconvinced that it is feasible to obtain databases sufficiently large to make quantitative predictions for experiments such as ours. Such a database would have to contain a large number of images for surfaces of many reflectances at many poses illuminated by light sources in many positions. The combinatorics do not seem favorable. Interposing explicit imaging models and inverse optics calculations between the characterization of natural scene statistics and predictions of human performance, as we do, mitigates the combinatoric explosion.

In recent work, both Gilchrist (1999) and Adelson (1999) have outlined approaches to lightness perception that emphasize the second part of the general modeling problem: these authors stress the need to understand how the visual system

segments the image into separate "Frameworks" (Gilchrist) or "Atmospheres" (Adelson) and the need to understand what image features in each "Framework"/"Atmosphere" affect how the visual system transforms surface luminance to perceived lightness. In this sense, these theories are broadly complementary to the modeling approach developed here, which focuses on a quantitative characterization of what parameters vary as context is changed.

Extending the Model

The challenge for the equivalent illuminant approach we advocate here is to pursue experiments and theories that predict the equivalent illuminant to the image. An attractive feature of the equivalent illuminant approach is that the model parameters are exactly the quantities estimated by computer vision algorithms designed to achieve constancy. Thus to solve the second part of the modeling problem, any equivalent illuminant model may be coupled with any algorithm that attempts to estimate its parameters from image data. We have begun to explore this in our work on color constancy, where we have coupled an algorithm for spectral illuminant estimation (Brainard & Freeman, 1997) with our equivalent illuminant model for successive color constancy (Brainard, 1998). In initial tests, this general approach has led to successful predictions across a wide range of experimental conditions (Brainard, Kraft, & Longère, 2003). As computational algorithms for estimating illumination geometry become available, these may be used in a similar recipe with the equivalent illuminant model presented here. In the meantime it seems of interest to explore experimentally how the equivalent illuminant parameters vary with more systematic changes in the scene (e.g. light position, relative strength of directional and ambient illumination), as these measurements may provide clues as to what image features affect the parameters.

Acknowledgements

This work was supported NIH EY 10016. We thank B. Backus, H. Boyaci, A. Gilchrist, L. Maloney, J. Nachmias, and S. Sternberg for helpful discussions.

References

- Adelson, E. H. (1999). Lightness perception and lightness illusions. In M. Gazzaniga (Ed.), *The New Cognitive Neurosciences, 2nd ed.* (pp. 339-351). Cambridge, MA: MIT Press.
- Arend, L. E., Reeves, A., Schirillo, J., & Goldstein, R. (1991). Simultaneous color constancy: papers with diverse Munsell values. *Journal of the Optical Society of America A, 8*, 661-672.
- Boyaci, H., Maloney, L. T., & Hersch, S. (2003). The effect of perceived surface orientation on perceived surface albedo in binocularly viewed scenes. *Journal of Vision, 3*, 541-553.
- Brainard, D. H. (1998). Color constancy in the nearly natural image. 2. achromatic loci. *Journal of the Optical Society of America A, 15*, 307-325.
- Brainard, D. H., Brunt, W. A., & 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*(7), 1393-1411.
- Brainard, D. H., Kraft, J. M., & Longère, P. (2003). Color constancy: developing empirical tests of computational models. In R. Mausfeld & D. Heyer (Eds.), *Colour Perception: Mind and the Physical World* (pp. 307-334). Oxford: Oxford University Press.
- 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*. 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*(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.
- Cornsweet, T. N. (1970). *Visual Perception*. New York: Academic Press.
- 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.
- Geisler, W. S., & Kersten, D. (2002). Illusions, perception, and Bayes. *Nature Neuroscience, 5*, 508-510.
- Gilchrist, A., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., Spehar, B., Annan, V., & Economou, E. (1999). An anchoring theory of lightness perception. *Psychological Review, 106*, 795-834.
- Gilchrist, A. L. (1980). When does perceived lightness depend on perceived spatial arrangement? *Perception and Psychophysics, 28*, 527-538.
- Gilchrist, A. L., Delman, S., & Jacobsen, A. (1983). The classification and integration of edges as critical to the perception of reflectance and illumination *Perception and Psychophysics 33*, 425-436.
- Gregory, R. L. (1968). *Visual Illusions*. *Scientific American*.
- Helmholtz, H. (1896). *Physiological Optics*. New York: Dover Publications, Inc.

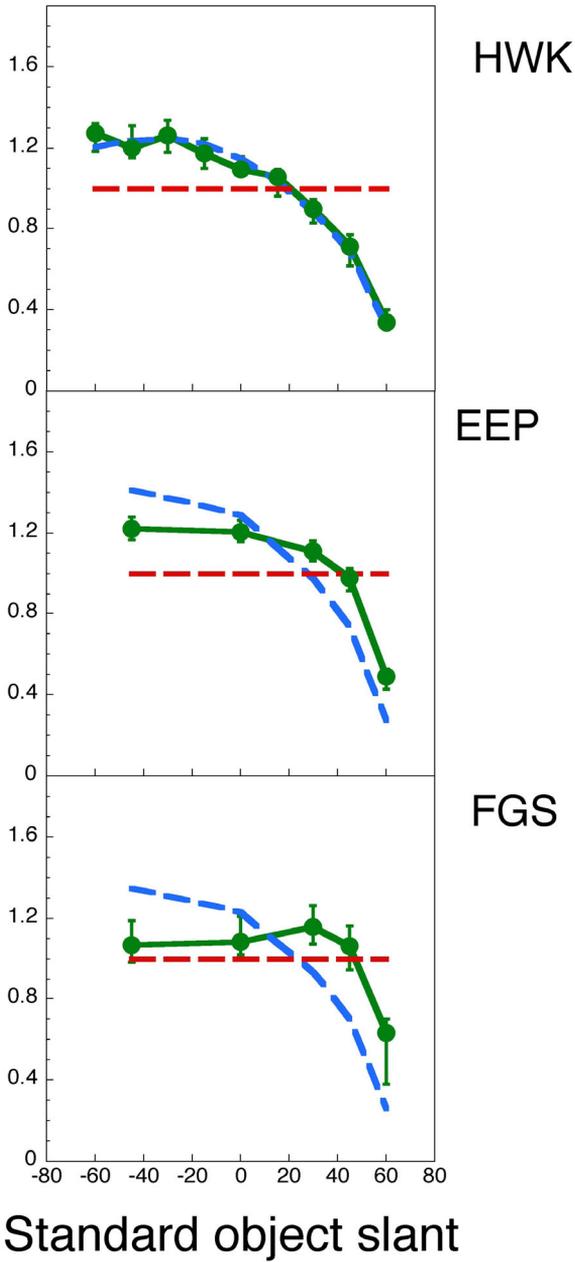
- Kaufman, J. E., & Christensen, J. F. (Eds.). (1972). *IES Lighting Handbook; The Standard Lighting Guide* (5 ed.). New York: Illuminating Engineering Society.
- Kraft, J. M., & Brainard, D. H. (1999). Mechanisms of color constancy under nearly natural viewing. *Proc. Nat. Acad. Sci. USA*, 96(1), 307-312.
- Krantz, D. (1968). A theory of context effects based on cross-context matching. *Journal of Mathematical Psychology*, 5, 1-48.
- Land, E. H., & McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America*, 61, 1-11.
- Landy, M. S., & Movshon, J. A. (Eds.). (1991). *Computational Models of Visual Processing*. Cambridge, MA: MIT Press.
- Maloney, L. T., & Yang, J. N. (2001). The illuminant estimation hypothesis and surface color perception. In R. Mausfeld & D. Heyer (Eds.), *Colour Perception: From Light to Object*. Oxford: Oxford University Press.
- Marr, D. (1982). *Vision*. San Francisco: W. H. Freeman.
- Purves, D., & Lotto, R. B. (2003). *Why We See What We Do: An Empirical Theory of Vision*. Sunderland, MA: Sinauer.
- Ripamonti, C., Bloj, M., Mitha, K., Greenwald, S., Hauck, R., Maloney, S. I., & Brainard, D. H. (2004). Measurements of the effect of surface slant on perceived lightness. *Journal of Vision*, submitted.
- Rutherford, M. D. (2000). *The role of illumination perception in color constancy*. Unpublished Ph.D., UC Santa Barbara, Santa Barbara.
- 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.
- Stiles, W. S. (1967). Mechanism concepts in colour theory. *Journal of the Colour Group*, 11, 106-123.
- Wandell, B. A. (1995). *Foundations of Vision*. Sunderland, MA: Sinauer.

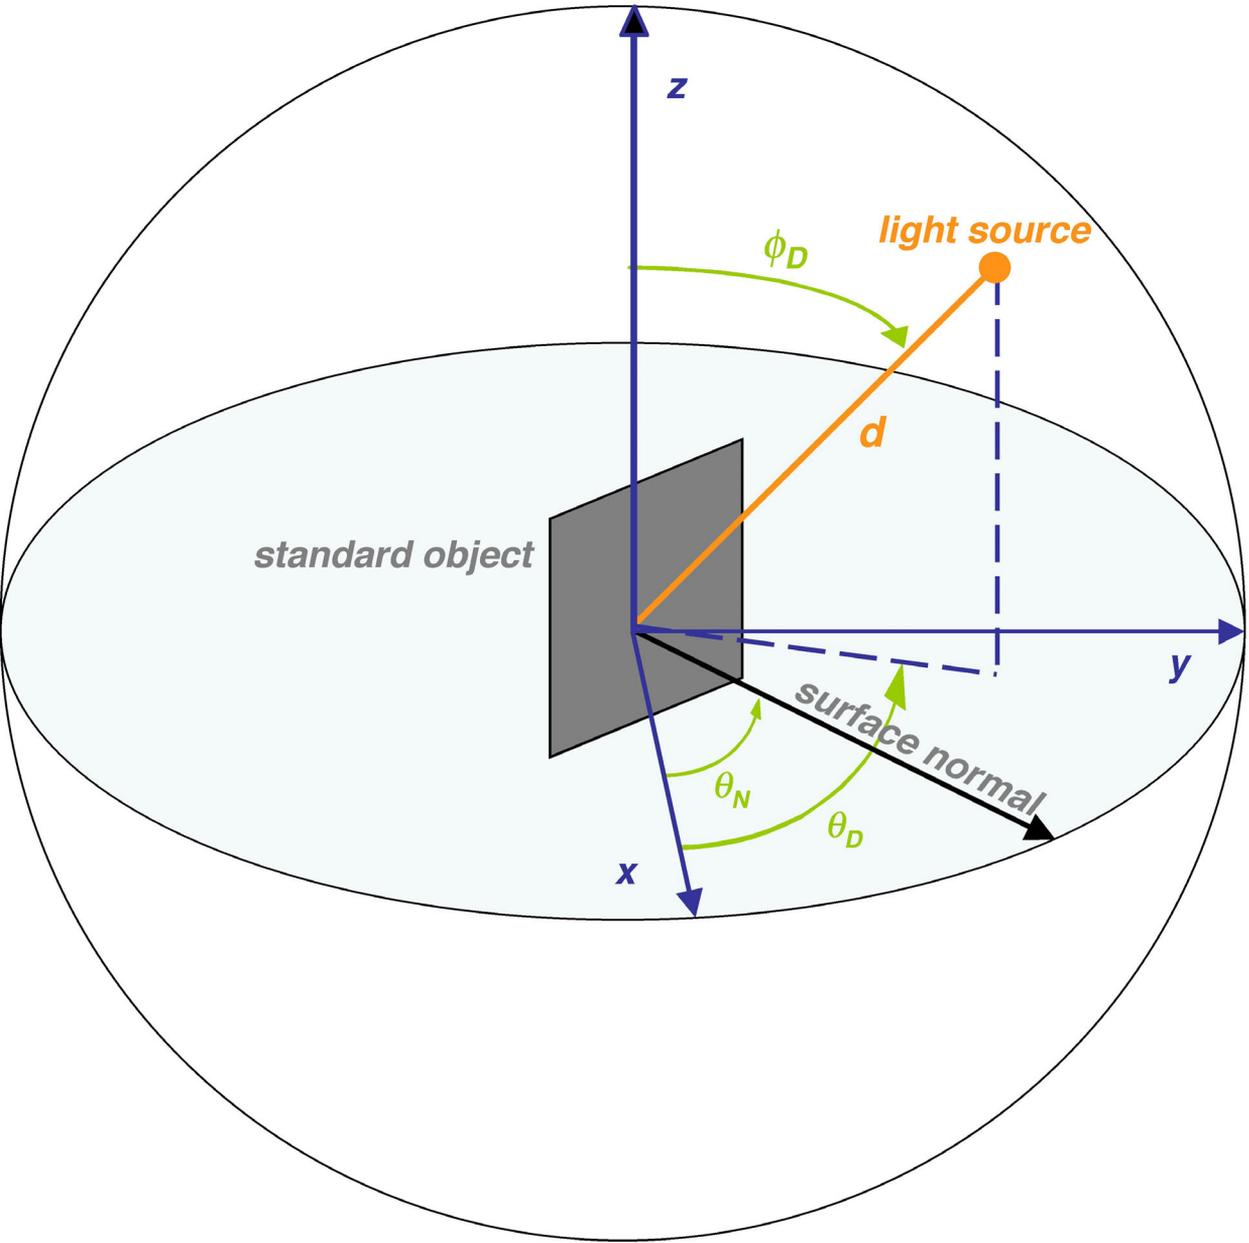
Notes

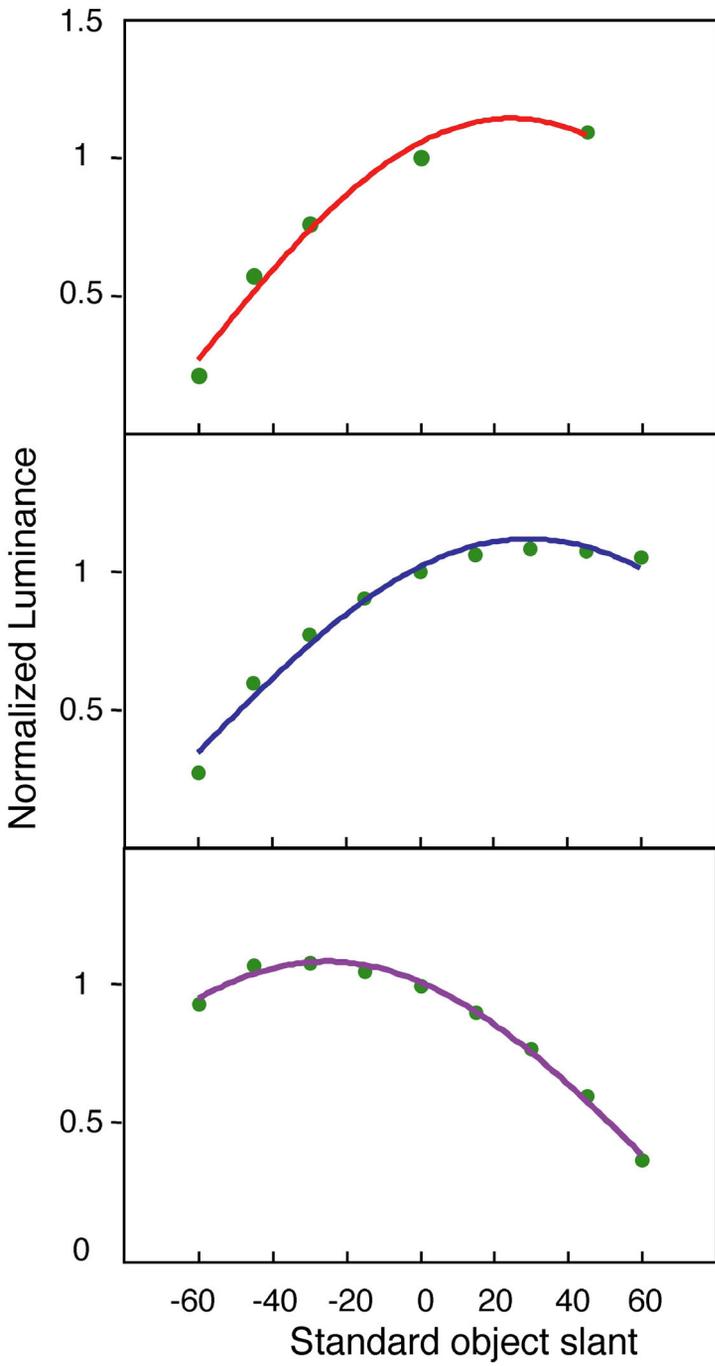
¹A Lambertian surface is a uniformly diffusing surface with constant luminance regardless of the direction from which is viewed (Wyszecki and Stiles, 1982).

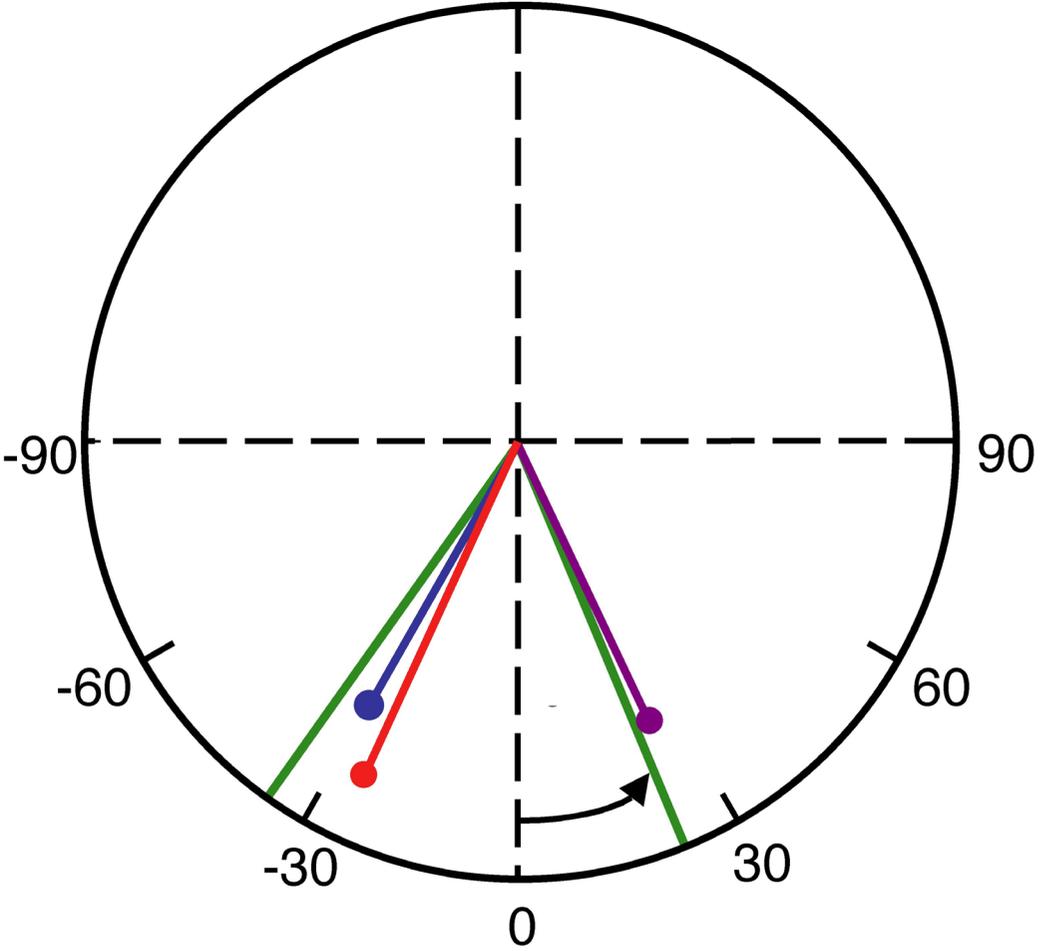
²A light source whose distance from the illuminated object is at least 5 times its main dimension is considered to be a good approximation of a point light source (Kaufman & Christensen, 1972).

Relative match reflectance

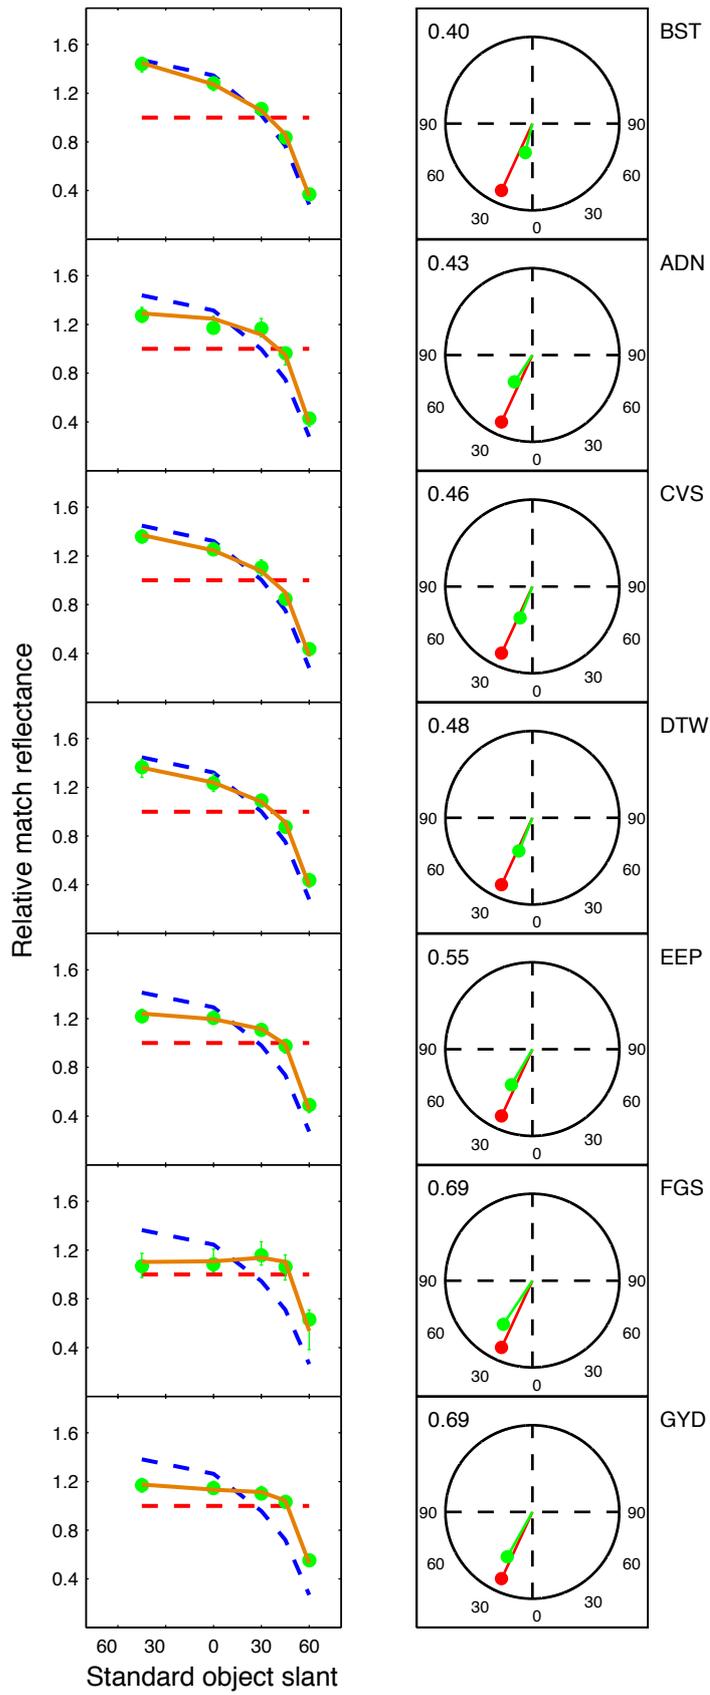






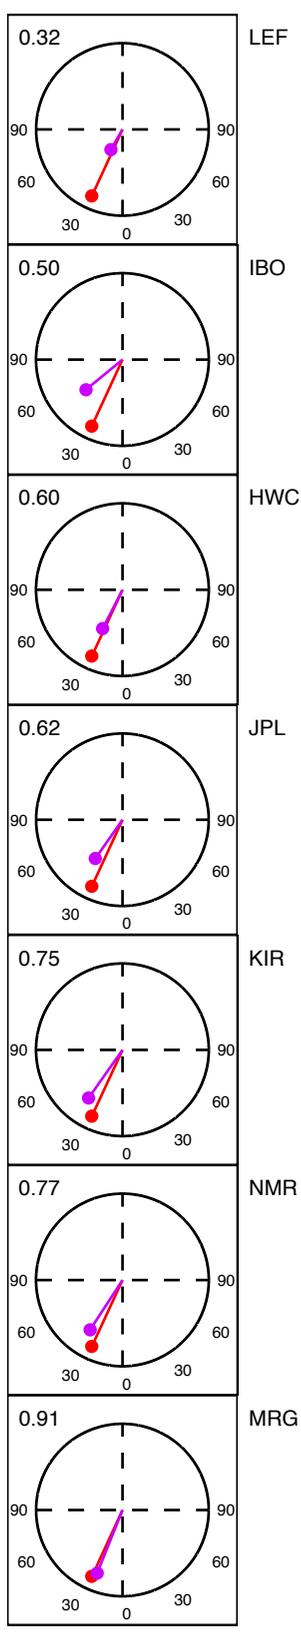
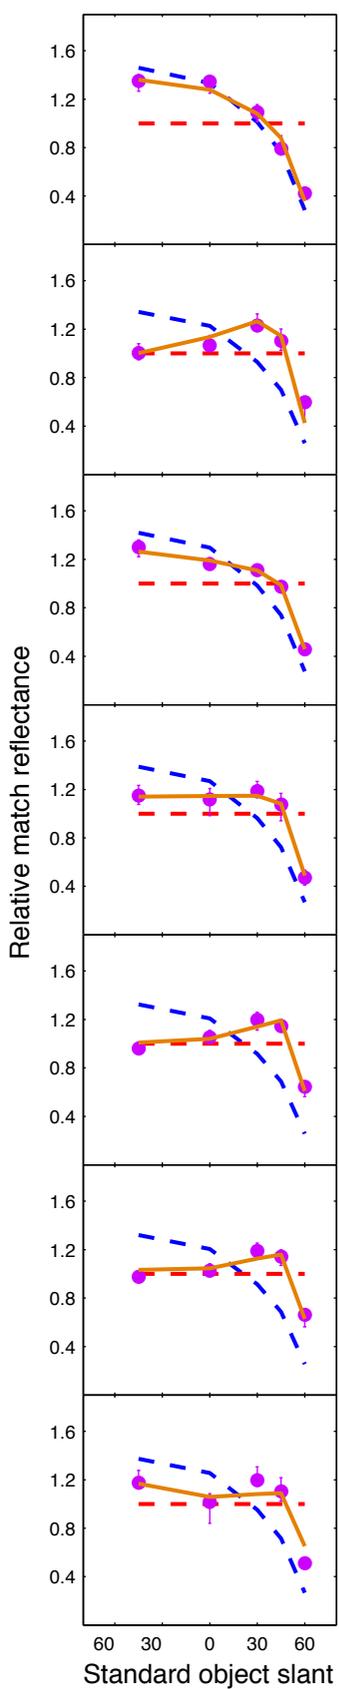


● Observer match
— Model Fit
- - - Luminance match
- - - Constancy
— Model estimation
— Physical estimation



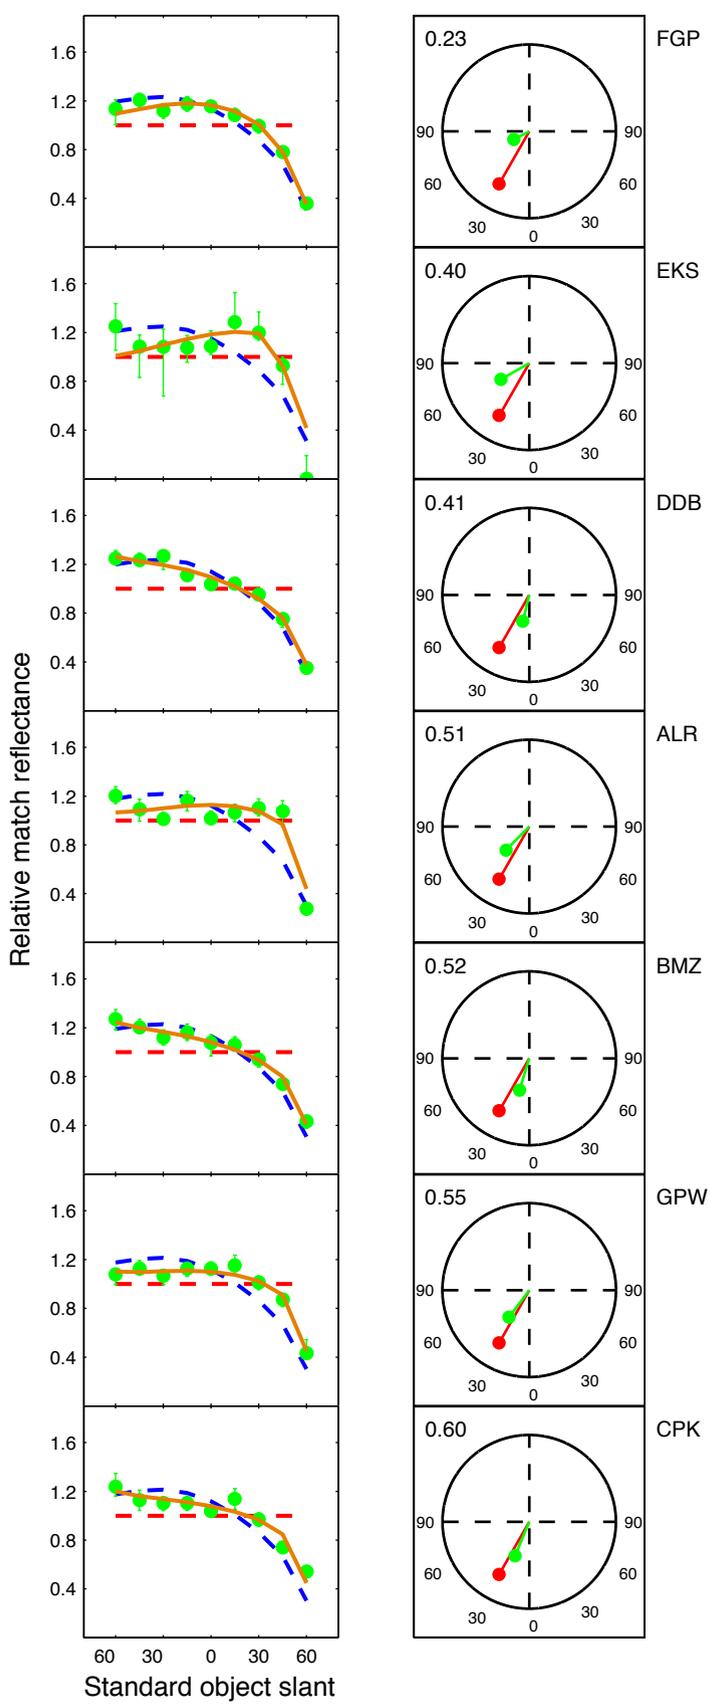
- Observer match
- ModelFit
- Luminance match
- Constancy

- Model estimation
- Physical estimation



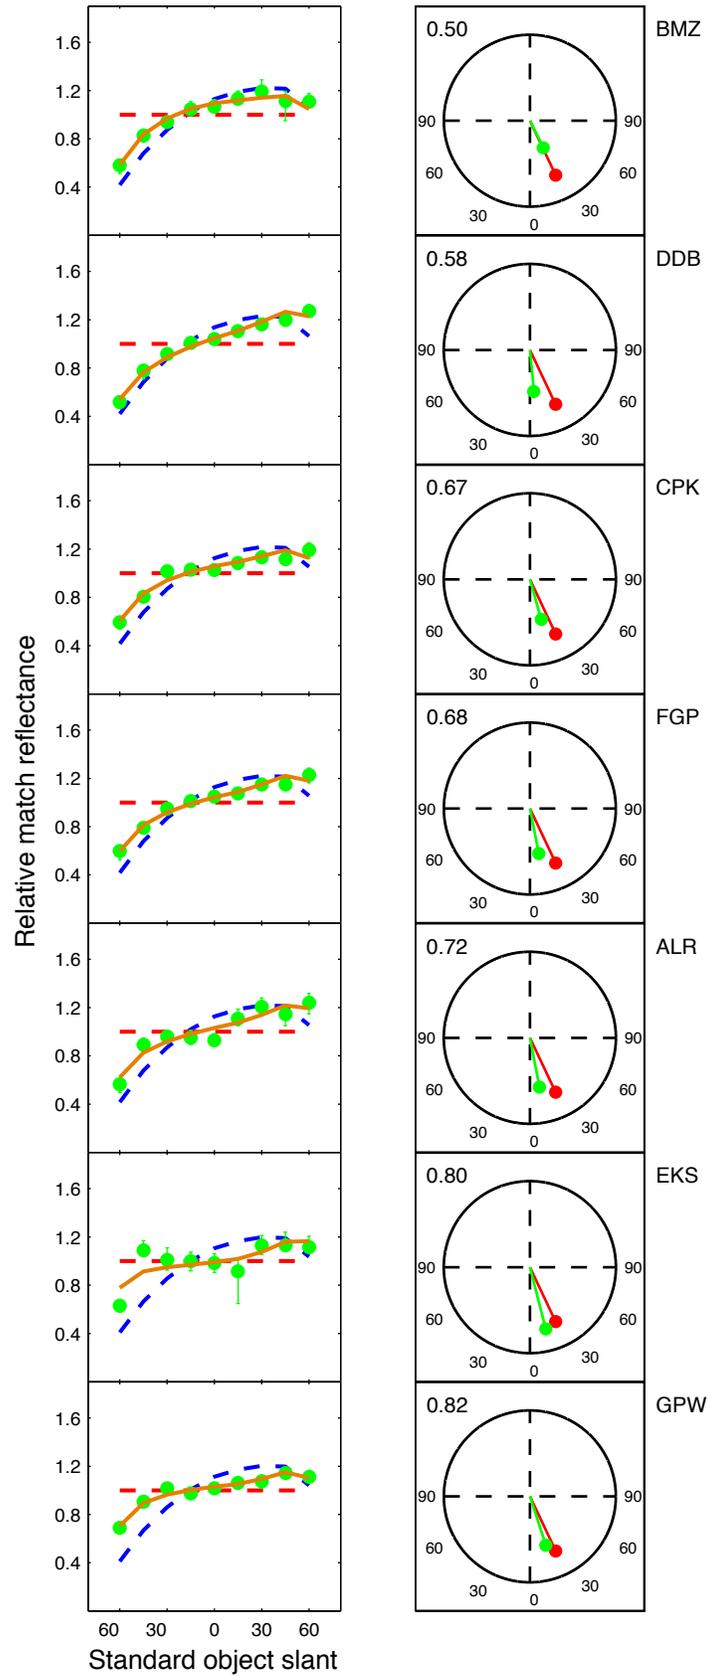
● Observer match
— ModelFit
— Luminance match
— Constancy

— Model estimation
— Physical estimation



- Observer match
- ModelFit
- Luminance match
- Constasy

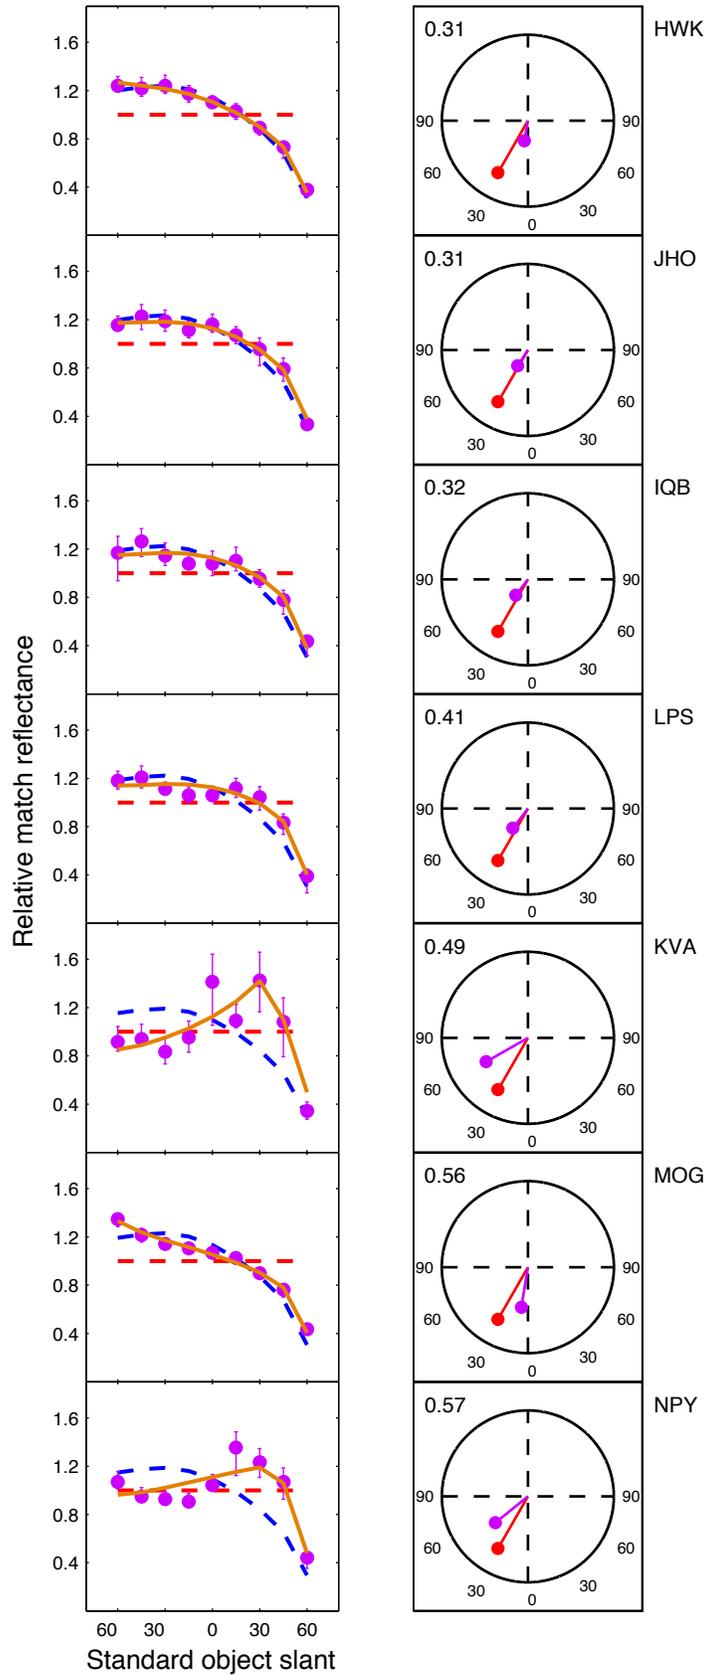
- Model estimation
- Physical estimation



Bloj et al., Figure 8

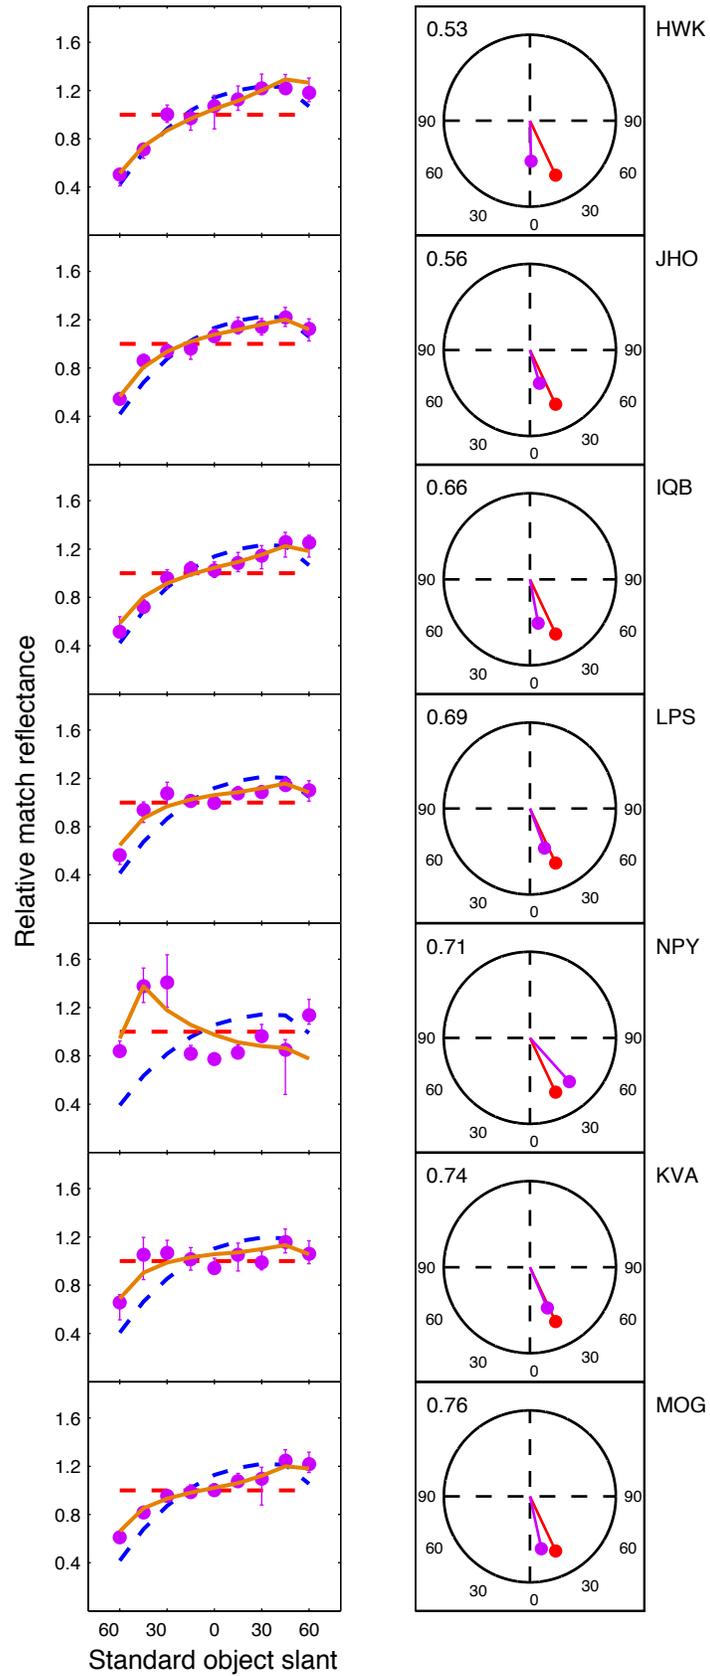
- Observer match
- ModelFit
- Luminance match
- Constancy

- Model estimation
- Physical estimation

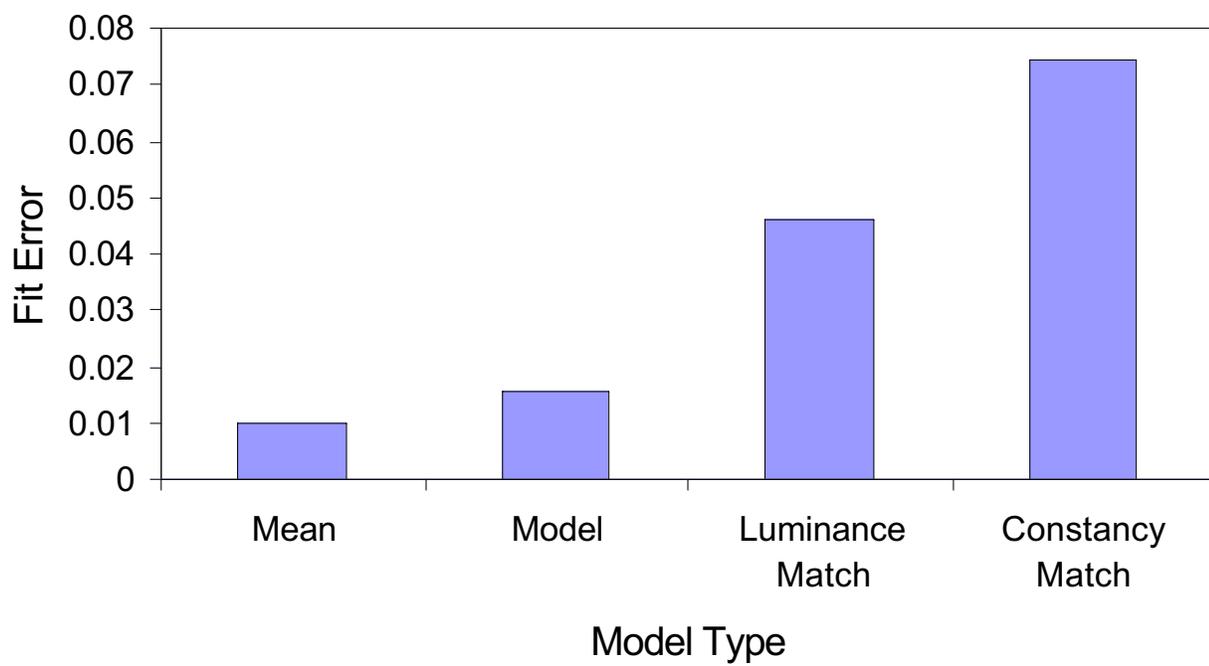


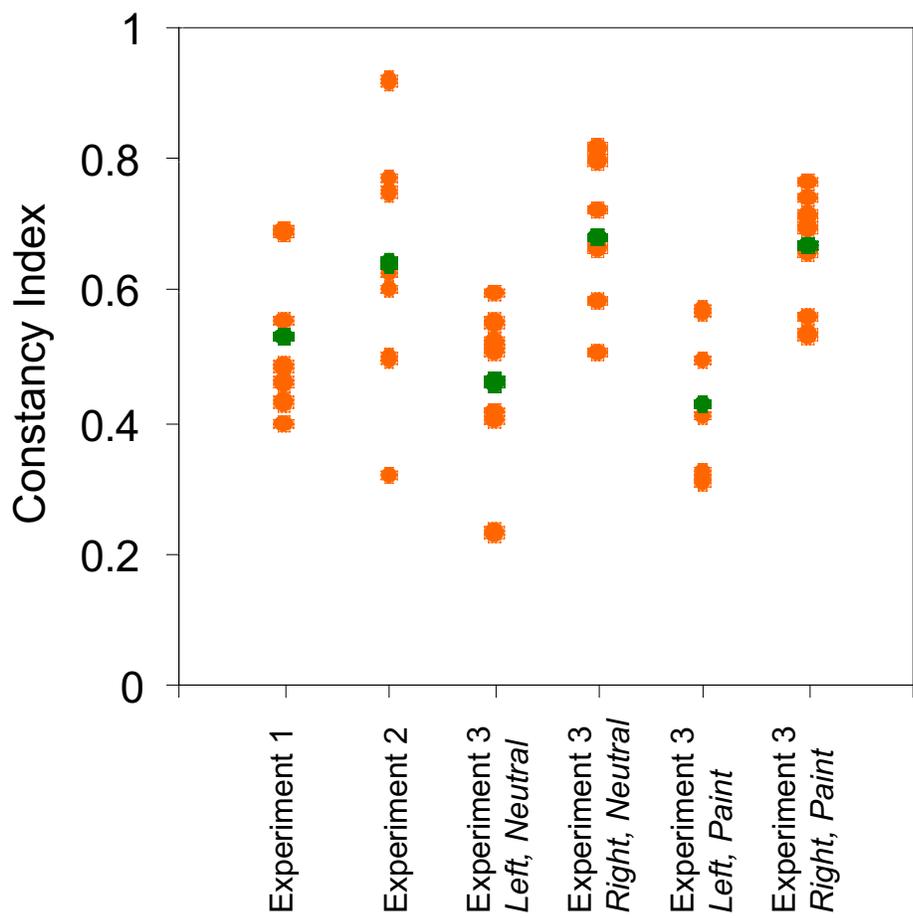
- Observer match
- ModelFit
- Luminance match
- Constancy

- Model estimation
- Physical estimation



Bloj et al., Figure 10





Bloj et al. Figure 12

