The Relation Between Color Discrimination and Color Constancy: When Is Optimal Adaptation Task Dependent?

Alicia Björnsdotter Abrams, James M. Hillis & David H. Brainard*

University of Pennsylvania, Department of Psychology

3401 Walnut St. 302C, Philadelphia, PA, 19104

*To whom correspondence should be addressed

Abstract

Color vision supports two distinct visual functions - discrimination and constancy. Discrimination requires that the visual response to distinct objects with a scene be different. Constancy requires that the visual response to any object be the same across scenes. Across changes in scene illumination, adaptation can improve discrimination by optimizing the use of the available response range. Similarly, adaptation can improve constancy by stabilizing the visual response to any fixed object across changes in illumination. Can common mechanisms of adaptation achieve these two goals simultaneously? This paper develops a theoretical framework for answering this question and presents several example calculations. For changes in illuminant spectral power distribution typical of variation in daylight and natural surface spectra, the answer is yes. For changes in the statistical ensemble of surfaces in scenes, the answer is no.

Submitted to Neural Computation, August 2005.
Comments welcome. Please do not quote.

INTRODUCTION

Color vision supports two distinct visual functions - discrimination and constancy (e.g. Jacobs, 1981; Mollon, 1982). Color discrimination, the ability to determine that two spectra differ, is useful for segmenting an image into regions corresponding to distinct objects. Effective discrimination requires that the visual response to distinct objects within a scene be different. Across changes in scene, adaptation can improve discrimination by optimizing the use of the available response range for objects in the scene (e.g. Walraven *et al.*, 1990).

Color constancy is the ability to identify objects on the basis of their color appearance (Brainard, 2004). Because the light reflected from an object to the eye depends both on the object's surface reflectance and on the illumination, constancy requires that some process stabilize the visual representation of surfaces across changes in illumination. Early visual adaptation can mediate constancy if it compensates for the physical changes in reflected light caused by illumination changes (e.g. Wandell, 1995).

Although there are large theoretical and empirical literatures concerned both with how adaptation affects color appearance and constancy on the one hand (e.g. Wyszecki, 1986; Zaidi, 1999; Foster, 2003; Shevell, 2003; Brainard, 2004) or discrimination on the other (e.g. Wyszecki & Stiles, 1982; Walraven *et al.*, 1990; Hood & Finkelstein, 1986; Lennie & D'Zmura, 1988; Kaiser & Boynton, 1996; Eskew *et al.*, 1999), it is rare that the two functions are considered simultaneously . Still, it is clear that they are intimately linked since they rely on the same initial representation of spectral information. In addition, constancy is only useful if color vision also supports some amount of discrimination performance; in the absence of any requirement for discrimination, constancy can be achieved trivially by a visual system that assigns the same single color to every object in every scene.[1] The recognition that constancy (or its close cousin appearance) and discrimination are profitably considered jointly has been exploited in a few recent papers (Robilotto & Zaidi, 2004; Hillis & Brainard, in press).

Here we ask whether applying the same adaptive transformations to visual responses can simultaneously optimize performance for both constancy and discrimination. If the visual system adapts to each of two environments so as to produce optimal color discrimination within each, what degree of constancy is achieved? How does this compare with what is possible if adaptation is instead tailored to optimize constancy, and what cost would such an alternative adaptation strategy impose on discrimination performance?

To address these questions we adopt the basic theoretical framework introduced by Grzywacz and colleagues (Grzywacz & Balboa, 2002; Grzywacz & Juan, 2003; also von der Twer & MacLeod, 2001; Foster *et al.*, 2004) by analyzing task performance using explicit models of the visual environment and early visual processing. Parameters in the model visual system specify the system's *state of adaptation*, and we study how these parameters should be set to maximize performance, where the expectation is taken across scenes drawn from a statistical model of the visual environment (see Grzywacz & Balboa, 2002; Grzywacz & Juan, 2003). Within this framework, we investigate the tradeoffs between optimizing performance for discrimination and for constancy. We begin with consideration of a simple one-dimensional example that illustrates the basic

[1] This is sometimes referred to as the Ford algorithm, after a quip attributed to Henry Ford: "People can have the Model T in any color - so long as it's black." (http://en.wikiquote.org/wiki/Henry_Ford).

ideas (Univariate Example), and then generalize to color (Contrast Adaptation, Chromatic Adaptation.) The work presented here complements our recent experimental efforts directed towards understanding the degree to which measured adaptation of the visual pathways mediates judgments of both color discrimination and color appearance (Hillis & Brainard, in press).

Univariate Example

We begin with the specification of a model visual system, a visual environment, and performance measures. The basic structure of our problem is well-illustrated for the case of lightness/brightness constancy and discrimination, and we begin with a treatment of this case.

Visual Environment

The model visual environment consists of achromatic matte surfaces lit by a diffuse illuminant. Each surface j is characterized by its reflectance r_j, which specifies the fraction of incident illumination that is reflected. Each illuminant is specified by its intensity e_i. The intensity of light $c_{i,j}$ reflected from surface j under illuminant i is thus given by

$$c_{i,j} = e_i r_j . \qquad (1)$$

At any given moment, we assume that the illuminant e_i is known and that the particular surfaces in the scene have been drawn from an ensemble of surfaces. The ensemble statistics characterize the regularities of the visual environment. In particular, we suppose that

$$r_j \sim \bar{N}(\mu_r, \sigma_r^2) \qquad (2)$$

where ~ indicates "distributed as" and $\bar{N}(\mu_r, \sigma_r^2)$ represents a truncated Normal distribution with mean parameter μ_r and variance parameter σ_r^2. The overbar in the notation indicates the truncation, which means that probability of obtaining a reflectance in the range $0 \leq r_j \leq 1$ is proportional to the standard Normal density function while the probability of obtaining a reflectance outside of this range is zero. The truncated distribution is normalized so that the total probability across all possible values of r_j is unity.

We are interested in a) how well a simple model visual system can discriminate and identify randomly chosen surfaces viewed within a single scene and b) how well the same visual system can discriminate and identify randomly chosen surfaces viewed across different scenes where the illumination, surface ensemble, and/or state of adaptation have changed.

MODEL VISUAL SYSTEM

The model visual system has a single class of photoreceptor. At each location, the information transmitted by this photoreceptor is limited in two ways. First, the receptor has a limited response range. We capture this by supposing that the deterministic component of the response to surface j under illuminant i is given by

$$u_{i,j} = \frac{\left(gc_{i,j}\right)^n}{\left(gc_{i,j}\right)^n + 1} \tag{3}$$

where $u_{i,j}$ represents the visual response, $c_{i,j}$ represents the intensity of incident light obtained through Eq. (1), g is a gain parameter, and n is a steepness parameter that controls the slope of the visual response function. For this model visual system, the *adaptation parameters* g and n characterize the system's state of adaptation. For each scene, the visual system may set g and n to optimize its performance, but we require that the same parameters be applied for the incident intensities reflected from every surface in the scene.[2]

The second limit on the transmitted information is that the responses are noisy. We can capture this by supposing that the deterministic component of the visual responses are perturbed by zero mean additive noise that is Normally distributed with variance σ_n^2.

DISCRIMINATION TASK AND PERFORMANCE MEASURE

To characterize discrimination performance, we need to specify a discrimination task. We consider a same-different task. On each trial, the observer either sees two views of the same surface (*same trials*) or one view each of different surfaces (*different trials*), all viewed under the same light e_i. The observer's task is to respond "same" on the same trials and "different" on the different trials. On same trials, a single surface is drawn at random from the surface ensemble and viewed twice. On different trials, two surfaces are drawn independently from the surface ensemble. Independently drawn noise is added to the response for each view of each surface. This task is referred to in the signal detection literature as a roving same-different design (Macmillan & Creelman, 2005).

---

[2] Treatments of adaptation are clarified by distinguishing two separate issues (Stiles, 1967; Krantz, 1968; Brainard & Wandell, 1992; Brainard, 2004). First, what are the adaptation parameters? Second, what aspects of the visual input determine the state of adaptation? Here our interest lies in understanding performance tradeoffs across various choices of adaptation parameters. We use parameter search to optimize adaptation parameters in a manner that provides access to information not available to actual visual systems. Previous authors have investigated how the visual input may be used to set adaptive parameters to optimize discrimination (Grzywacz & Juan, 2003) or achieve constancy (see e.g. Brainard & Freeman, 1997; Hurlbert, 1998; Maloney, 1999).

The observer's performance is characterized by a hit rate (fraction of "same" responses on same trials) and a false alarm rate (fraction of "same" responses on different trials).

It is well-known that the hit and false alarm rates obtained by an observer in a same-different task depend both on the quality of the information supplied by the visual responses (i.e. signal-to-noise ratio) and on how the observer chooses to trade off hits and false alarms (Green & Swets, 1966; Macmillan & Creelman, 2005). By tolerating more false alarms an observer can increase his or her hit rate. Indeed, by varying the response criterion used in the hit/false-alarm tradeoff, an observer can obtain performance denoted by a locus of points in what is referred to as an ROC diagram (see Figure 1 and its caption). A standard criterion-free measure of the quality of information available in the visual responses is A', the area under the ROC curve (Green & Swets, 1966; Macmillan & Creelman, 2005). In this paper, we use A' as our measure of performance, both for discrimination (as is standard) and for constancy (see below).

Figure 1. ROC diagram. The ROC (receiver operating characteristic) diagram plots hit rate versus the false rate. An observer can maximize hit rate by responding "same" on evary trial. This will lead to a high false alarm rate and performance will plot at (1,1) in the diagram. An observer can minimize false alarms by responding "different" on every trial can and achieve performance at (0,0). Varying criteria between these two extremes produces a trade-off between hits and false alarms. The exact locus traced out by this tradeoff depends on the information used at the decision stage. Better information leads to performance curves that tend more towards the upper left of the plot (solid curve indicates better information than the dashed curve.) The area under the ROC curve, referred to as A', is a task-specific measure of information that does not depend on criterion. The hatched area is A' for the dashed ROC curve. The ROC curves shown were computed for two surfaces with reflectances $r_1 = 0.15$ and $r_2 = 0.29$ presented in a roving same-different design. The illuminant had intensity $e = 100$ and the deterministic component of the visual responses was computed from Equation (3) with $g = 0.02$ and $n = 2$. The solid line corresponds to $\sigma_n = 0.05$ and $A' = 0.85$ while the dashed line corresponds to $\sigma_n = 0.065$ and $A' = 0.76$. Hit and false alarm rates were computed using the decision rule described in the next section.

Effect of Adaptation on Discrimination

To understand the effect of adaptation, we ask how the average A' depends on the adaptation parameters, given the surface ensemble, illuminant, and noise. For a roving design, a near-optimal strategy is to compute the difference between the visual responses for two surfaces and compare this difference to a criterion C (Macmillan & Creelman, 2005). Let $u_{i,j}$ be the visual response to one surface under the given illuminant e_i, $u_{i,k}$ to the other surface. The observer responds "same" if the squared response difference $\Delta u_{i,jk}^2 = \left\| u_{i,j} - u_{i,k} \right\|^2$ is less than C and "different" if $\Delta u_{i,jk}^2 \geq C$.

For any pair of surfaces r_j and r_k, we can compute the values of the deterministic component of the corresponding visual responses ($u_{i,j}$ and $u_{i,k}$), once we know the illuminant e_i and the adaptation parameters g and n. Because of noise, the observed response difference $\Delta u_{i,jk}^2$ varies from trial-to-trial. If the variance of the noise is σ_n^2, the distribution of the quantity $\left(\Delta u_{i,jk}'\right)^2 = \left(\Delta u_{i,jk} / \sqrt{2}\sigma_n\right)^2$ is non-central chi-squared with 1 degree of freedom and non-centrality parameter $\left(\Delta u_{i,jk} / \sqrt{2}\sigma_n\right)^2$.[3] Because scaling the visual response for both same and different trials by a common factor $1 / \sqrt{2}\sigma_n$ does not affect the information contained in these responses, a decision rule based on comparing $\left(\Delta u_{i,jk}'\right)^2$ criterion $C' = C / 2\sigma^2$ leads to the same performance as one that compares $\Delta u_{i,jk}^2$ to C. Thus the known non-central chi-square distributions on same and different trials may be used, along with standard signal detection methods, to compute hit and false

[3] On same trials the difference $\left(\Delta u_{i,jk}'\right)^2$ is 0 and the distribution reduces to ordinary chi-squared.

alarm rates for a set of criteria. The resultant ROC curve may then be numerically integrated to find the value of $A'_{i,jk}$.

To evaluate overall discrimination performance, we compute $A'_{i,jk}$ for many pairs of surfaces drawn according to the surface ensemble. We sort the resulting $A'_{i,jk}$ and average the lower 50th percentile of values in the ordered list to obtain an aggregate measure of discrimination performance, $\overline{A}'_{50}$.[4]

Figure 2 illustrates how the gain parameter g affects discrimination performance for a single illuminant and surface ensemble, when the steepness parameter is held fixed at $n = 2$. The top panel shows a plot of $\overline{A}'_{50}$ for two noise levels. There is an optimal choice of gain for each noise level. The solid line plotted in the middle right panel shows the response function obtained for the optimal choice of gain when $\sigma_n = 0.05$. The histogram below the x-axis shows the distribution of reflected light intensities, while that to the left of the y-axis shows the distribution of visual responses. The histogram of responses appears more uniform than the histogram of light intensities. This general effect is expected from standard results in information theory, where maximizing the information transmitted by a channel occurs when the distribution of channel responses is uniform (Cover & Thomas, 1991). The response histogram is not perfectly uniform because varying the gain alone cannot produce this result, and because our performance

---

[4] The decision to average only the lower 50th percentile of the $A'_{i,jk}$ is somewhat arbitrary. Doing so increases the sensitivity of the aggregate measure to changes in the adaptation parameters, as the larger of the $A'_{i,jk}$ tend to cluster near their ceiling value of 1.0. We have not systematically explored the effect of varying the range of $A'_{i,jk}$ averaged.

measure is $\overline{A}'_{50}$ rather than bits transmitted. Figure 5 below shows response histograms when both gain and steepness parameters are allowed to vary.

Figure 2. Effect of gain on discrimination performance. The top panel plots the discrimination measure $\overline{A}'_{50}$ as a function of the gain parameter, for two noise levels (solid line, $\sigma_n = 0.05$; dashed line $\sigma_n = 0.065$). In the calculations, we set $n = 2$, $\mu_r = 0.5$, $\sigma_r = 0.3$, and $e = 100$. The bottom panel shows the response function for the optimal choice of gain when the noise is $\sigma_n = 0.05$ ($g = 0.0205$, $\overline{A}'_{50} = 0.73$). Below the response function is a histogram of the light intensities c_{ij} reaching the eye, while to the left is a histogram of the resultant visual responses. Calculations were performed for 500 draws from the surface ensemble, and $A'_{i,jk}$ was evaluated for all possible 124750 surface pairs formed from these draws. Choices of gain less than or greater than the optimum shift would shift the response function right or left. For these non-optimal choices, visual responses would tend to cluster near the floor or ceiling of the response range, resulting in poorer discrimination performance.

EFFECT OF ILLUMINANT CHANGE ON DISCRIMINATION

We can also investigate the effect of illumination changes on performance, and how adaptation can compensate for such changes. First consider the case where the adaptation parameters are held fixed and the illuminant is varied. The solid curve shown in Figure 3 plots $\overline{A}'_{50}$ as a function of the illuminant intensity. Not surprisingly, performance falls off with the change of illuminant: Increasing the illuminant intensity pushes the intensity of the reflected light towards the saturating region of the visual response function and compresses the response range used. The effect of increasing the illuminant intensity is multiplicative, so this effect can clearly be compensated for by decreasing the gain (which also acts multiplicatively) so as to keep the distribution of responses constant. Perfect compensation is possible in this example because of the match between the physical effect of an illuminant change (multiplication of all reflected light intensities by the same factor) and the effect of a gain change (also multiplication of the same intensities by a common factor). In general such perfect compensation is not possible.

CONSTANCY

Suppose that instead of discriminating between surfaces seen under a common illuminant, we ask the question of constancy: how well can the visual responses be used to judge whether two surfaces are the same, when one surface is viewed under a *reference illuminant* and the other is viewed under a *test illuminant*? Imagine that on same trials of an experiment, the observer sees a surface, drawn at random from the surface ensemble, under the reference illuminant and the same surface under the test illuminant. On different trials, two surfaces are drawn from the surface ensemble and one is seen under the reference illuminant and the other under the test illuminant. As in the one illuminant discrimination task, the observer must respond "same" or "different". We assume that the observer continues to employ the same basic distance decision rule introduced previously, with the decision variable evaluated across the change of illuminant: $\vec{\Delta} u^2_{jk} = \left\| u_{ref,j} - u_{test,k} \right\|^2$. On same trials, the expression is evaluated for a single surface across the change ($r_k = r_j$), and on different trials the expression is evaluated for two draws from the surface ensemble. In the notation, the arrow indicates the change of illuminant.[5] This decision rule models the case where the observer has no explicit knowledge of the illuminant or state of adaptation.[6]

The quantity A' remains an appropriate measure of performance for the illuminant change case, as it continues to characterize how well hits and false alarms trade off as a

[5] More generally, we will use the notation $\vec{\Delta} u^2_{jk}$ to denote the squared difference in response for two surfaces across a change in illuminant or in adaptation parameters.
[6] This choice may be contrasted with work where the measure of performance is bits of information transmitted (e.g. Foster *et al.*, 2004). Measurements of information transmitted are silent about what subsequent processing is required to extract the information. Here we are explicitly interested in the performance supported directly by the visual response representation.

function of a decision criterion. As previously, we can compute an aggregate measure of A' across many pairs of surfaces drawn from the ensemble, with the only change being that the distribution of the response difference on same trials is no longer guaranteed to have a mean of zero (because of the illuminant change). We will denote the expected performance value here as $\vec{A}'_{50}$, where the arrow indicates that A' has been computed across the illuminant change.[7] The dashed curve in Figure 3 shows how $\vec{A}'_{50}$ falls off with a change in illuminant, when there is no compensatory change in gain. The reference illuminant had intensity 100, and the test illuminant intensity is plotted along the x-axis.[8]

As with the effect of the illuminant change on within-illuminant discrimination performance, the deleterious effect of the illuminant change on constancy may be eliminated by an appropriate gain change. This is because changing the gain with the illumination can restore the responses under the test light back to their values under the reference light.

Figure 3. Effect of illuminant change on discrimination and constancy performance. The plot shows how discrimination performance ($\bar{A}'_{50}$, solid line) and constancy performance ($\vec{A}'_{50}$, dashed line) decrease when the illumination is changed and the state of adaptation is held constant. All calculations performed with adaptation parameters held fixed ($g = 0.0205$, $n = 2$) and for $\sigma_n = 0.05$. These values maximize $\bar{A}'_{50}$ for an illuminant intensity of 100. The surface distribution had parameters $\mu_r = 0.5$ and $\sigma_r = 0.3$. For the computation of $\vec{A}'_{50}$, a reference illuminant intensity of 100 was used.

[7] More generally we can evaluate $\vec{A}'_{50}$ across changes in illuminant, changes in adaptation parameters, or both.

[8] Those familiar with standard signal detection analyses may initially find it odd that $\vec{A}'_{50}$ can drop below 0.5. Here this indicates that in some cases the effect of the illuminant change is to make the response difference smaller on different trials than on same trials.

TRADEOFFS BETWEEN DISCRIMINATION AND CONSTANCY

When the adaptation parameters include a gain change, adaptation can compensate perfectly for changes in illumination so that discrimination performance $\bar{A}'_{50}$ remains unchanged and constancy performance is at ceiling given discrimination ($\vec{A}'_{50} = \bar{A}'_{50}$.) More generally there will be cases where the adaptation parameters available within a given model are not able to compensate completely for environmental changes. This raises the possibility that the adaptation parameters that optimize discrimination performance may differ from those that optimize constancy performance.

Consider the case of an illuminant change where the gain parameter is held fixed and the steepness parameter n is allowed to vary. Figure 4 plots $\vec{A}'_{50}$ against $\bar{A}'_{50}$ for various choices of the steepness parameter. The closed circle plots performance when the steepness parameter was chosen to maximize $\bar{A}'_{50}$, while the open circle plots performance when the steepness parameter was chosen to maximize $\vec{A}'_{50}$. The curve connecting the closed and open circle represents performance obtained when various weighted averages of the two measures were optimized. The figure shows that in this case, there is a tradeoff between the two performance measures – optimizing for constancy results in decreased discrimination performance and vice-versa. If there were no tradeoff between the two measures, performance at the point in the plot indicated by the star would be possible.

Figure 4. Tradeoff between discrimination and constancy. The plot shows $\vec{A}'_{50}$ versus $\bar{A}'_{50}$ for different choices of steepness parameter n. Discrimination performance $\bar{A}'_{50}$ was evaluated for $\mu_r = 0.5$, $\mu_r = 0.5$, $e_i = 150$, $\sigma_n = 0.05$, and $g = 0.0203$. Constancy performance $\vec{A}'_{50}$ was evaluated for the same surface ensemble, noise level, and gain, with the reference illuminant $e_{ref} = 100$ and test illuminant $e_{test} = 150$. The steepness

parameter under the reference illuminant was $n = 4.272$. Points on the plotted curve were obtained by optimizing n to maximize $\overline{A}'_{50}$ (closed circle), $\vec{A}'_{50}$ (open circle), or various weighted averages of the two measures. The star indicates the joint performance that could be achieved if the same state of adaptation could simultaneously optimize both discrimination and constancy.

INTERMEDIATE DISCUSSION

The example above illustrates our basic approach to understanding how adaptation affects both discrimination and constancy. The example illustrates a number of key points. First, as is well known, adaptation is required to maintain optimal discrimination performance across changes in the state of the visual environment (e.g. Walraven *et al.*, 1990). Second, adaptation is necessary to optimize performance on a constancy task as well as for discrimination, when we require that surface identity is judged directly in terms of the visual responses. The link between adaptation and constancy has also been explored previously (e.g. Burnham *et al.*, 1957; D'Zmura & Lennie, 1986; Wandell, 1995). What is new about our approach is that we have set our evaluations of both discrimination and constancy in a common framework by using an A' measure for both. This allows us to ask questions about how any given adaptation strategy affects both tasks, and whether common mechanisms of adaptation can simultaneously optimize performance for both. The theory we develop is closely related to measurements of lightness constancy made by Robilotto and Zaidi (2004), who used a forced choice method to measure both discrimination within and identification across changes in illuminant.

Below we employ our basic approach to analyze two more interesting examples of adaptation: a) a generalization of contrast adaptation for a univariate visual system and b) adaptation for a trichromatic visual system.

Contrast Adaptation

Changing the illuminant is not the only way to change the properties of the environment. Within the context of the univariate case introduced above, we can also vary both the mean and variance of the surface ensemble. Such variation might occur as a person travels from, say, the city to the suburbs during an afternoon commute. We can use the framework developed above to investigate the effect of adaptation to such changes in the visual environment. There is good evidence that the visual system adapts not only to changes in the overall intensity of the reflected light (as considered above in the context of illumination intensity changes) but also to changes in the variance of the reflected light. This is called contrast adaptation (Krauskopf *et al.*, 1982; Webster & Mollon, 1991; Chubb *et al.*, 1989; Zaidi & Shapiro, 1993; Jenness & Shevell, 1995; Schirillo & Shevell, 1996; Brown & MacLeod, 1997; Bindman & Chubb, 2004; Chander & Chichilnisky, 2001; Solomon *et al.*, 2004). Here we generalize by considering changes to the mean and variance of the surface ensemble, and explore the effect of adapatation to such changes on both discrimination and constancy, using the approach developed above.

We consider the same visual environment as previously, but now ask what happens when the properties of the surface ensemble vary. Given a particular illuminant, we used numerical search to find the values of g and n that optimized $\overline{A}'_{50}$ for two fixed surface ensembles (Surface Ensemble 1 and Surface Ensemble 2) under a single illuminant. This calculation tells us how adaptation should be chosen under a discrimination criterion. Figure 5 shows the results. We see (central panel) that the visual response function under Surface Ensemble 2 has shifted to the right and become steeper. The effect of this adaptation is to distribute the visual responses fairly evenly across the available response range for each ensemble.

Figure 5. Adaptation to surface ensemble change for discrimination. Numerical search was used to optimize $\bar{A}'_{50}$ for two surface ensembles and an illuminant intensity of 100. In surface Ensemble 1, $\mu_r = 0.5$ and $\sigma_r = 0.3$. In surface Ensemble 2, $\mu_r = 0.7$ and $\sigma_r = 0.1$. The histogram under the graph shows the distributions of reflected light intensities for the two ensembles. The graph shows the resultant visual response function for each case. The solid line corresponds to Surface Ensemble 1 ($g = 0.0203$, $n = 4.272$, $\bar{A}'_{50} = 0.79$) and the dashed line to Surface Ensemble 2 ($g = 0.0144$, $n = 16.214$, $\bar{A}'_{50} = 0.80$). The histogram to the left of the graph shows the response distribution for Surface Ensemble 1 under the Surface Ensemble 1 response function, while the histogram to the right shows the response distribution for Surface Ensemble 2 under the Surface Ensemble 2 response function. All calculations done for $\sigma_n = 0.05$ and $e = 100$. The gray long-dashed lines show how the visual response to the light intensity reflected from a fixed surface varies with the change in adaptation parameters. (Since the illuminant is held constant, a fixed surface corresponds to a fixed light intensity.)

What is the cost for constancy of adaptation to the changed surface ensemble for discrimination? The gain and steepness parameters that optimize discrimination for Surface Ensemble 1 and Surface Ensemble 2 are different. Suppose we draw surfaces from Surface Ensemble 2, and evaluate $\vec{A}'_{50}$ across the change in adaptation parameters required to optimize discrimination. The resultant value is very low: $\vec{A}'_{50} = 0.1377$ (see Figure 6, solid circle.) This low value occurs because the change in adaptation parameters remaps the relation between surface reflectance and visual response (see Figure 5.)

Figure 6. Tradeoff between discrimination and constancy. The plot shows $\vec{A}'_{50}$ versus $\bar{A}'_{50}$, evaluated for Surface Ensemble 2, for different optimization criteria. When the adaptation parameters are chosen to optimize discrimination ($\bar{A}'_{50}$), discrimination performance is maximized and constancy performance is poor (sold circle). When the adaptation parameters are chosen to optimize constancy ($\vec{A}'_{50}$), constancy performance is maximized and discrimination performance is poor (open circle). The solid line shows how performance on the two tasks trades off as varying weight is placed on the two performance criteria during optimization. In evaluating $\vec{A}'_{50}$, the adaptation parameters

were those that optimize discrimination performance for Surface Ensemble 1. Surface ensemble, illuminant, and noise parameters are given in the caption for Figure 5.

Rather than choosing adaptation parameters for Surface Ensemble 2 to optimize discrimination performance, one can instead choose them to optimize constancy performance. This choice leads to quite different adaptation parameters, and to different values of $\bar{A}'_{50}$ and $\vec{A}'_{50}$ as shown by the open circle in Figure 6. Here the best adaptation parameters for Surface Ensemble 2 are very similar (but not identical) to their values for Surface Ensemble 1, and constancy performance is better.[9] Discrimination performance suffers, however. It is also worth noting that maximized constancy performance is not particularly good, with the value of $\vec{A}'_{50}$ less than 0.6. For this example, it is not possible to simultaneously make good use of the response range for Surface Ensemble 2 and preserve constant responses across the change from Surface Ensemble 1 to 2.

More generally, the visual system can choose adaptation parameters that trade $\bar{A}'_{50}$ off against $\vec{A}'_{50}$. The performance boundary that may be obtained as the relative weight on these two criteria is varied is shown as the solid curve in Figure 6. If there were no tradeoff between $\bar{A}'_{50}$ and $\vec{A}'_{50}$, the curve would follow the dashed lines in the plot, and performance at the point indicated by the star would be possible. The calculation shows that, for the model visual system considered, adapting to optimize discrimination in the

[9] One might initially intuit that that the best adaptation parameters for constancy would be identical to the reference parameters in this case, since the illuminant does not change. The reason that a small change in parameters helps is that the cost of variation in visual response to a fixed surface caused by the shift in parameters is offset by an improved use of the available response range. The fact that constancy can sometimes be improved by changing responses to fixed surfaces is an insight that we obtain by assessing constancy with a signal detection theoretic measure.

face of changes in the distribution of surfaces in the environment is not always compatible with adapting to maximize constancy.

CHROMATIC ADAPTATION

The same approach we used above may be applied to study chromatic adaptation, both for discrimination and constancy. We begin with a standard description (e.g. Wandell, 1987; Brainard, 1995) of the color stimulus and its initial encoding by the visual system. Each illuminant is specified by its spectral power distribution, which we represent by a column vector $\mathbf{e}$. The entries of $\mathbf{e}$ provide the power of the illuminant in a set of N_λ discretely sampled wavelength bands. Each surface is specified by its spectral reflectance function, which we represent by a column vector $\mathbf{s}$. The entries of $\mathbf{s}$ provide the fraction of incident light power reflected in each wavelength band. The light reflected to the eye has a spectrum described by the column vector

$$\mathbf{c} = \mathbf{diag}(\mathbf{e})\,\mathbf{s}\,, \tag{4}$$

where $\mathbf{diag}()$ is a function that returns a square diagonal matrix with the elements of its argument along the diagonal. The initial encoding of the reflected light is the quantal absorption rate of the L-, M-, and S-cones. We represent the spectral sensitivity of the three cone types by a $3 \times N_\lambda$ matrix $\mathbf{T}$. Each row of $\mathbf{T}$ provides the sensitivity of the corresponding cone type (L, M, or S) to the incident light. The quantal absorption rates may then by computed as

$$\mathbf{q} = \mathbf{T}\,\mathbf{c} = \mathbf{T}\mathbf{diag}(\mathbf{e})\mathbf{s} \tag{5}$$

where $\mathbf{q}$ is a 3-dimensional column vector whose three entries represent the L-, M-, and S-cone quantal absorption rates.

As with the univariate example, we model visual processing as transformation between cone quantal absorptions ($\mathbf{q}$) to visual responses. Here we model the deterministic component of this transformation as

$$\mathbf{u} = \mathbf{f}(\mathbf{M}\mathbf{D}\mathbf{q} - \mathbf{q}_0) \tag{6}$$

where $\mathbf{u}$ is a 3-dimensional column vector representing trivariate visual responses, $\mathbf{D}$ is a 3×3 diagonal matrix whose entries specify multiplicative gain control applied to the cone quantal absorbtion rates, $\mathbf{M}$ is a fixed 3×3 matrix that describes a post-receptoral recombination of cone signals, $\mathbf{q}_0$ is a 3-dimensinal column vector that describes subtractive adaptation, and the vector-valued function $\mathbf{f}()$ applies the function $f_i()$ to the i^{th} entry of its vector argument. Because incorporation of subtractive adaptation allows the argument to the non-linearity to be negative, we used a modified form of the non-linearity used in the univariate example:

$$f_i(x) = \begin{cases} \dfrac{(x+1)^{n_i}}{(x+1)^{n_i}+1} & x > 0 \\ 0.5 & x = 0 \\ 1 - \dfrac{(1-x)^{n_i}}{(1-x)^{n_i}+1} & x < 0 \end{cases} \quad . \tag{7}$$

This non-linearity maps input x in the range $[-\infty, \infty]$ from the real line into the range $[0,1]$. We allow the exponent n_i to vary across entries. The matrix $\mathbf{M}$ was chosen to model, in broad outline, the post-receptoral processing of color information (Wandell, 1995; Kaiser & Boynton, 1996; Eskew *et al.*, 1999; Brainard, 2001):

$$\mathbf{M} = \begin{bmatrix} 0.33 & 0.33 & 0.33 \\ 0.5 & -0.5 & 0 \\ -0.25 & -0.25 & 0.5 \end{bmatrix}. \tag{8}$$

This choice of $\mathbf{M}$ improves discrimination performance by approximately decorrelating the three entries of the visual response vector prior to application of the non-linearity and the injection of noise (Buchsbaum & Gottschalk, 1983; Wandell, 1995).

As with the univariate example, we assume that each entry of the deterministic component of the visual response vector is perturbed by independent zero-mean additive noise that is Normally distributed with variance σ_n^2.

We characterized the surface ensemble using the approach developed by Brainard and Freeman (Brainard & Freeman, 1997; Zhang & Brainard, 2004). We assumed that the spectral reflectance of each surface could be written as a linear combination of N_s basis functions via

$$\mathbf{s} = \mathbf{B}_s \mathbf{w}_s . \tag{9}$$

Here $\mathbf{B}_s$ is an $N_\lambda \times N_s$ matrix whose columns provide the basis functions and $\mathbf{w}_s$ is a N_s dimensional vector whose entries provide the weights that describe any particular surface as a linear combination of the columns of $\mathbf{B}_s$. We then assume that surfaces are drawn from an ensemble where $\mathbf{w}_s$ is drawn from a truncated multivariate Normal distribution with mean vector $\bar{\mathbf{w}}$ and covariance matrix $\mathbf{K}_w$. The truncation is chosen so that the reflectance in each wavelength band lies within the range $[0,1]$. We computed $\mathbf{B}_s$ by computing the first 8 principal components of the reflectance spectra measured by Vrhel et al. (1994). We computed $\bar{\mathbf{w}}$ and $\mathbf{K}_w$ by taking the mean and covariance of the set of $\mathbf{w}_s$ required to best approximate each of the measured spectra with respect to $\mathbf{B}_s$.

Given the visual system model and surface ensemble defined above, we can proceed as with the univariate case and ask how the adaptation parameters affect $\overline{A}'_{50}$ and $\vec{A}'_{50}$, the discrimination and constancy performance measures respectively. The only modification required is that the decision rule now operates on the difference variable $\Delta u^2_{i,jk} = \left\| \mathbf{u}_{i,j} - \mathbf{u}_{i,k} \right\|^2$, and this variable is distributed as a non-central chi-squared distribution with 3 degrees of freedom rather than 1. The adaptation parameters are the three diagonal entries of $\mathbf{D}$, the three entries of $\mathbf{q}_0$, and the three exponents n_i. Figure 7 shows how $\overline{A}'_{50}$ and $\vec{A}'_{50}$ trade off when the illuminant is changed from CIE illuminant D65 to a CIE daylight with correlated color temperature of 20000° K. The figure shows that here discrimination and constancy are highly compatible.

Figure 7. Tradeoff between discrimination and constancy, color case for illuminant change. The plot shows $\vec{A}'_{50}$ versus $\overline{A}'_{50}$, evaluated for the surface ensemble described in the text under a CIE daylight with correlated color temperature of 20000° K, for different optimization criteria. In evaluating $\vec{A}'_{50}$, the reference illuminant was CIE D65 and the adaptation parameters for this illuminant were those that optimize discrimination performance under it. Note the scale of the plot.

We also investigated discrimination constancy tradeoffs for the color case when the surface ensemble is changed. Here we find that, as with the univariate case, there is not a single set of parameters that simultaneously optimize constancy and discrimination for the two tasks (Figure 8).[10]

[10] The numerical optimization for this case proved difficult because the searches tended to converge on local rather than global minima. This in turn made it difficult to identify many distinct points on the tradeoff curve. The data shown represent essentially three distinct "islands" of performance, one at the open circle, one at the closed circle, and one

Figure 8. Tradeoff between discrimination and constancy, color case for surface ensemble change. The plot shows $\vec{A}'_{50}$ versus $\overline{A}'_{50}$ under D65 across a surface ensemble change. The reference surface ensemble was that used above and described in the text. The changed surface ensemble was obtained from the reference ensemble by shifting the mean of the weights of the first two basis functions, decreasing the standard deviation of the weights of the first basis function, increasing the standard deviation of the weights for the second basis function, and forcing the covariance matrix for the basis function weights to be diagonal. The reference adaptation parameters were those that optimize discrimination for the reference surface ensemble under D65. Note the scale of the plot.

SUMMARY AND DISCUSSION

The theory and calculations presented here lead to several broad conclusions. First, we note that constancy cannot be evaluated meaningfully without considering discrimination. By using a signal detection theoretic measure (A') to quantify constancy, we explicitly incorporate discrimination into our treatment of constancy.

When the environmental change is a change in illuminant, then the dual goals of discrimination and constancy are compatible: a common change in adaptation parameters optimizes performance for both our discrimination and constancy performance measures. This is trivially true for our univariate example, where a change in gain can compensate perfectly for a change in illuminant intensity. It is true to good approximation for our chromatic example, where the adaptation parameters allow approximate compensation for the effect of a change in illuminant spectrum on the cone quantal absorption rates of reflected light (Figure 7). Our conclusion here is consistent with previous analyses that have suggested that independent gain changes in the L-, M-, and S-cone pathways can stabilize responses to a fixed set of surfaces across many changes in illuminant (West & Brill, 1982; Brainard & Wandell, 1986; Foster & Nascimento, 1994; Foster *et al.*, 2004; also Finlayson *et al.*, 1994; Finlayson & Funt, 1996).

---

at the "knee" of the plot. We were unable to find adaptation parameters that produced performance intermediate between the "islands" shown. This latter effect also occurred (to a lesser extent) for the tradeoff curve shown in Figure 6.

When the environmental change is a change in the surface ensemble, discrimination and constancy are not always compatible: Figures 6 and 8 show that optimizing the adaptation parameters of the models we have studied for discrimination comes at a substantial cost for constancy and vice-versa. Thus the analysis suggests that stimulus conditions where the surface ensemble changes may provide data that are diagnostic of whether the early visual system optimizes for discrimination, for constancy, or whether it has evolved separate sites that mediate performance on the two tasks. For example, if we imagine the limit on discrimination to occur at the level of retinal ganglion cells, whose performance is very roughly described by the models we consider here, then one might also imagine that cortical processing could act on the retinal output to produce an appearance representation that is stable across adaptation to changes in surface ensemble. For this to be possible, the cortical processing would need access to the values of the adaptation parameters as well as to the retinal output. We have started to develop an experimental framework for approaching this question (see Hillis & Brainard, in press; also Robilotto & Zaidi, 2004).

As noted in the introduction, our approach is similar to that of Grzywacz and colleagues (Grzywacz & Balboa, 2002; Grzywacz & Juan, 2003; also Foster *et al.*, 2004). Previous authors have considered the nature and adaptation of the visual response function required to optimize discrimination performance (e.g. Laughlin; Buchsbaum & Gottschalk, 1983; Twer & MacLeod, 2001) as well as nature of adaptive transformations that can mediate constancy (e.g. von Kries, 1902; Buchsbaum, 1980; West & Brill, 1982; Brainard & Wandell, 1986; D'Zmura & Lennie, 1986; Maloney & Wandell, 1986; Foster & Nascimento, 1994; Foster *et al.*, 2004; Finlayson *et al.*, 1994; Finlayson & Funt, 1996).

Here we bring analysis of the two tasks together in a common signal detection theoretic framework, and treat the compatibility of adaptation for discrimination and constancy.

ACKNOWLEDGMENTS

Dan Lichtman helped with early versions of the computations. We thank M. Kahana and the Penn Psychology Department for access to computing power. Supported by NIH RO1 EY10016.

REFERENCES

Bindman, D., & Chubb, C. (2004). Mechanisms of contrast induction in heterogeneous displays. *Vision Research, 44*, 1601-1613.

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

Brainard, D. H. (2001). Color vision theory. In N. J. Smelser & B. P.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 L. Chalupa & J. Werner (Eds.), *The visual neurosciences* (Vol. 1, pp. 948-961). Cambridge, MA, U.S.A.: MIT Press.

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., & 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.

Brown, R. O., & MacLeod, D. I. A. (1997). Color appearance depends on the variance of surround colors. *Current Biology, 7*, 844-849.

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

Buchsbaum, G., & Gottschalk, A. (1983). Trichromacy, opponent colours coding and optimum colour information transmission in the retina. *Proceedings of the Royal Society of London B, 220*, 89-113.

Burnham, R. W., Evans, R. M., & Newhall, S. M. (1957). Prediction of color appearance with different adaptation illuminations. *Journal of the Optical Society of America, 47*, 35-42.

Chander, D., & Chichilnisky, E. J. (2001). Adaptation to temporal contrast in primate and salamander retina. *Journal of Neuroscience, 21*, 9904-9916.

Chubb, C., Sperling, G., & Solomon, J. A. (1989). Texture interactions determine perceived contrast. *PNAS, 86*, 9631-9635.

Cover, T. M., & Thomas, J. A. (1991). *Elements of informtion theory*. New York: John Wiley and Sons.

D'Zmura, M., & Lennie, P. (1986). Mechanisms of color constancy. *Journal of the Optical Society of America A, 3*, 1662-1672.

Eskew, R. T., McLellan, J. S., & Giulianini, F. (1999). Chromatic detection and discrimination. In K. Gegenfurtner & L. T. Sharpe (Eds.), *Color vision: From* *molecular genetics to perception* (pp. 345-368). Cambridge: Cambridge University Press.

Finlayson, G. D., Drew, M. S., & Funt, B. V. (1994). Color constancy - generalized diagonal transforms suffice. *Journal of The Optical Society of America A, 11*(11), 3011-3019.

Finlayson, G. D., & Funt, B. V. (1996). Coefficient channels - derivation and relationship to other theoretical studies. *Color Research and Application, 21*(2), 87-96.

Foster, D. H. (2003). Does colour constancy exist? *Trends in Cog Sci, 7*(10), 493-443.

Foster, D. H., & Nascimento, S. M. C. (1994). Relational colour constancy from invariant cone-excitation ratios. *Proc. R. Soc. Lond. B, 257*, 115-121.

Foster, D. H., Nascimento, S. M. C., & Amano, K. (2004). Information limits on neural identification of colored surfaces in natural scenes. *Visual Neuroscience, 21*, 1-6.

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

Grzywacz, N. M., & Balboa, R. M. (2002). A bayesian framework for sensory adaptation. *Neural Computation, 14*, 543-559.

Grzywacz, N. M., & Juan, J. d. (2003). Sensory adaptation as kalman filtering: Theory and illustration with contrast adaptation. *Network: Computation in Neural Systems, 14*, 465-482.

Hillis, J. M., & Brainard, D. H. (in press). Do common mechanisms of adaptation mediate color discrimination and appearance? Uniform backgrounds. *Journal of the Optical Society of America A*.

Hood, D. C., & Finkelstein, M. A. (1986). Senstivity to light. In K. R. Boff, L. Kaufman & J. P. Thomas (Eds.), *Handbook of perception and human performance: Sensory processes and perception* (Vol. 1): John Wiley & Sons.

Hurlbert, A. C. (1998). Computational models of color constancy. In V. Walsh & J. Kulikowski (Eds.), *Perceptual constancy: Why things look as they do* (pp. 283-322). Cambridge: Cambridge University Press.

Jacobs, G. H. (1981). *Comparative color vision*. New York: Academic Press.

Jenness, J. W., & Shevell, S. K. (1995). Color appearance with sparse chromatic context. *Vision Research, 35*(6), 797-805.

Kaiser, P. K., & Boynton, R. M. (1996). *Human color vision* (2 ed.). Washington, D.C.: Optical Society of America.

Krantz, D. (1968). A theory of context effects based on cross-context matching. *Journal of Mathematical Psychology, 5*, 1-48.

Krauskopf, J., Williams, D. R., & Heeley, D. W. (1982). Cardinal directions of color space. *Vision Research, 22*, 1123-1131.

Laughlin, S. (1989). The reliability of single neurons and circuit design: A case study. In R. Durbin, C. Miall & G. MItchison (Eds.), *The computing neuron* (pp. 322-335). New York: Addison-Wesley.

Lennie, P., & D'Zmura, M. (1988). Mechanisms of color vision. *CRC Critical Reviews in Neurobiology, 3*(4), 333-400.

Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory: A users guide* (2cd edition ed.). Mahwah, New Jersey: Lawrence Erlbaum Associates.

Maloney, L. T. (1999). Physics-based approaches to modeling surface color perception. In K. T. Gegenfurtner & L. T. Sharpe (Eds.), *Color vision: From genes to perception* (pp. 387-416). Cambridge: Cambridge University Press.

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.

Mollon, J. D. (1982). Color vision. *Annual Review of Psychology, 33*, 41-85.

Robilotto, R., & Zaidi, Q. (2004). Limits of lightness identification for real objects under natural viewing conditions. *J Vis, 4*(9), 779-797.

Schirillo, J. A., & Shevell, S. K. (1996). Brightness contrast from inhomogeneous surrounds. *Vision Research, 36*(12), 1783-1796.

Shevell, S. K. (2003). Color appearance. In S. K. Shevell (Ed.), *The science of color, 2cd edition*. Washington, D.C.: Optical Society of America.

Solomon, S. G., Peirce, J. W., Dhruv, N. T., & Lennie, P. (2004). Profound contrast adaptation early in the visual pathway. *Neuron, 42*, 155-162.

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

Twer, T. v. d., & MacLeod, D. I. A. (2001). Optimal nonlinear codes for the perception of natural colors. *Network: Computation in Neural Systems, 12*, 395-407.

von der Twer, T., & MacLeod, D. I. A. (2001). Optimal nonlinear codes for the perception of natural colours. *Network: Computation in Neural Systems, 12*, 395-407.

von Kries, J. (1902). Chromatic adaptation. In: Originally published in *Festschrift der Albrecht-Ludwigs-Universitat*, 145-148. In MacAdam, D.L., ed., *Sources of Color Vision*. (MIT Press: Cambridge, 1970).

Vrhel, M. J., Gershon, R., & Iwan, L. S. (1994). Measurement and analysis of object reflectance spectra. *Color Research And Application, 19*(1), 4-9.

Walraven, J., Enroth-Cugell, C., Hood, D. C., MacLeod, D. I. A., & Schnapf, J. L. (1990). The control of visual sensitivity. Receptoral and postreceptoral processes. In L. Spillman & J. S. Werner (Eds.), *Visual perception: The neurophysiological foundations* (pp. 53-101). San Diego: Academic Press.

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

Wandell, B. A. (1995). *Foundations of vision*. Sunderland, MA: Sinauer.

Webster, M. A., & Mollon, J. D. (1991). Changes in colour appearance following post-receptoral adaptation. *Nature, 349*(6306), 235-238.

West, G., & Brill, M. H. (1982). Necessary and sufficient conditions for von kries chromatic adaptation to give color constancy. *J. Theor. Biolog., 15*, 249.

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

Wyszecki, G., & Stiles, W. S. (1982). *Color science - concepts and methods, quantitative data and formulae* (2 ed.). New York: John Wiley & Sons.

Zaidi, Q. (1999). Color and brightness inductions: From mach bands to three-dimensional configurations. In K. R. Gegenfurtner & L. T. Sharpe (Eds.), *Color vision: From* *genes to perception* (Vol. 1, pp. 317-344). Cambridge, U.K.: Cambridge University Press.

Zaidi, Q., & Shapiro, A. G. (1993). Adaptive orthogonalization of opponent-color signals. *Biological Cybernetics, 69*(5-6), 415-428.

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.

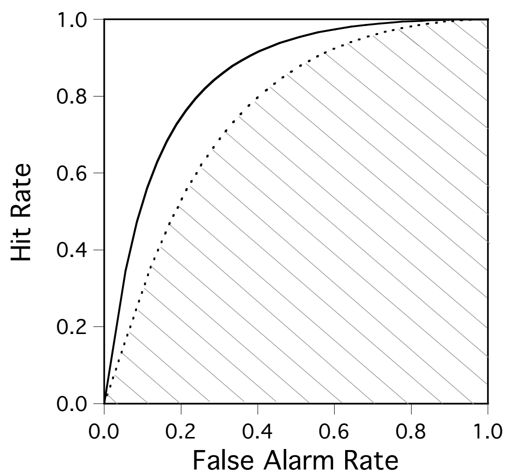

Figure 1

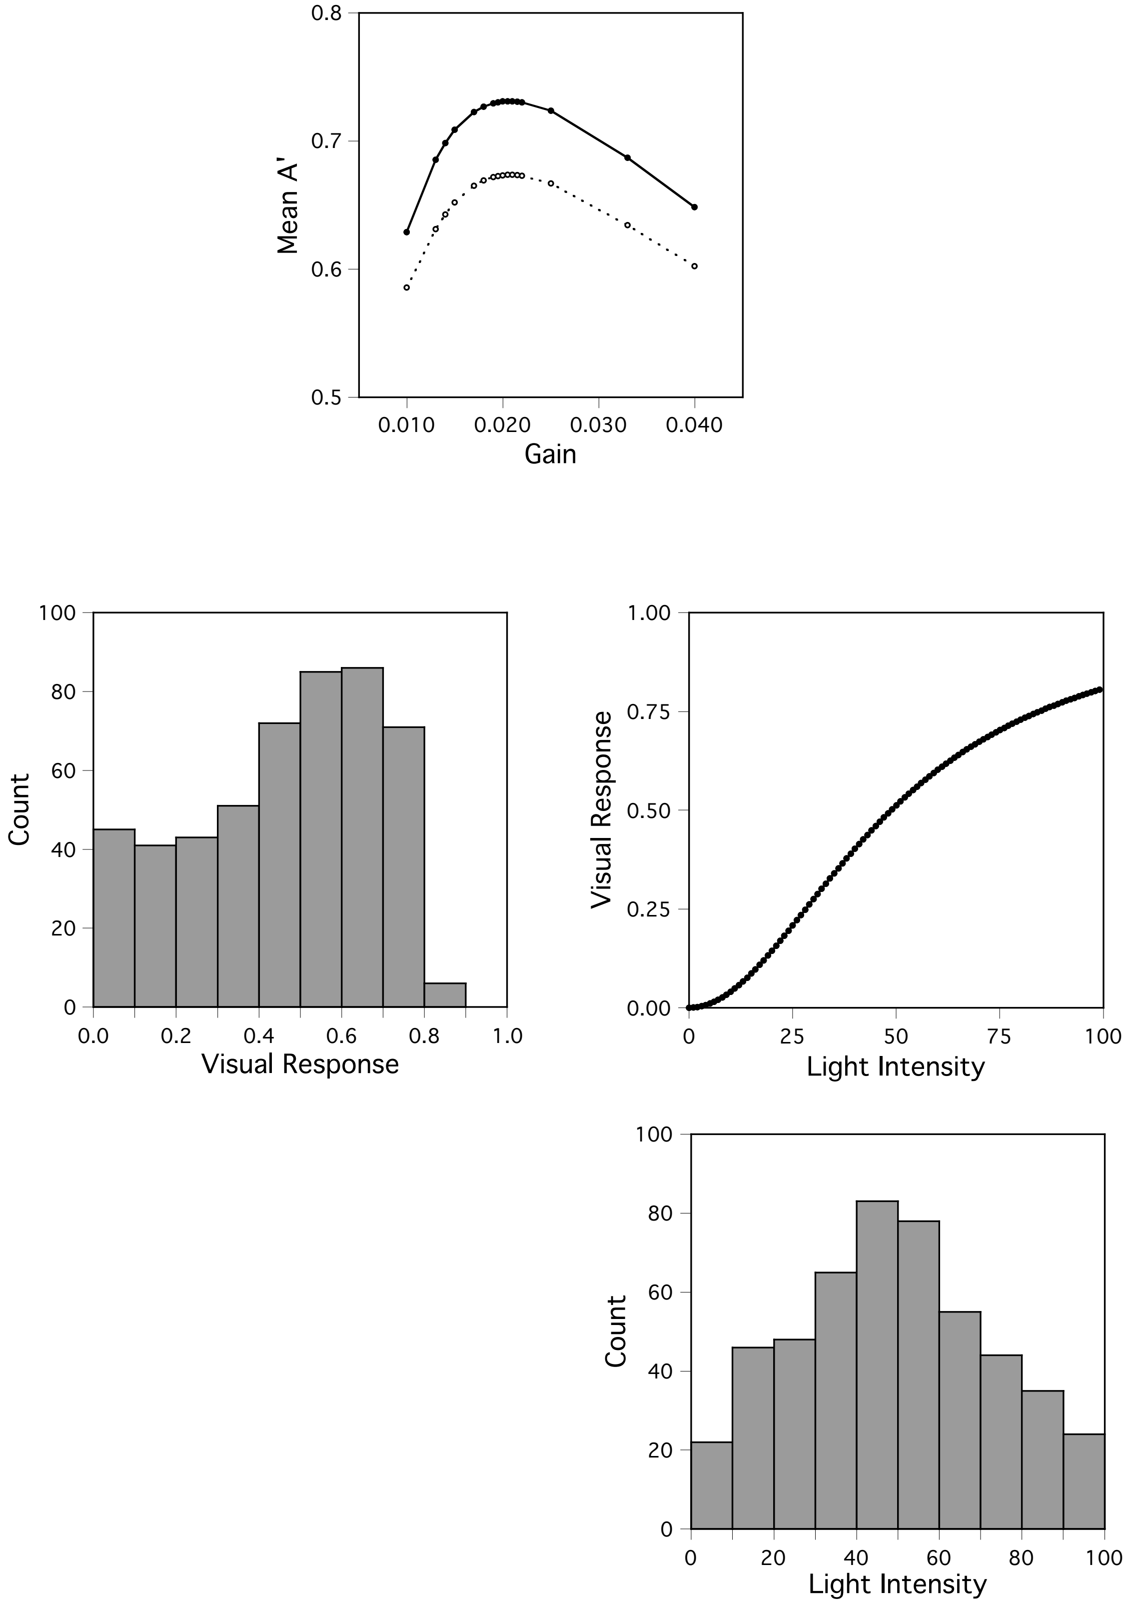

0.8
0.7
0.6
0.5
Mean A'
0.010
0.020
0.030
0.040
Gain
100
80
60
40
20
0
Count
0.0
0.2
0.4
0.6
0.8
1.0
Visual Response
1.00
0.75
0.50
0.25
0.00
Visual Response
0
25
50
75
100
Light Intensity
100
80
60
40
20
0
Count
0
20
40
60
80
100
Light Intensity

Figure 2

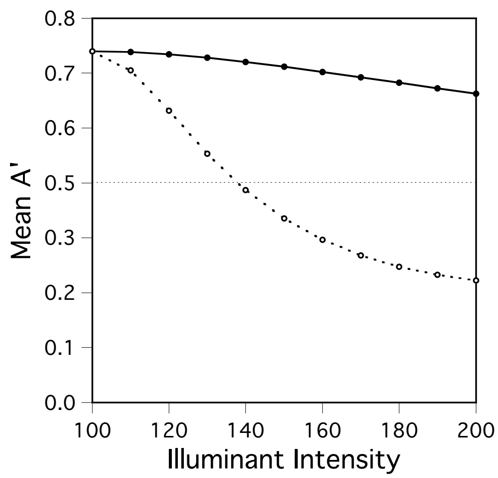

0.8
0.7
0.6
0.5
0.3
0.2
0.1
0.0
Mean A'
100
120
140
160
180
200
Illuminant Intensity

Figure 3

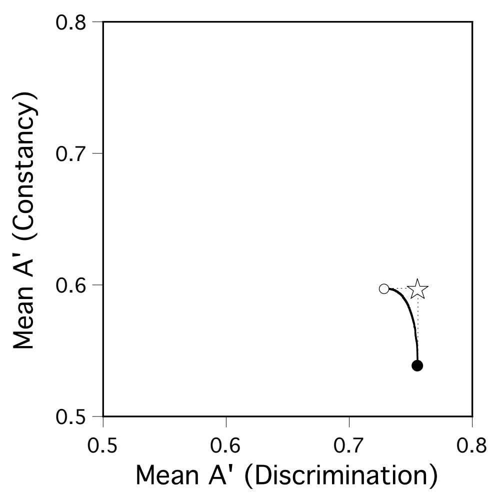

Figure 4

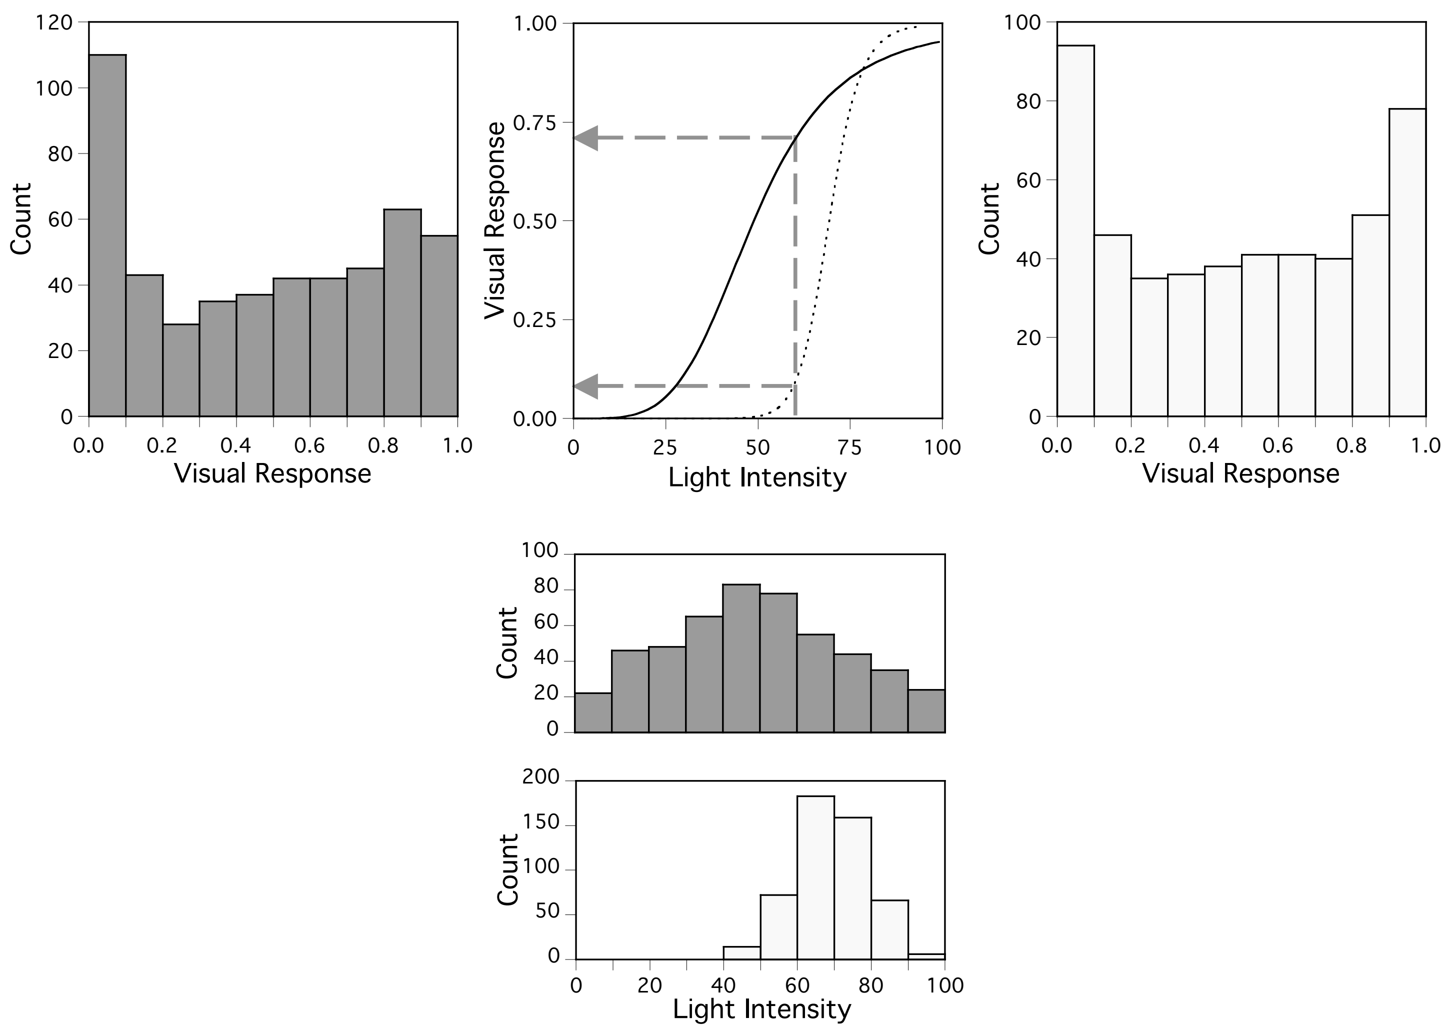

Figure 5

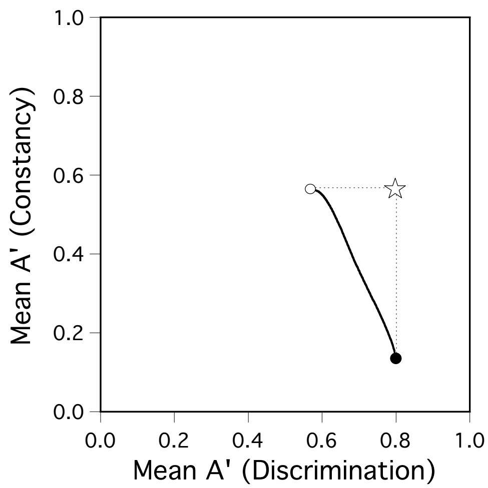

1.0
0.8
0.6
0.4
0.2
0.0
Mean A' (Constancy)
0.0
0.2
0.4
0.6
0.8
1.0
Mean A' (Discrimination)

Figure 6

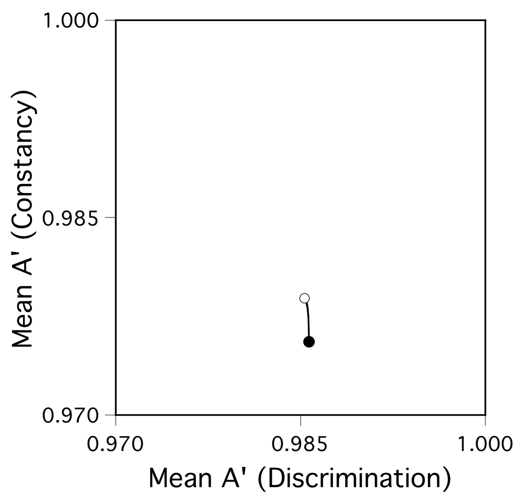
1.000
0.985
0.970
Mean A' (Constancy)
0.970
0.985
1.000
Mean A' (Discrimination)

Figure 7

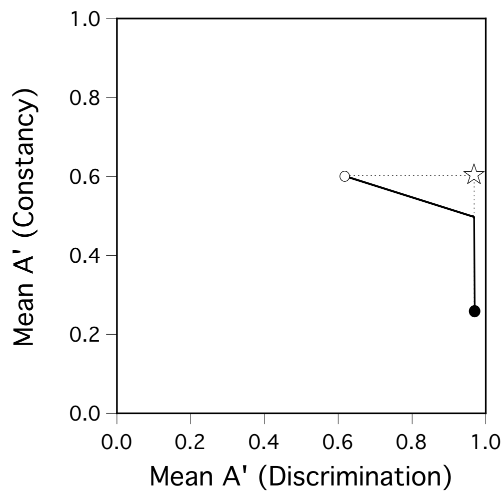
Mean A' (Constancy)
1.0
0.8
0.6
0.4
0.2
0.0
0.0
0.2
0.4
0.6
0.8
1.0
Mean A' (Discrimination)

Figure 8