Chapter 11

Approaching Color with Bayesian Algorithms

Sarah Allred

What is the goal of color vision? How ought we to think of color appearance? Under one view, the goal of vision is to maintain a stable representation of object properties across changes in the environment. This poses a challenge to the visual system, because the sensory signal on which visual perception is based is ambiguous with respect to the physical properties of objects in the world. Thus, to maintain stable color appearance, the visual system must estimate what object was most likely to have caused the ambiguous sensory signal (Helmholtz 1910/1924). In this chapter, I present a Bayesian approach to solving this estimation problem that relies on statistical regularities in the world to resolve the sensory ambiguity. This is a sensible idea: the human visual system evolved in this world, and thus its statistical regularities are likely to be of functional importance to vision (Attneave 1954; Geisler and Kersten 2002; Mamassian et al. 2002; Weiss et al. 2002).

Any functional understanding of vision has to deal with the problem of information loss between the distal object and its visual representation. In color vision, there are several salient sources of information loss. First, there is ambiguity inherent in image formation, illustrated in Figure 11.1. In color vision, that ambiguity takes the form that the light reaching the eye from a surface j [$C_j(\lambda)$, sensory signal] confounds reflectance properties of the surface [$S_j(\lambda)$] with the properties of the illuminant [$E_j(\lambda)$]. Color perception typically follows the properties of a surface rather than the illuminant. However, the sensory signal that gives rise to color perception is ambiguous with respect to which surface gave rise to it; that is, infinitely many combinations of surfaces and illuminants could have caused any given sensory signal. Despite this ambiguity in the sensory signal, our perception is (usually) not ambiguous; that is, at any given time, we only perceive one color at a particular spatial location.

A second kind of information loss occurs because of the physiological properties of the photoreceptors. As illustrated in Figure 11.1, under normal light levels, human vision relies upon three different photoreceptor classes, L-, M-, and S-cones. Each cone type is most sensitive to a different wavelength of light; color vision relies on the comparison of the responses of these three cone types to a given sensory signal. Thus, rather than representing the intensity of each wavelength in the sensory signal, the photoreceptors can be characterized as giving three discrete intensity values (r_j). Despite this trichromatic representation of a continuous color spectrum at the earliest level of the physiology, a continuous color circle still captures many aspects of color experience.

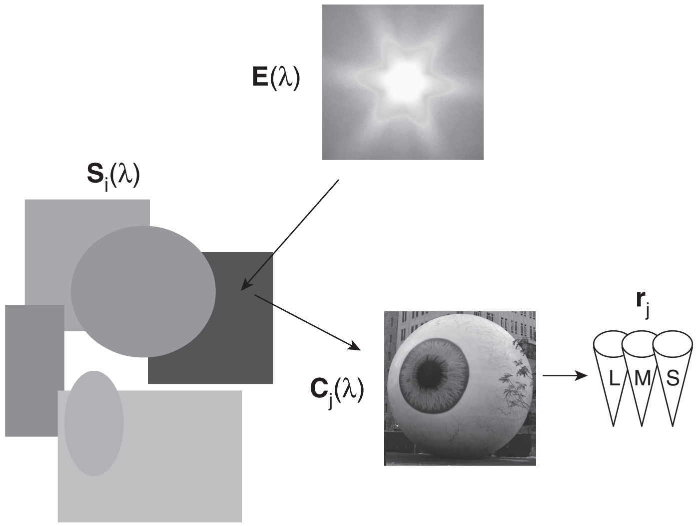

Fig. 11.1 Illustration of sensory ambiguity. (© Maria Olkkonen, 2011, reproduced here with permission.) (See also Plate 10.)

Third, information loss occurs because of noise in the physiological system. Photoreceptors signal the presence of light to subsequent parts of the visual system through a change in electrical current flowing through the photoreceptor; that change is mediated by photopigment isomerization. When a photopigment absorbs light, its shape changes. This change in shape initiates a cascade of events that eventually alters the amount of neurotransmitter released by the photoreceptor, which in turn affects the responses of other neurons in the visual system. However, the responses of photoreceptors are noisy; that is, subsequently presented identical stimuli can cause different numbers of photopigment isomerizations. In addition, sometimes photopigments isomerize in the absence of light. In high-light conditions, photopigment isomerizations are dominated by the light incident on the photoreceptors; however, at low illumination levels, isomerizations due to chance may be closer in magnitude to the response evoked by a visual stimulus. Particularly at low-light levels, photoreceptor noise can have interesting perceptual effects.

To perceive, one must guess

Although the ambiguities described here are particular to color perception, sensory ambiguity is fundamental to all visual perception. For any visual property (e.g., color, size, shape, distance, depth, gloss), reconstruction of the distal stimulus from the sensory signal is an underdetermined problem. In other words, there is no theoretical way to recover the distal stimulus from the sensory signal without applying some sort of guessing rule, heuristic, or external constraint. Although some descriptions of visual processing leave these guessing rules unstated, any explanation of visual perception tied to a reconstruction of the distal stimulus requires guessing rules. These rules may be stated or unstated, formulated implicitly or explicitly, described functionally or mechanistically, and implemented qualitatively or quantitatively, but they are necessary for a complete characterization of visual perception. To perceive, one must guess. Moreover, one set of guessing rules is not enough; visual constancy (or the ability to perceive an object the same way in different environments) requires that guessing rules change from environment to environment.[1]

How ought the visual system to deal with this information loss? Given that the visual system must guess, how ought it to guess best? There are a number of possible answers. In this chapter, I describe a Bayesian computational approach to formulating these guessing rules. To motivate the description and evaluation of the Bayesian approach, I first outline some perennial problems of visual perception. Second, I outline the general principles of Bayesian algorithms. Importantly, the Bayesian approach to visual perception as described here does not aim for veridicality per se; rather, it aims to describe human perception. Third, I illustrate the utility of the Bayesian approach to two specific cases, color constancy and color appearance. Finally, I compare and contrast the Bayesian approach to other ways of describing visual perception and I discuss the empirical, theoretical and philosophical advantages that accrue from a Bayesian approach.

The puzzle of constancy

As already outlined, the ambiguity in the sensory signal makes unambiguous perception a computational puzzle. However, visual constancy also raises a number of philosophical puzzles. For example, consider naïve realism. This is the idea that humans perceive directly the properties of objects in the world. Naïve realism says that perception of an object's color and size should remain constant across changes in that object's environment. After all, what has changed is the environment, not the object (or its properties). A banana seen on a green banana plant outdoors should appear to have the same color as a banana under the fluorescent light of a grocery store. A newspaper under partial shadow should appear to have uniform color. If naïve realism were correct—or, in other words, if it were true that humans had complete visual constancy because they directly perceive object properties themselves—visual constancy would be a computational puzzle but not a philosophical one.

However, a first philosophical puzzle arises from noting that humans demonstrably do not achieve complete visual constancy. In the case of color, empirical studies suggest that color constancy in relatively realistic environments is around 60% (see Shevell and

[1] As an example, consider the "guessing rule" of using local contrast rather than luminance to encode surface lightness of an object. Then consider two changes to the visual environment. In A, the illumination changes; for example, the intensity of the light source in the room is increased. In B, the illumination stays the same, but the object is placed on a darker surface, one with lower reflectance. Contrast coding would support lightness constancy in A, since the luminance ratio between the object and its surround has not changed. However, in B, application of contrast coding would make appearance inconstant; the object would appear lighter because the luminance ratio between the object and its surround has increased.

Kingdom 2008 for review). Furthermore, the degree to which humans (and animals) exhibit color constancy is highly dependent on the environment. Variables that affect the degree of color constancy between environments include the complexity of the environment (Fleming et al. 2003; Kraft et al. 2002; Schirillo 1999), scene geometry (Allred and Brainard 2009; Boyaci et al. 2003; Schirillo et al. 1990), material properties such as gloss (Xiao and Brainard 2008), and whether the environmental changes are in the illumination or in the surrounding surfaces (Allred and Brainard 2009; Arend and Spehar 1993a, b; Kraft and Brainard 1999; Delahunt and Brainard 2004). Why does the visual system have incomplete constancy? And why is it so variable?

The puzzle of incomplete constancy is exacerbated by the fact that there are conditions under which individuals can make multiple perceptual judgments about the same stimulus. The possibility of multiple judgments is clear through introspection, and is verified by empirical studies (Arend and Spehar 1993b; Blakeslee et al. 2008). For example, when one looks at a surface like a book under a sharp shadow, there is a sense in which it is clear that the cover of the book is all the same color, but there is also a sense in which the part of the book under shadow seems darker. Thus, multiple judgments are made: the color of the book looks to be the same and also appears to vary. A silver dollar seen at an angle can be judged both as elliptical and as circular. A large object (like a car) seen at a distance is experienced both as being large and as being smaller than it is when it is viewed up close.

That multiple judgments about the same stimulus can be made under some conditions is clear, both from introspection and from empirical studies. Less clear is the characterization of when multiple judgments can be made—are there only certain circumstances when this can happen, or is it always possible? Consider again the book under shadow. Are there some shadows under which the book appears only one color? Are there some distances or objects for which the object looks exactly the same despite changes in its distance or orientation? The literature is murky on this point. Particularly in the cases of color, there are times when it seems only one judgment can be made; that is, even when observers are instructed to make multiple judgments, they can make only one (Arend and Spehar, 1993a; Ripamonti et al. 2004). However, there are also circumstances in which observers can comfortably make multiple judgments (Arend and Spehar 1993b; Blakeslee et al. 2008). Why should it be the case that multiple judgments are sometimes, but not always, possible? Is it possible to predict the circumstances under which multiple judgments are possible?

As outlined from both computational and perceptual perspectives, the puzzles of constancy are threefold. First, given the ambiguity in the sensory signal, how ought the visual system to arrive at an unambiguous percept? Second, is there a principled way to predict when that ambiguity will be resolved veridically? That is, is there a principled way to predict the extent to which humans achieve color constancy in different environments? Finally, is there a principled way to predict the circumstances under which multiple judgments of color appearance can be made; and, when multiple judgments can be made, can one assert whether a given judgment is perceptual or cognitive? In the following sections, I describe an answer to these questions that relies on Bayesian statistical theory. In the end, I argue that using a Bayesian approach facilitates answers to the first two questions by providing a quantitative account of empirical data. I also argue that the success with which Bayesian algorithms resolve sensory ambiguity and predict variable constancy points to several ways one might think about answering the latter question of multiple judgments. To do so, I first describe the general principles of Bayesian algorithms.

General principles of Bayesian algorithms

Any Bayesian algorithm makes an optimal interpretation of some set of data. Here Bayesian algorithms are applied to the problem of color: the optimal interpretation is a judgment about color and the data are some parameterization of a sensory stimulus (e.g., values for the isomerization rates of different photopigments in response to incident light reflected from a colored surface). Importantly, as described here, the Bayesian approach aims to predict human perception rather than veridicality per se. In this sense, the Bayesian approach to color perception is distinct from many other computational approaches. Many computational approaches aim for veridicality (either implicitly or explicitly); that is, they aim to recover accurately the distal object from the sensory signal (Retinex is an example, see Land and McCann 1971). However, as we will see, the Bayesian framework is absent the notion that optimality requires veridicality. As formulated specifically in the context of color perception, we can think of a Bayesian algorithm as guessing what state of the world (the *optimal* interpretation, as explained below) caused the sensory signal that reaches the eyes (the data). Keep in mind, however, that this optimal guess may not be the one that correctly describes the physical properties of the stimulus.

To illustrate the process by which the Bayesian algorithm makes this guess, consider the following simple example. Suppose the world consists of a single object comprised of a single surface that reflects only one wavelength of light. In this scenario, which wavelength is reflected can be ignored, but the percentage of light reflected (surface reflectance) is the relevant characteristic of the state of the world, and one that can vary from surface to surface. The datum from which the algorithm must guess this object's surface reflectance is a single value, the intensity of light that reaches the eye. As illustrated in Figure 11.1, the intensity of light reaching the eye (the datum) is a product of the reflectance of the surface and the light that is available to be reflected (the illuminant). The Bayesian algorithm must take the single datum (the intensity value) and make its best guess about the state of the world (the reflectance of the object and the intensity of the illumination) that caused the datum. In other words, the Bayesian algorithm must take one value (intensity) and guess two values (surface reflectance, illumination). Now suppose, in this single surface world, that illumination changes while the reflectance of the object stays the same. This would change the datum (the intensity of the light reaching the eye). If the algorithm predicts the reflectance to be the same in both illumination contexts, then the algorithm predicts constancy; if, on the other hand, the algorithm predicts different reflectance values in the two illumination contexts, then the algorithm shows a failure of constancy.

How do Bayesian algorithms arrive at an optimal interpretation of the data? First, the algorithm computes the probability of each possible state of the world conditional on the data at hand. In the one-surface world, this is every combination of reflectance and illuminant (state of the world) that could have caused the observed intensity of light (datum). For example, one possible state of the world is a high illuminant and low surface reflectance, another possible state of the world is a low illuminant and a high surface reflectance. This probability distribution—the probability of each possible state of the world given the data—is called the posterior distribution.

This posterior distribution is generated by multiplying two other probability distributions—the *likelihood* and the *prior*. The *likelihood* is the probability that the datum would be observed given each possible state of the world. For example, suppose my datum is 10 units. How likely is a sensory signal of 10 units if the state of the world is a reflectance of 0.1 and an illuminant of 100? The value is calculated using a likelihood function; conceptually, the likelihood function captures both sensory noise and the ambiguity inherent in image formation. In the one-surface world, for example, likelihood values will be high for states of the world that are consistent with the data (e.g., reflectance of 0.1 and illuminant of 100, or reflectance of 1 and illuminant of 10) and low for states of the world that are inconsistent with the data (e.g., a reflectance of 1 and an illuminant of 10,000, or a reflectance of 0.01 and an illuminant of 1). The exact value of the likelihood depends on the sensory noise. If there were no sensory noise (for example, if photoreceptors responded exactly the same every time they were stimulated by a given number of photons), only states of the world perfectly consistent with the datum would have a non-zero likelihood. Furthermore, all consistent states of the world would have equal likelihood. In contrast, large sensory noise permits higher likelihood values for reflectance-illumination combinations that are not consistent with the data. The precise formulation of the likelihood varies with the complexity of the algorithm, but for color perception, likelihood functions might utilize information about phototransduction noise, cone isomerization rates, or cone metamerism. Importantly, there are often multiple states of the world with identical likelihood values; this is another way of saying that the data alone are ambiguous with respect to what object(s) caused them.

The *prior* is the probability that any given state of the world exists to begin with, independent of the data. This part of the Bayesian algorithm captures information we have about the statistics of the world. In our one-surface world, the prior would contain two one-dimensional probability distributions, one for surface reflectance and one for illuminants. If there were more dark surfaces than light surfaces in the world, for example, the prior probability for surfaces with low reflectance would be greater than the prior probability for surfaces with high reflectance. Intuitively, the prior can be thought of as the guessing rules to resolve the sensory ambiguity of the likelihood. Next, the posterior probability—the probability of each possible state of the world, given the data—is computed by multiplying the likelihood of each state of the world by the prior probability of each state of the world.[2]

---

[2] The exact values of the likelihood and prior depend on their specific characterization; this is often a statistical distribution with a known probability density function. In principle, values are normalized

The final step in making an optimal estimation is to convolve the posterior distribution with a loss function, sometimes called a cost function or utility function. This function describes the cost associated with incorrectly estimating any particular state of the world. In the one-surface world, for example, if there were a high cost for guessing that a surface was very white when it was actually closer to black, the cost function would shift the optimal estimate towards a darker surface than the posterior distribution alone would yield. Intuitively, the loss function captures information about task strategy. The convolution of the posterior distribution with the loss function yields the expected Bayes' utility. The optimal interpretation of the data (in the one-surface world, the algorithm's final guess about reflectance and illuminant) is the state of the world with the highest expected Bayes' utility. Many Bayesian models, including all of those presented in this paper, have a simple, uniform loss function, and thus the optimal interpretation is driven entirely by the posterior. However, strategic effects can be modeled using the loss function.

A Bayesian model of color constancy

To illustrate the value of the Bayesian approach, consider the case of color constancy. For the color appearance of an object to be useful, it should be approximately constant; that is, it should correlate with reflectance properties of that object across changes in the environment. Such changes could include the intensity and chromaticity of the incident illumination and the reflectance properties of surrounding objects. As discussed previously, the human visual system is approximately color constant, but it has been difficult to predict the extent of that constancy in different environments. Recently, Brainard and colleagues (2006) proposed a Bayesian model of color constancy that provided a good explanation for the color appearance of surfaces under different illuminants. This model predicted both successes and failures of color constancy. An overview of the model is presented here; for detailed information, see Brainard et al. (2006).

To understand the model, recall that any Bayesian algorithm estimates what state of the world most likely caused a particular set of data. For the color constancy algorithm, the "state of the world" is characterized as a mathematical approximation of the reflectance of a series of surfaces and the chromaticity of one uniform illumination. The data are conceived as the isomerization rates of the three photoreceptor types. The likelihood function then captures the mapping between states of the world (surfaces/illuminant) and the data (isomerization rates). This mapping took the form of modeling the sensory signal as the multiplication of the surface reflectance vectors by the illuminant chromaticity, using a biological model of the absorption spectra of the different photoreceptor classes to estimate average isomerization rates, and then perturbing those isomerization rates with noise (see Brainard et al. 2006 for the mathematics). The likelihood function characterizes the probability with which any combination of reflectance and illuminants would have caused the modeled isomerization rates of the cone classes. Intuitively, this

by a constant. In practice, we are often more interested in knowing relative probabilities than absolute probabilities, and the constant is ignored.

likelihood function captures knowledge of image formation and early sensory processing of the cones.

As in the one-surface world, however, multiple combinations of reflectance and illuminant are equally consistent with the data. To make an optimal estimate or guess, the algorithm uses the prior; that is, the probability with which any particular reflectance or illuminant is likely to occur in the first place. To create the prior for surfaces, Brainard et al. (2006) first assumed that surface reflectance functions are well approximated by a three-dimensional linear model. This assumption is frequently used in rendering applications (e.g., 3D computer graphics in video games or movies). The basis functions for the linear model were obtained via statistical methods from empirically measured surface reflectance functions. The same process was repeated to characterize illuminant spectral power functions. The average chromaticity of the illuminant prior was set as the CIE illuminant D65.[3] The remaining parameters determining the relative probability of other chromaticities compared to the mean were estimated from a set of measured daylight spectra. There are several steps of mathematical simplification, but these priors capture the intuition that the visual system expects a distribution of illuminants that occurs under various daylight conditions, and it expects surface reflectances that are likely to occur in the natural world.

To get the posterior distribution, the prior is multiplied by the likelihood, and the estimate of scene illuminant and reflectance of the surfaces within the scene is the maximum of the posterior distribution. Thus, for any arbitrary scene, the algorithm can estimate the illuminant and reflectance of surfaces within the scene.

The computational challenges involved in arriving at an unambiguous prediction of reflectance and illuminant chromaticity are non-trivial. However, the fact that an algorithm can pick one interpretation out of an infinite number is not sufficient to recommend the algorithm; the chosen interpretation should be sensible. To evaluate the success of the algorithm, it is important to compare its predictions to human judgments of color appearance. To do so, Brainard et al. (2006) compared model predictions to data collected by Delahunt and Brainard (2004). In that experiment, observers successively viewed seventeen simulated scenes (see Fig. 11.2) and adjusted a test spot within the scene until it appeared achromatic (that is, lacking in color, except for white, gray, or black). The background reflectance and the illuminant were both varied between scenes. In general, Delahunt and Brainard (2004) reported much higher degrees of color constancy when the illuminant alone was varied than when both the illuminant and the background reflectance were varied.

Any comparison of algorithm estimates to human perception involves an explicit linking hypothesis; that is, it requires a way to relate algorithm estimates of scene parameters

[3] This is considered a "standard daylight illuminant"; purchasers of light bulbs aimed at full-spectrum or daylight illumination may see this notation on their packaging. Please note: in the text "CIE" stands for Commission Internationale de l'Éclairage.

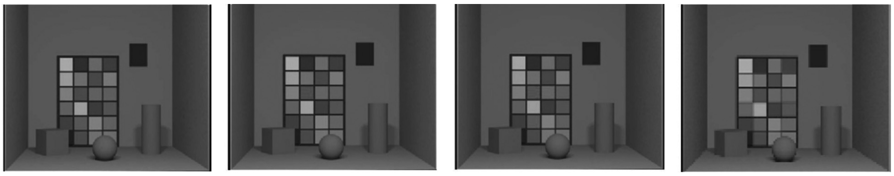

Fig. 11.2 Examples of four stimulus configurations viewed by human observers in Delahunt and Brainard (2004). Subjects adjusted a test patch at the location indicated by the black square until it appeared achromatic. (Reproduced from Brainard, ~~David H.,~~ David R. Williams, and Heidi Hofer, Trichromatic reconstruction from the interleaved cone mosaic: Bayesian model and the color appearance of small spots, *Journal of Vision* 8(5): pp. 11–23 © 2008 Association for Research in Vision and Ophthalmology, with permission.) (See also Plate 11.)

to human estimates of color perception. Here I discuss model estimates of illuminant reflectance and illuminant chromaticity estimated from human achromatic settings.[4]

Model estimates of illuminant intensity tended to agree well with those inferred from human observers' achromatic points. Interestingly, the model's failures to predict veridical illuminant intensity were often accompanied by failures of human color constancy. There are an infinite number of combinations of surface reflectance and illuminants that could give rise to the sensory signal as modeled in the paper. On average, the algorithm predicted illuminant chromaticity correctly in the same contexts that elicited veridical judgments of illuminant chromaticity from human observers.

One of the puzzles outlined above is the variability of color constancy in different environments. Here the Bayesian algorithm predicts that variability. Interestingly, the way it does so provides some insight. The algorithm predicts failures of constancy in contexts where the illumination profile is relatively improbable according to the prior distribution. One experimental example is an environment with purple illumination. According to the prior, purple illumination is very improbable. When a "purple" sensory signal is

[4] The measure of constancy in these experiments is complicated and illustrates the philosophical issues at hand. Constancy implies that an achromatic surface will continue to appear achromatic under different illuminations. Constancy indices typically reflect the degree to which experimental settings conform to that actual achromatic surface. However, there is no principled reason to stipulate the illumination under which achromatic settings reflect "real" achromatic reflectance properties. To increase the chances that any reported failures of constancy were not simply failures to characterize the "real" achromatic surface, Brainard et al. used numerical search to infer the reflectance properties of an achromatic surface that minimized the constancy errors across all 17 contexts simultaneously. Although observers actually adjusted knobs until the light reaching their eyes appeared achromatic, this process of inferring the reflectance properties of an achromatic surface allows characterization of the human observers' estimates of the illuminant chromaticity. The algorithm estimates illuminant chromaticity. Achromatic settings in different contexts can be predicted from these estimated illuminant chromaticities. The differences between actual and predicted achromatic chromaticities are the data plotted in Brainard et al. (2006). This analysis is complicated. For simplicity and because estimates of illuminant chromaticity are at the heart of both model and human calculations, I refer to model and human comparisons of illuminant chromaticity rather than predicted and actual achromatic characteristics.

encountered, therefore, odds are that the "purpleness" arose from a purple surface rather than a purple light. The prior quantitatively implements that probability. Because of this, the Bayesian algorithm predicts relatively low constancy for purple illumination; that is, it predicts that observers will interpret surfaces under a purple light as being more purple than they really are. This prediction agrees with observer estimates. Thus, the Bayesian algorithm suggests that observers will exhibit low degrees of constancy when the actual state of the world is an improbable one.

To summarize, the Bayesian model of color constancy does four things. First, the model solves the computational puzzle of arriving at one estimate of color from the infinite number of possibilities presented by the sensory signal. Second, the solutions at which the model arrives predict human perception well. Third, the model predicts the variability in human color constancy. Finally, the model provides a principled reason for that variability; that is, constancy is worse when the visual system encounters relatively improbable events.

A Bayesian model of color appearance

There are a number of unresolved issues surrounding the mapping between the color appearance of stimuli and the physiological responses early in the visual system to those stimuli. One example is the way in which the retinal photoreceptor tiling in an individual is related to that individual's phenomenal color experience. For example, it is an empirical observation that the ratio of L to M cones varies substantially between subjects, as does the pattern in which they are arranged (Hofer et al. 2005a, b). The ratio also varies by location in the retina even with individual subjects. The comparison of different cone classes is necessary for color vision. If color appearance at a particular retinal location is dependent on the comparison of L and M cone responses in that location, then this inter- and intra-subject variability raises the perplexing question of how color appearance is relatively stable across retinal location and between observers (see MacLeod 2010 for discussion).

How should the variability in photoreceptor mosaic affect the color appearance in an ideal observer? Brainard et al. (2008) took advantage of recent advances in the experimental technique of adaptive optics to address this question. Adaptive optics has made it possible to present very small spots of light to the retina of awake, behaving human observers. These spots are so small and well characterized that there is a high probability that a single stimulus falls on a single cone. Further, extensive imaging of individual human retinas has made it possible to estimate the photoreceptor tiling in individual observers. These techniques are complicated and explaining them is beyond the scope of this chapter (see Hofer et al. 2005a for discussion). Here the results are summarized.

Hofer et al. (2005a) presented monochromatic stimuli to observers after imaging their photoreceptor mosaics. Examples of differences between photoreceptor mosaics are shown in Figure 11.3, where the variation is qualitatively obvious. For example, observer HS had primarily M-cones (colored green), while observer BS had primarily L-cones (colored red).

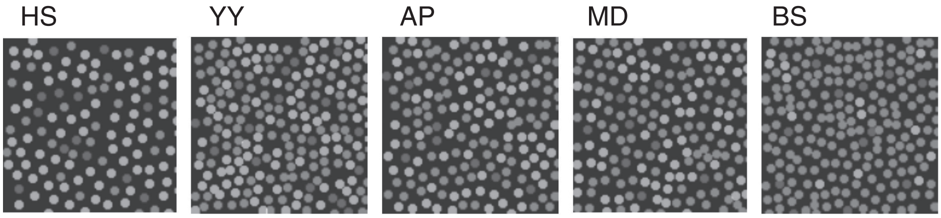

Fig. 11.3 Illustration of a small part of the photoreceptor mosaic for five subjects. Retinas were imaged using adaptive optics techniques described in Hofer et al. (2005a). (Reproduced from Brainard, ~~David H.,~~ David R. Williams, and Heidi Hofer, Trichromatic reconstruction from the interleaved cone mosaic: Bayesian model and the color appearance of small spots, *Journal of Vision* 8(5): pp. 11–23 © 2008 Association for Research in Vision and Ophthalmology, with permission.) (See also Plate 12.)

After the flashes were presented, observers reported the apparent color of any flashes that were sensed. Not all flashes were detected. Hofer et al. (2005a) observed both intra-observer and inter-observer variability in reported color. For example, the retinal location of the presented stimulus affected the response of individual observers, such that the same stimulus could be named a different color depending on where in the retina it was presented. Figure 11.4 (left panel) shows the variability in color names for a 550-nm light. Monochromatic 550-nm light presented under normal conditions is named green, yet under conditions that stimulated a single cone, observers gave up to nine different color names to the same physical stimulus. Moreover, the relative frequency of different color names varied between observers. An example of this variability is seen in Figure 11.4, where observer YY used blue much more frequently than did observer BS, though they both "saw" the same 550-nm stimulus.

What can account for this pattern of results? To further understand the complexity of the pattern, first consider that although the portions of the retina imaged contained some S-cones, 550-nm light does not elicit responses from S-cones. Thus, it is likely that at most two different physiological responses (either M- or L-cone) elicited the full range of color names. Second, a salient feature of the data is that a number of stimuli were reported as achromatic (see Fig. 11.4, top panel), although every stimulus was monochromatic and stimulated only one cone. The traditional view in color science is that stimuli appear achromatic when all three cone classes are stimulated equally. Taken together, as Hofer et al. noted (2005a), these results show that stimulating a single cone type resulted in phenomenally different experiences. This provides a challenge for theories of color appearance that equate stimulation of a particular cone class with an elemental sensation.

Comparing the individual photoreceptor mosaics and the color names (Fig. 11.3 and Fig. 11.4) does not present an immediately obvious explanation for the results. To account for the results, Brainard et al. (2008) formulated the color naming task as the output of a Bayesian model that estimated the most likely color from a noisy photoreceptor array.

To understand this Bayesian algorithm, consider again the process of computing the optimal interpretation (chromaticity of the stimulus, mapped to a color name) from a set of data (in this case, the expected mean isomerization rates of the cones in each individual's retina after accounting for the expected blur due to the optics of the eye). The known features of the retina and physiological properties of different photoreceptors were used to model the likelihood. As with the color constancy model, the key feature of the model is the prior. The prior here incorporates several observations about natural images. First, the intensity in images tends to vary slowly over space. Second, the isomerization rates of L-, M-, and S-cones at a given location tend to be highly correlated. This occurs both because of the properties of images and the properties of the cones themselves. The math required to implement these ideas is complicated, but the essence is that the prior incorporates intuitions about correlations across space and color.

The Bayesian algorithm then estimated the most likely ideal image given the photoreceptor responses. To convert the algorithm output to a form that could be directly compared with psychophysical observations, Brainard et al. (2008) used a standard map of chromaticity to color names. The model estimates are shown in Figure 11.4 (right panel). The model predicted the variability between subjects both in terms of the colors named and the proportion of colors called achromatic.

In this case, a new experimental technique (adaptive optics) revealed situations in which our intuitions about color appearance are incorrect; the appearance of small spots is *not* constant across retinal location within an observer, and the appearance of small spots is *not* constant across observers. The Bayesian model that estimates the most likely ideal image from expected isomerization rates of cones is in agreement with the empirical data.

At first glance, it may seem surprising that intuitions about color appearance are so wrong; why do we believe color appearance to be stable across retinal position if it is not?

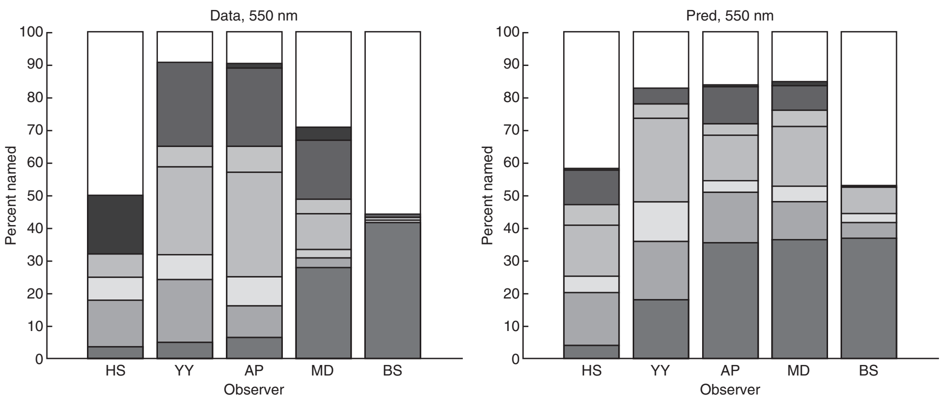

Fig. 11.4 Actual (left) and predicted (right) color names for presentation of 550-nm spots. The observers' task was to report the color of each spot they judged nameable. The histogram shows the proportion of total presentations named the illustrated color. (Reproduced from Brainard, ~~David H.,~~ David R. Williams, and Heidi Hofer, Trichromatic reconstruction from the interleaved cone mosaic: Bayesian model and the color appearance of small spots, *Journal of Vision* 8(5): pp. 11–23 © 2008 Association for Research in Vision and Ophthalmology, with permission.) (See also Plate 13.)

The answer lies again with probability. Spots of light this small are very unlikely. They are so unlikely, in fact, that they would never have been encountered over the course of our evolutionary history. Although the priors in this algorithm were designed using the statistics of common natural images, here they are used to predict the appearance of very unlikely stimuli. A good test of a model is whether it can generalize to predict the appearance of novel stimuli; in this case, Brainard et al. (2008) simulated the cone responses to more typical, larger stimuli. The same model that predicted variability in color appearance to "unlikely," highly localized, monochromatic stimuli revealed stability of color appearance across less localized stimuli. For example, the ideal images for simulated isomerization rates of cones to larger spots of uniform color, or to low-frequency sinusoidal color gratings (wide stripes of color) were stable across subjects and close to veridical.

To summarize, the Bayesian model of color appearance of very small spots also addresses the puzzles described in the introduction. First, it solves the computational problem of arriving at a single prediction of color from an ambiguous sensory signal, in this case the isomerization of a single cone. Second, the color names chosen as optimal by the model are in agreement with the color names utilized by human observers. Third, this algorithm also predicts the inter- and intra-observer variability in color names. Fourth, as with the previous model, the way in which the model predicts this variability provides useful information. In order to predict the variability, the model assumes that the visual system is attempting to guess what stimulus was in the world. There is always ambiguity in the sensory stimulus. The model assumes that the visual system applies the same priors (guessing rules) to judging the appearance of very small spots as it would apply to judging the appearance of larger, more typical stimuli. In other words, the prior says that highly localized stimulation of a single cone is very, very unlikely. When it receives sensory input consistent with a single cone being stimulated, the visual system will guess that parts of the sensory signal were due to noise. The visual system will guess (because of the priors) that the stimulus presented was larger and more chromatically correlated than it actually was. An important test is that the priors that predict the appearance of very small spots also predict the veridical appearance of more typical stimuli. Thus, the priors that predict the appearance of small spots tell us something about the guessing rules that humans employ under more normal viewing conditions; they provide a principled explanation for the variability in human color perception.

Discussion

There are multiple ways to approach the problem of predicting visual perception from the sensory signal. The Bayesian approach described here is an example of the computational approach as characterized most famously by Marr (1982). Other common approaches are the mechanistic approach and one I will term the piecework approach. In the mechanistic approach, parts of the model are thought of as abstractions of pieces in a real visual system. One understands color vision (or visual perception) by understanding the physiology that underlies it in the mammalian system. Typically, the foundation of mechanistic models is the responses of neurons in early visual areas (such as the retina, lateral geniculate nucleus (LGN), and primary visual cortex) to simple, well-controlled visual stimuli whose dimensions can be well characterized. Mathematical approximations of typical neural processes are used to guide model development (at least conceptually; for an example, see ODOG—oriented difference of Gaussian filters—in Blakeslee and McCourt 1999). This approach has achieved remarkable success in predicting perception for certain types of scenes (e.g., with uniform illumination). However, it is unclear as yet whether it will generalize to more complicated scenes where the image data are less easily parameterized with respect to neural responses. One can think of this type of approach as a ground-up model that takes as input an idealized retinal image, and uses the early physiological treatment of the sensory signal as the "guessing rules" to solve the ambiguity in image formation.

A second approach, which I have termed piecework, is to first segment a scene (or an idealized retinal image) using a higher level representation of the image or heuristics and then apply segment-specific rules to solving sensory ambiguity. A scene could be divided into a region of low illumination or high illumination, or segregated into regions representing different depths. An example of this is anchoring theory (Gilchrist et al. 1999). The Gestalt principles of perceptual organization can be considered part of this tradition. Theories like this can explain a number of perceptual phenomena (Gilchrist et al. 1999). One drawback to this type of approach is that it is usually post hoc, and it is not clear how a particular approach should generalize to novel images. As Weiss et al. (2002, 603) state with reference to the motion stimuli, "Because these rules are not formulated directly on image measurements, it is not clear how one should generalize them for application to arbitrary spatiotemporal stimuli."

A third approach is computational. The idea here is to consider perception as an information processing task and to consider the optimal way to solve that task. This approach is agnostic as to physiological implementation. The particular computational approach described here is Bayesian, although some of the advantages described here accrue to any well-conceived computational approach. One advantage of most computational models is that they provide testable predictions. For example, most computational models take as input some set of numbers characterizing a visual image, and yield as output another set of numbers describing the estimated characteristics of that scene. In theory, at least, these estimates can then be compared quantitatively to human perception. Unlike more qualitative theories of visual perception, this allows for interpretable evaluation of models.

In addition to the more general advantage of providing testable predictions, two more advantages accrue to a specifically Bayesian computational approach: (1) compared to other models, Bayesian models tend to be more easily generalizable; (2) Bayesian models require an explicit mathematical representation of the assumptions that go into every modeling approach. The "guessing rules" that must underlie every modeling approach (even if they remain unstated) are formulated explicitly as the priors.

In discussing the Bayesian approach, it is important to recognize that it is not a refutation of the ideas underlying either the mechanistic approach or the piecework approach. In fact, many likelihood functions are based on the same building blocks that go into mechanistic models. With regard to the piecework approach, one can think of the priors as a way of mathematically formulating heuristics or as an alternative description of Gestalt principles. All three approaches are complementary.

The Bayesian modeling approach is used in many different fields. In the end, Bayesian statistical theory provides an optimal estimate given a set of assumptions. Those assumptions are mathematically formulated as the likelihood function, the prior, and the loss function. In the context of Bayesian modeling of color, the priors have been interpreted here as a formalization of the guessing rules used to resolve sensory ambiguity. The priors used in these models capture something interesting about the statistics of the natural world. The priors, thus framed, are meant to characterize the salient features of the world in which the human visual system evolved. However, any of these Bayesian models as constituted could be generalized to include a prior that is less evolutionary in nature. For example, priors could be modified to include such variables as the prevalence of a particular stimulus within an experimental session. In the models presented here, the loss function was very simple and never invoked explicitly. However, the loss function could be modified to include the experimental cost or benefit of making different kinds of errors within an experiment. Manipulating the expected frequency of a stimulus and the reward associated with different patterns of responses have been shown to change behavioral responses (Hudson et al. 2007; Knill 2007). Bayesian modeling that incorporates such stimulus changes into the prior and such reward contingencies into the loss function has successfully explained behavioral responses to certain types of visual stimuli (Knill 2007), although experiment by experiment priors and loss functions have not yet been applied to color perception.

The computational approach described here differs from the mechanistic approach in one fundamental way. Although the Bayesian approach can functionally solve the problem of sensory ambiguity, it does not speak directly to the question of neural implementation or processing. Should a Bayesian model perfectly describe perception, it would not constitute evidence that the brain has circuits devoted to multiplying probability distributions, or that quantities such as prior probabilities are represented explicitly at some neural locus (although some research has suggested this, see Ma et al. 2006). Of course, the models described here do incorporate knowledge of the human visual system in the likelihood. However, the Bayesian approach is best thought of as a way to articulate cogently the relationship between assumptions and results, and as a good way to frame a complicated problem.

Although the algorithms presented provide estimates of color appearance that are in broad agreement with human data, the agreement is not perfect. There are several possibilities for the discrepancy. To compute likelihood and prior probabilities in a timely fashion, convenient forms of probability distributions are sometimes used. In each model, choices are made for how to parameterize scene variables. For example, in the constancy model, the data are weights on basis functions that describe reflectance and illuminant. In the color appearance model, the data are expected isomerization rates of cones. Each of these parameterizations requires assumptions and these affect the results. However, in the case of the likelihood, the data generation process and scene parameters are often well understood; for example, isomerization rates of cones have been measured, and the optics of the eye are well characterized (Wandell 1995; Wyszecki and Stiles 1982). Given that the likelihood function is usually well constrained, it is often the choice of the priors that determines the model estimates; thus, lack of agreement between model estimates and human data may lie with the choice of priors. In each model presented here, the priors are highly simplified characterizations of the natural world. In many cases, they contain known errors that are tolerated for computational reasons. In fact, given the simplicity of the priors and the statistical complexity of images, the degree to which the models do explain human perception is somewhat surprising.

However, despite the simplifications and limitations described above, the Bayesian approach to characterizing visual perception and constancy is appealing on broad empirical, theoretical, and philosophical grounds.

Empirically, Bayesian estimates are in good agreement with human visual perception in many different domains. Here I have described in detail results from two models. Bayesian models can also account for the speed and direction of motion perception (Stocker and Simoncelli 2006), perceptual phenomena associated with vision at low-light levels (Manning and Brainard 2009), shape perception (Knill 2007), perceptual bistability (Van Ee et al. 2003), and Bayesian approaches are being applied to a host of perceptual phenomena (see Kersten and Yuille 2003; Knill and Richards 1996). Moreover, Bayesian models are typically more parsimonious than other models, requiring fewer parameters to describe human perception under a wide range of conditions.

Theoretically, Bayesian models are easily generalizable. To make predictions about more complex stimuli, it is only necessary to add more distributions to the likelihood and probability functions, rather than engaging in *de novo* model construction. This modularity is useful. Estimates of object properties can be compared quantitatively to human performance and the effect of each intuition instantiated in the prior can be measured. For example, in the Bayesian model of color constancy (Brainard et al. 2006), the initial model that drew on the measured daylight prior exhibited some notable failures to predict human performance. Subsequently, Brainard et al. (2006) derived the prior that provided the best possible fit to the data. This prior was considerably broader than the daylight prior that was originally used. Given that the likelihood is well constrained, and the functional form of the prior is defined, the strategy of determining what prior parameters best account for the data can provide insights into visual system processing (Stocker and Simoncelli 2006). The empirical success of Bayesian models does not require that the brain represent probabilities in the same way that the models do. However, if the priors that work best in different models all seem to share a critical feature (viz., representation of statistics of natural images), this may suggest that those statistics are relevant to the neural substrates underlying perception.

Finally, although Bayesian algorithms are themselves agnostic to philosophical questions, an interested party might make some observations about the nature of visual perception from the assumptions that are required for Bayesian algorithms to make estimates of color appearance that agree with human estimates of color appearance. Specifically, it is interesting to note how the Bayesian algorithms presented here predict both successes and failures of constancy (Brainard et al. 2006; Stocker and Simoncelli 2006; Weiss et al. 2002). Given that the likelihood is well constrained, it is the parameters of the priors that control model success. In multiple cases, perceptual phenomena that have seemed to indicate flawed or suboptimal perception have been recast as the output of a rational perceptual system. The model that predicts color constancy also predicts failures of constancy when they occur (Brainard et al. 2006). The model that predicts the seemingly incomprehensible pattern of responses to very small spots of light also explains perception of more typical color stimuli. A model that explains veridical speed perception in high-light conditions also predicts various motion illusions under low-light conditions (Weiss et al. 2002). The critical importance of the priors in predicting human perception in a wide range of contexts suggests three interrelated points.

First, when some property of a prior shows itself to be empirically useful in very different circumstances, the philosophical implications of that property are worth considering. Specifically, because the required priors in each of the presented models have the common feature of capturing the statistics of natural images, a philosophical account of visual perception should take seriously the notion that the visual system is "designed" for veridical perception in environments considered important from an evolutionary point of view. An important goal of philosophical and psychological accounts of perception has been to find correspondence between parameters of physical objects (e.g., reflectance, spatial extent) and parameters of the relevant perceptual space (e.g., color, size). Illusions and failures of constancy are seen as obstacles to discovering this correspondence. Color has been seen as particularly puzzling because the relationship between physical parameters is not related in an obvious way to perceptual dimensions. Taking priors, likelihood functions, and cost functions seriously means rethinking the notion that there exists a direct correspondence between some parameters that span a physical space (e.g., reflectance of an object) and some parameters that span a perceptual space (e.g., perceived color). Under this view, the visual system does not care about getting things right, per se. It cares about getting things right enough to survive. Veridicality is important in the environments in which veridicality would have been essential for an organism's survival. Thus, the correspondence between physical properties of objects in the world and perception of those properties ought to be mediated both by the probability with which the properties occurred and the costs of getting those properties wrong.

It may seem troubling to abandon the notion of direct correspondence between the properties of objects and the properties of perception; after all, if object properties do not uniquely determine perceptual properties, and perceptual properties of a given object vary from scene to scene and from observer to observer, then deriving a principled account of perceptual properties may seem impossible. The appeal of the Bayesian approach is that it provides an empirical account of perception in a wide variety of circumstances that does not rely on hand-waving. Failures of constancy and visual illusions can be given a unified explanation and do not require an ecumenical approach to perception.

Second, if one accepts the Bayesian formulation that the optimal interpretation of a set of data (perception) is shaped by prior probabilities as well as the data themselves, then the set of properties it is possible to perceive veridically is constrained. There are states of the world that a Bayesian algorithm would predict will never be perceived veridically because they are considered too unlikely. For example, in the model of color constancy, some of the variability in constancy is explained by the degree to which the actual state of the world is an improbable state of the world (according to the prior). In the Bayesian formulation, because purple illumination is very unlikely, surfaces illuminated by purple light will be perceived as more purple than they actually are.

A third related point is that the Bayesian framework provides an empirically parsimonious and philosophically appealing account of both visual illusions (Geisler and Kersten 2002) and variability in constancy (Brainard et al. 2006). One model containing one set of priors can predict both illusions and veridical perception; there is no need to think of illusions as errors or puzzles in a philosophical sense. Others have noted the importance of errors in understanding the human visual system (Brainard et al. 2006; Geisler and Kersten 2002; Gilchrist 2006). As Geisler and Kersten (2002, 508) noted when describing one Bayesian model, "These perceptual errors seem to reveal a rational (but automatic) perceptual system designed to correctly interpret the retinal images evoked by the world." Thus, visual illusions and variability in constancy can be thought of as the output of a rational visual system rather than as epiphenomena of some strange kind.

The Bayesian approach suggests one possible way to frame the discussion of multiple judgments, although there is not yet empirical data to support this idea. The Bayesian approach suggests that to make an estimate of what object is in the world given the sensory signal, we combine the likelihood, the prior, and the cost function to determine what object was most probable. But what if the expected Bayes' utility is multimodal; that is, what if two objects are both equally likely (or relatively so)?

Perhaps multiple judgments are possible only when the expected Bayes' utility is multimodal. This multimodality could arise for at least two reasons: (1) The utility/cost function could push you in one direction while the posterior pushed you in another. This is akin to saying that the visual system estimates that a surface is white, but an experimenter has said you will win a million dollars if you correctly identify a dark surface, and you will lose 10 cents if you misidentify a white surface. (2) The posterior itself could be multimodal. This could happen if the likelihood pushes estimates in one direction while the prior pushes estimates in another direction. An example of this could be when a purple surface is presented under purple illumination. Depending on the formulation, the likelihood could be high for states of the world that include purple illumination, but the prior could be very low for those states. Thus, the posterior could have high probabilities for both purple illumination and also for daylight illumination.

I speculate that if the expected Bayes' utility is not multimodal, then only one judgment is possible. Furthermore, I speculate that when expected Bayes' utility is multimodal, two judgments are possible, and the different causes of multimodality implicate different kinds of judgments. When the multimodality occurs in the posterior, then both judgments are perceptual or apparent. That is, the book is experienced as being both white and not-white. Alternatively, when multimodality arises from differential pulls between the cost function and the posterior, I propose that one judgment is perceptual (the left half of the book is darker) and the other is cognitive (but I know it must be painted the same as the right half of the book).

These ideas are entirely speculative; designing a Bayesian algorithm that is simple enough to generate this kind of multimodality but rich enough to suggest interesting psychophysical experiments is not a trivial endeavor, but one that could provide worthwhile empirical predictions.

In summary, any account of visual perception must solve the computational and philosophical puzzles of sensory ambiguity. There are multiple ways to do so. Bayesian approaches, while not perfect, provide an intuitive, generalizable, and philosophically interesting solution to the problems of sensory ambiguity.

References

Allred, Sarah R., and David H. Brainard (2009). Contrast, constancy, and measurements of perceived lightness under parametric manipulation of surface slant and surface reflectance. *Journal of the Optical Society of America A, Optics, Image Science, and Vision* 26(4): 949–61.

Arend, Lawrence E., and Branka Spehar (1993a). Lightness, brightness, and brightness contrast: 1. Illuminance variation. *Perception and Psychophysics* 54(4): 446–56.

———(1993b). Lightness, brightness, and brightness contrast: 2. Reflectance variation. *Perception and Psychophysics* 54: 457–68.

Attneave, Fred (1954). Some informational aspects of visual perception. *Psychological Review* 61(3): 183–93.

Blakeslee, Barbara, Daniel Reetz, and Mark E. McCourt (2008). Coming to terms with lightness and brightness: Effects of stimulus configuration and instructions on brightness and lightness judgments. *Journal of Vision* 8(11): 3.1–14.

Blakeslee, Barbara, and Mark E. McCourt (1999). A multiscale spatial filtering account of the White effect, simultaneous brightness contrast and grating induction. *Vision Research* 39(26): 4361–77.

Boyaci, Hussein, Lawrence T. Maloney, and Sarah Hersh (2003). The effect of perceived surface orientation on perceived surface albedo in binocularly viewed scenes. *Journal of Vision* 3(8): 541–53.

Brainard, David H., David R. Williams, and Heidi Hofer (2008). Trichromatic reconstruction from the interleaved cone mosaic: Bayesian model and the color appearance of small spots. *Journal of Vision* 8(5): 11–23.

Brainard, David H., Philippe Longere, Peter B. Delahunt, William T. Freeman, James M. Kraft, and Bei Xiao (2006). Bayesian model of human color constancy. *Journal of Vision* 6(11): 1267–81.

Delahunt, Peter B., and David H. Brainard (2004). Does human color constancy incorporate the statistical regularity of natural daylight? *Journal of Vision* 4(2) 57–81.

Fleming, Roland W., Ron O. Dror, and Edward H. Adelson (2003). Real-world illumination and the perception of surface reflectance properties. *Journal of Vision* 3(5): 347–68.

Geisler, William S., and Kersten Daniel (2002). Illusions, perception and Bayes. *Nature Neuroscience* 5(6): 508–10.

Gilchrist, Alan (2006). *Seeing Black and White*. Oxford: Oxford University Press.

Gilchrist, Alan, Christos Kossyfidis, Frederick Bonato, Tiziano Agostini, Joseph Cataliotti, Xiaojun Li (1999). An anchoring theory of lightness perception. *Psychological Review* 106(4): 795–834.

Helmoltz, Hermann V. (1924). *Treatise on Physiological Optics*, tr. James P. C. Southall. Milwaukee: Optical Society of America. Translation of third German edition, 1910.

Hofer, Heidi, Ben Singer, and David R. Williams (2005a). Different sensations from cones with the same photopigment. *Journal of Vision* 5(5): 444–54.

Hofer, Heidi, Joseph Carroll, Jay Neitz, Maureen Neitz, and David R. Williams (2005b). Organization of the human trichromatic cone mosaic. *Journal of Neuroscience* 25(42): 9669–79.

Hudson, Todd E., Lawrence T. Maloney, and Michael S. Landy (2007). Movement planning with probabilistic target information. *Journal of Neurophysiology* 98(5): 3034–46.

Kersten, Daniel, and Alan Yuille (2003). Bayesian models of object perception. *Current Opinion in Neurobiology* 13(2): 150–8.

Knill, David C. (2007). Learning Bayesian priors for depth perception. *Journal of Vision* 7(8): 13.

Knill, David C., and Whitman Richards (1996). *Perception as Bayesian Inference*. Cambridge: Cambridge University Press.

Kraft, James M., Shannon I. Maloney, and David H. Brainard (2002). Surface-illuminant ambiguity and color constancy: Effects of scene complexity and depth cues. *Perception* 31(2): 247–63.

Kraft, James M., and David H. Brainard (1999). Mechanisms of color constancy under nearly natural viewing. *Proceedings of the National Academy of Science USA* 96(1): 307–12.

Land, Edwin H., and John J. McCann (1971). Lightness and retinex theory. *Journal of the Optical Society of America* 61(1): 1–11.

Ma, Wei J., Jeffrey M. Beck, Peter E. Latham, and Alexandre Pouget (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience* 9(11): 1432–8.

MacLeod, Donald I. A. (2010). Into the neural maze. In Jonathan Cohen and Mohan Matthen (eds.), *Color Ontology and Color Science*, 151–78. Cambridge: MIT Press.

Mamassian, Pascal, Michael S. Landy, and Lawrence T. Maloney (2002). Bayesian modelling of visual perception. In R. P. N. Rao, B. A. Olshausen, and M. S. Lewicki (eds.), *Probabilistic Models of the Brain: Perception and Neural Function*, 13–36. Cambridge: MIT Press.

Manning, Jeremy R., and David H. Brainard (2009). Optimal design of photoreceptor mosaics: Why we do not see color at night. *Visual Neuroscience* 26(1): 5–19.

Marr, David (1982). *Vision*. San Francisco: Freeman.

Ripamonti, Caterina, Marina Bloj, Kiran Mitha, Robin Hauck, Scott Greenwald, and David H. Brainard (2004). Measurements of the effect of surface slant on perceived lightness. *Journal of Vision* 4(9): 747–63.

Schirillo, James A. (1999). Surround articulation. II. Lightness judgments. *Journal of the Optical Society of America A, Optics, Image Science, and Vision* 16(4): 804–11.

Schirillo, James A., Adam Reeves, and Lawrence Arend (1990). Perceived lightness, but not brightness, of achromatic surfaces depends on perceived depth information. *Attention, Perception and Psychophysics* 48(1): 82–90.

Shevell, Steven K., and Fred A. Kingdom (2008). Color in complex scenes. *Annual Review of Psychology* 59: 143–66.

Stocker, Alan A., and Eero P. Simoncelli (2006). Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience* 9(4): 578–85.

van Ee, Raymond, Wendy J. Adams, and Pascal Mamassian (2003). Bayesian modeling of cue interaction: bistability in stereoscopic slant perception. *The Journal of the Optical Society of America A, Optics, Image Science, and Vision* 20(7): 1398–406.

Wandell, Brian A. (1995). *Foundations of Vision*. Sunderland, MA: Sinauer Associates.

Weiss, Yair, Eero P. Simoncelli, and Edward H. Adelson (2002). Motion illusions as optimal percepts. *Nature Neuroscience* 5(6): 598–604.

Wyszecki, Gunter, and W. S. Stiles (1982). *Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd edn. New York: Wiley.

Xiao, Bei, and David H. Brainard (2008). Surface gloss and color perception of 3D objects. *Visual Neuroscience* 25(3): 371–85.

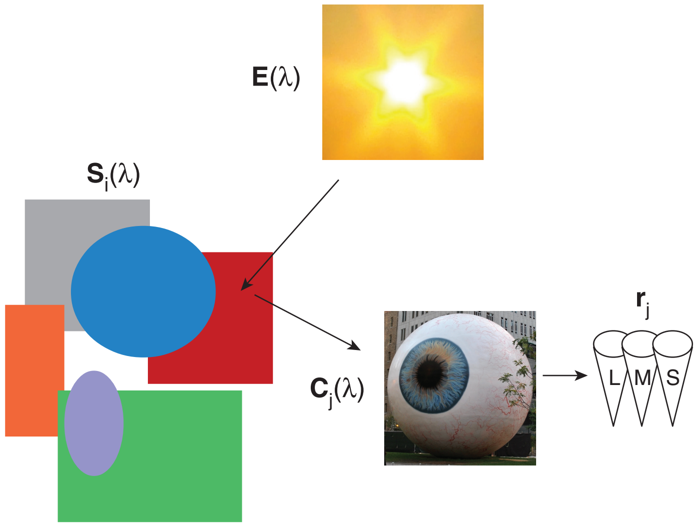

Plate 10 (see also Fig. 11.1) Illustration of sensory ambiguity. (© Maria Olkkonen, 2011, reproduced here with permission.)

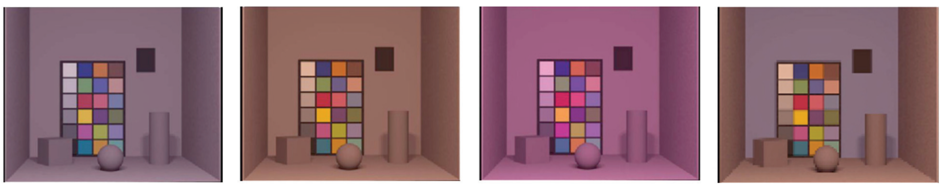

Plate 11 (see also Fig. 11.2) Examples of four stimulus configurations viewed by human observers in Delahunt and Brainard (2004). Subjects adjusted a test patch at the location indicated by the black square until it appeared achromatic. (Reproduced from Brainard, David H., David R. Williams, and Heidi Hofer, Trichromatic reconstruction from the interleaved cone mosaic: Bayesian model and the color appearance of small spots, *Journal of Vision* 8(5): pp. 11–23 © 2008 Association for Research in Vision and Ophthalmology, with permission.)

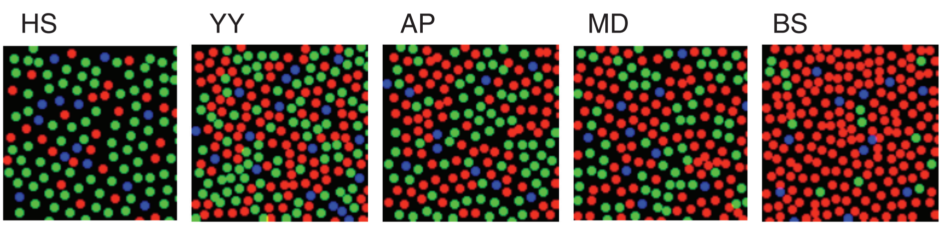

Plate 12 (see also Fig. 11.3) Illustration of a small part of the photoreceptor mosaic for five subjects. Retinas were imaged using adaptive optics techniques described in Hofer et al. (2005a). (Reproduced from Brainard, David H., David R. Williams, and Heidi Hofer, Trichromatic reconstruction from the interleaved cone mosaic: Bayesian model and the color appearance of small spots, *Journal of Vision* 8(5): pp. 11–23 © 2008 Association for Research in Vision and Ophthalmology, with permission.)

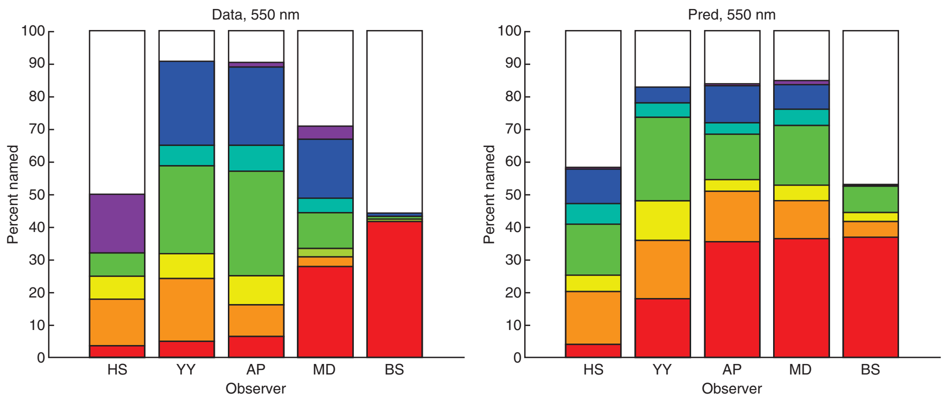

Plate 13 (see also Fig. 11.4) Actual (left) and predicted (right) color names for presentation of 550-nm spots. The observers' task was to report the color of each spot they judged nameable. The histogram shows the proportion of total presentations named the illustrated color. (Reproduced from Brainard, David H., David R. Williams, and Heidi Hofer, Trichromatic reconstruction from the interleaved cone mosaic: Bayesian model and the color appearance of small spots, *Journal of Vision* 8(5): pp. 11–23 © 2008 Association for Research in Vision and Ophthalmology, with permission.)