52 Ideal Observer Analysis

WILSON S. GEISLER

VISUAL SYSTEMS, AND the developmental and learning mechanisms that shape them during the life span, have evolved because they enhance performance in those tasks relevant to survival and reproduction, such as detecting and localizing predators or prey, navigating through the environment, identifying materials, estimating the three-dimensional geometry of the environment, recognizing specific objects encountered before, and so on. Thus, the proper study of a visual system must include an analysis of those specific tasks that the system evolved to perform. An ideal observer analysis provides a principled approach for understanding a visual task, the stimulus information available to perform the task, and the anatomical and physiological constraints that limit performance of the task.

The central concept in ideal observer analysis is the *ideal observer*, a theoretical device that performs a given task in an optimal fashion, given the available information and some specified constraints. This is not to say that ideal observers perform without error, but rather that they perform at the physical limit of what is possible in the situation. In general, ideal observers make mistakes because of the complexity and uncertainty that exist in the visual environment and because of the inherent noise in light or in whatever signal serves as input to the ideal observer. The fundamental role of uncertainty and noise in limiting possible performance implies that ideal observers must be derived and described in probabilistic (statistical) terms.

Ideal observer analysis involves determining the performance of the ideal observer in a given task and then comparing its performance to that of the biological system under consideration, which (depending on the application) might be the organism as a whole, some neural subsystem, or an individual neuron. In vision science, ideal observer analyses have been carried out for many different tasks, ranging from photon detection, to pattern discrimination, to information coding in neural populations, to shape estimation, to recognition of complex objects. Here the focus is on detection, discrimination, and identification, with an emphasis on what has been learned through ideal observer analysis about the retina, lateral geniculate nucleus (LGN), and primary visual cortex. Other applications of ideal observer analysis are described in other chapters within this volume (see also Knill and Richards, 1996; Simoncelli and Olshausen, 2001).

Basic concepts and formulas

The purpose of deriving an ideal observer is to determine the optimal performance in a task, given the physical properties of the environment and stimuli. Organisms generally do not perform optimally, and hence one should not think of an ideal observer as a potentially realistic model of the actual performance of the organism. Rather, the value of an ideal observer is to provide a precise measure of the stimulus information available for performing the task, a computational theory of how to perform the task, and an appropriate benchmark against which to compare the performance of the organism (Green and Swets, 1966). In addition, the ideal observer can serve as a useful starting point for developing realistic models (e.g., Schrater and Kersten, 2001). With an appropriate ideal observer in hand, one knows how the task should be performed. Thus, it becomes possible to explore in a principled way what the organism is doing right and what it is doing wrong. This can be done by degrading the ideal observer in a systematic fashion by including, for example, hypothesized sources of internal noise (Barlow, 1977), inefficiencies in central decision processes (Barlow, 1977; Green and Swets, 1966; Pelli, 1990), or known anatomical or physiological factors that would limit performance (Geisler, 1989).

Obviously, an ideal observer analysis is sensible only for objective tasks with well-defined performance goals, such as identifying as accurately as possible the physical category to which an object belongs or estimating some physical property of an object. Ideal observer analysis is neither possible nor sensible for subjective tasks such as judging the apparent hue of a stimulus or judging whether a stereoscopic display appears fused or diplopic.

BAYESIAN IDEAL OBSERVERS   Most forms of ideal observer analysis are based on the concepts of Bayesian statistical decision theory. To illustrate the Bayesian approach, consider a categorization task where there are n possible stimulus categories, $c_1, c_2 \ldots, c_n$, and the observer's task on each trial is to identify the category correctly, given the particular stimulus $\mathbf{S}$ arriving at the eye.[1] If there is substantial stimu-

---

[1] Experts should note that in this example the utility/loss function is degenerate and does not appear; it will be introduced shortly.

lus noise or overlapping of categories, then the task will be inherently probabilistic. As might be expected intuitively, performance is maximized on average by computing the probability of each category, given the stimulus, and then choosing the category C that is most probable:[2,3]

$$C = \arg \max_{c_i}[p(c_i|\mathbf{S})] \qquad (1)$$

Note that "arg max" is just a shorthand notation for a procedure that finds and then returns the category that has the highest probability, given the stimulus.[4] In practice, the probability of a category given the stimulus is often computed by making use of Bayes' formula:

$$p(c_i|\mathbf{S}) = \frac{p(\mathbf{S}|c_i)p(c_i)}{p(\mathbf{S})} \qquad (2)$$

where $p(c_i|\mathbf{S})$ is the *posterior probability*, $p(\mathbf{S}|c_i)$ is the *likelihood*, and $p(c_i)$ is the *prior probability*.[5] The probability in the denominator, $p(\mathbf{S})$, is a constant that is the same for all the categories and hence plays no role in the optimal decision rule. Furthermore, it is completely determined by the likelihoods and prior probabilities:

$$p(\mathbf{S}) = \sum_{j=1}^{n} p(\mathbf{S}|c_j)p(c_j) \qquad (3)$$

Substituting Bayes' formula into equation 1, the optimal response is given by

$$C = \arg \max_{c_i}[p(\mathbf{S}|c_i)p(c_i)] \qquad (4)$$

In other words, one can identify a stimulus with maximum accuracy by combining the prior probability of the different categories and the likelihood of the stimulus given each of the possible categories.

In the laboratory, maximizing accuracy is a common goal defined by the experimental design. For this goal, all errors are equally costly, because all errors have the same effect on the accuracy measure. However, this is rarely the case for natural situations, where the costs and benefits associated with different stimulus-response outcomes have a more complex structure. For example, if the goal is survival, then some errors are more costly than others—mistaking a

poisonous snake for a branch is more costly than mistaking a branch for a poisonous snake. Within the framework of Bayesian statistical decision theory, more complex goals are represented with a utility function, $u(\mathbf{r}, \boldsymbol{\omega})$, which specifies the cost or benefit associated with making response $\mathbf{r}$ when the state of the environment is $\boldsymbol{\omega}$ (e.g., Berger, 1985). In this more general case, the optimal decision is to make the response, $\mathbf{R}$, that maximizes the average utility over all the possible states of the environment (see footnote 3):

$$\mathbf{R} = \arg \max_{r}\left[\sum_{\omega} u(\mathbf{r},\boldsymbol{\omega})p(\mathbf{S}|\boldsymbol{\omega})p(\boldsymbol{\omega})\right] \qquad (5)$$

In this decision rule, $p(\boldsymbol{\omega})$ is the prior probability of a given state of the environment, and $p(\mathbf{S}|\boldsymbol{\omega})$ is the stimulus likelihood given a state of the environment. Note that equation 4 is special case of equation 5, where the possible states of the environment are the stimulus categories, $c_1, c_2 \ldots, c_n$, the possible responses are the category names, the benefits for all correct responses are equal, and the costs for all incorrect responses are equal.

CONSTRAINED BAYESIAN IDEAL OBSERVERS   The class of Bayesian ideal observers considered so far operates directly on the stimulus $\mathbf{S}$ that arrives at the eye. However, for many applications, it is useful to incorporate some biological constraints into an ideal observer analysis. For example, if good estimates are available for the optics of the eye, the spatial arrangement of the photoreceptors, and their spectral sensitivities, then these estimates can serve as plausible constraints on an ideal observer. In this case, the ideal observer would show the maximum performance possible in the given task, using the photons caught in the photoreceptors. Such an ideal observer must perform worse than one designed to use the photons arriving at the cornea. To the extent that the constraints are accurate, the difference in performance between the ideal observer at the cornea and the one at the level of photon absorptions would provide a precise measure of the information (relevant to the task) lost in the process of image formation and photon capture. Furthermore, the difference in the performance of the ideal observer at the level of photon capture and the performance of the organism as a whole would provide a precise measure of the information lost in the neural processing subsequent to photon capture (e.g., Geisler, 1989).

Another useful way to use constrained ideal observers is to allow some free parameters in the biological constraints and then determine what parameter values produce the best-performing ideal observer. For example, in an ideal observer at the level of photon absorptions, one can allow the peak wavelengths of the receptors to be free parameters and then determine what peak wavelengths would produce the best-performing ideal observer. This would be a precise

---

[2] In the case of ties at the highest probability, one can pick arbitrarily from those tied categories.

[3] Throughout this chapter, capital letters refer to random quantities and boldface letters refer to *vector* quantities, where the term *vector* refers to an ordered list of properties (generally, integer- or real-valued quantities).

[4] More simply, arg max$[f(x)]$ is the value of x (the argument) for which $f(x)$ reaches its maximum value.

[5] Bayes' formula follows directly from the definition of conditional probability: $p(c_i \& \mathbf{S}) = p(c_i|\mathbf{S})p(\mathbf{S}) = p(\mathbf{S}|c_i)p(c_i)$

way of determining how close an organism's photoreceptors are to the optimum for the given task (Regan et al., 2001). Alternatively, the free parameters might represent the receptive field shapes (the configuration of weights placed on each receptor) for some given number of postreceptor neurons. This would be a precise way of determining how close an organism's receptive field shapes are to the optimum for the given task.

A general class of constrained Bayesian ideal observers can be represented by introducing a *constraint function*, $g_\theta(\mathbf{S})$, which maps (either deterministically or probabilistically) the stimulus $\mathbf{S}$ at the cornea into an intermediate signal $\mathbf{Z} = g_\theta(\mathbf{S})$. For example, $\mathbf{S}$ might be a vector representing the number of photons entering the pupil from each pixel on a display screen, and $\mathbf{Z}$ might be a vector representing the number of photons absorbed in each photoreceptor; hence $g_\theta(\mathbf{S})$ would specify the combined effect of the optics, photoreceptor lattice, and photoreceptor absorption spectra. Alternatively, $\mathbf{Z}$ might represent the spike count for each ganglion cell, and $g_\theta(\mathbf{S})$ would specify the combined effect of the optics and all retinal processing. Any free parameters, such as the peaks of the photoreceptor absorption spectra or the shapes of the receptive fields, are represented in the constraint function by a parameter vector $\boldsymbol{\theta}$.

For any given parameter vector, the optimal decision rule has the same structure as before:

$$\mathbf{R} = \arg\max_r \left[\sum_\omega u(\mathbf{r}, \boldsymbol{\omega}) p_\theta(\mathbf{Z}|\boldsymbol{\omega}) p(\boldsymbol{\omega}) \right] \qquad (6)$$

The only difference is that the stimulus $\mathbf{S}$ is replaced by the intermediate signal $\mathbf{Z}$. Applying this optimal decision rule typically requires determining the intermediate-signal likelihood $p_\theta(\mathbf{Z}|\boldsymbol{\omega})$, by combining the constraint function $g_\theta(\mathbf{S})$ with the stimulus likelihood distribution $p(\mathbf{S}|\boldsymbol{\omega})$. If there are free parameters, then the optimal parameter vector is given by the following formula (e.g., Geisler and Diehl, 2002):

$$\boldsymbol{\theta}_{opt} = \arg\max_\theta \left[\sum_Z \max_r \left[\sum_\omega u(\mathbf{r}, \mathbf{w}) p_\theta(\mathbf{Z}|\mathbf{w}) p(\mathbf{w}) \right] \right] \qquad (7)$$

Using $\boldsymbol{\theta}_{opt}$ in equation 6 gives the decision rule for the best-performing ideal observer over the free-parameter space.

The concepts and basic formulas of Bayesian ideal observer analysis are relatively straightforward. However, in specific applications, it can be very difficult to determine or compute the likelihoods, prior probabilities, utility functions, or sums over possible states of the environment. Indeed, there are many situations for which it is not yet possible to determine the performance of the ideal observer. Nonetheless, the number and range of successes have been growing over the years, and the prospects for continued success are good.

Detection, discrimination, and identification

Detection, discrimination, and identification are fundamental visual tasks that have been investigated extensively since the beginning of vision science. In the detection task, the observer is presented with either a background pattern or a background pattern plus a target pattern, and must decide whether or not the background pattern contains the target. The background pattern can range from a simple, uniform field of light to a complex natural scene. Similarly, the target can range from a simple, uniform patch of light to a complex natural object. In the discrimination task, the observer is presented with either a background plus a target or a background plus a modified target, and must decide whether the background contains the modified or unmodified target. Formally, detection and discrimination tasks are equivalent, because the discrimination task can be regarded as a detection task, where the "background" is the background plus the unmodified target and the "target" is the difference between the modified and unmodified targets. In the identification task, the observer is presented with a background plus one of n possible targets and must decide which target is contained in the background. Thus, the discrimination task is a special case of the identification task where the number of possible targets is two.

During the last half century, ideal observer analysis has played an important role in the development of our understanding of the physical, physiological, and cognitive factors that underlie detection, discrimination, and identification performance. Before discussing these applications of ideal observer analysis, I introduce the ideal observer for detection and discrimination tasks, where all sources of information that the ideal observer receives are statistically independent. Many of the results described later are based on this simple kind of ideal observer.

OPTIMAL DISCRIMINATION GIVEN STATISTICALLY INDEPENDENT SOURCES OF INFORMATION   On each trial of a detection or discrimination task, a stimulus, $\mathbf{S}$, from one of the two categories is received by the ideal observer, and it must pick a category. From equation 4, we see that the optimal decision rule is to compute the likelihood ratio, $p(\mathbf{S}|c_2)/p(\mathbf{S}|c_1)$, and compare this ratio to a criterion, $p(c_1)/p(c_2)$, which is the ratio of the prior probabilities of the two categories. If the likelihood ratio exceeds this criterion, then the ideal observer picks c_2; otherwise, it picks c_1. For present purposes, suppose that the prior probabilities are equal (criterion = 1.0).

Consider a situation where the stimulus consists of a set of stimulus components or information sources, $\mathbf{S} = \langle S_1, \ldots, S_n \rangle$. For example, S_i might represent the number of photons entering the pupil from the ith pixel on a video monitor or the number of spikes generated by the ith ganglion cell in the retina. If all the components are statistically

independent, then the probability of the whole stimulus is the product of the probabilities of the individual components. In this case, the performance of the ideal observer can be determined by considering the performance of the ideal observer separately for each component. Assuming approximate normality, an ideal observer that uses only the ith component will perform with an accuracy (percent correct) of

$$PC_i = \Phi\left(\frac{d_i'}{2}\right) \tag{8}$$

where

$$d_i' = \frac{|E(S_i|c_2) - E(S_i|c_2)|}{\sqrt{\dfrac{Var(S_i|c_2) + Var(S_i|c_1)}{2}}} \tag{9}$$

and $\Phi(\cdot)$ is the standard normal integral function. The quantity in equation 9 is called *d-prime*, and it is the absolute value of the difference in the expected values (means) for the two categories divided by the square root of the average of the variances for the two categories. Intuitively, *d-prime* is a signal-to-noise ratio; the signal is the difference in the means, and the noise is the square root of the average variance. These formulas for ideal observer performance are often quite accurate, even when the probability distributions for components deviate substantially from the normal distribution (although formulas become inaccurate for severe deviations).

It can be shown (e.g., Green and Swets, 1966) that the accuracy of an observer that optimally combines all the stimulus components is given by

$$PC_{ideal} = \Phi\left(\frac{d_{ideal}'}{2}\right) \tag{10}$$

where

$$d_{ideal}' = \sqrt{\sum_{i=1}^{n}(d_i')^2} \tag{11}$$

In words, the *d*-prime for an ideal observer combining independent sources of information is simply the square root of the sum of the squared *d*-primes for each source alone.

In psychophysical experiments, the measured performance accuracy can be converted into a *d*-prime value, d'_{real}. One useful measure of the difference between real and ideal performance is the "efficiency" η, which is defined as the square of the ratio of the real and ideal values of *d*-prime (Tanner and Birdsall, 1958):[6]

---

[6] This is a generalization of Tanner and Birdsall's definition.

$$\eta = \frac{(d'_{real})^2}{(d'_{ideal})^2} \tag{12}$$

PHOTON NOISE The first useful applications of ideal observer analysis in vision were directed at understanding the limits to visual performance imposed by the randomness of light and then determining how closely the human visual system approaches those limits (Barlow, 1957, 1958b; De Vries, 1943; Hecht et al., 1942; Rose, 1948). It was well known at the beginning of the twentieth century that the number of photons emitted by a light source (or absorbed by a material) in a fixed time interval is generally described by the Poisson probability density:

$$p(z) = \frac{e^{-a}a^z}{z!} \tag{13}$$

where z is the specific number of photons emitted (or absorbed) and a is the mean number of photons emitted (or absorbed).

In one of the earliest studies of how photon noise might affect visual performance, Hecht et al. (1942) measured threshold for detection of a spot of light in the dark. In this detection task, one stimulus category (c_1) is a completely dark background and the other (c_2) is a spot of light against the dark background. The investigators chose conditions likely to yield the lowest possible thresholds: a small spot with a wavelength at the peak of the rod spectral sensitivity function (507 nm), presented briefly at the eccentricity where rod receptor density is greatest. They found that threshold for this stimulus was approximately 100 photons at the cornea. Depending on one's estimates of the transmittance of the ocular media, light collection area of the rods, optical density of the rod photopigment, and isomerization probability given photon absorption, this threshold translates into something like 10 to 20 effective photon absorptions scattered among a few hundred rods (Barlow, 1977; Hecht et al., 1942). If this is the average number of effective photon absorptions at detection threshold, then the probability of absorbing no photons is quite small, and hence the human observer must be performing considerably below the level of the ideal observer (whose performance is given by combining equations 4 and 8). Nonetheless, this threshold is small enough to imply that photon noise may be an important factor limiting human vision.

It was realized that if photon noise does limit human vision, then intensity discrimination (i.e., contrast detection) should follow the square-root relation implied (approximately) by the photon-noise-limited ideal observer (DeVries, 1943; Rose, 1948):

$$\Delta a = k a^{0.5} \tag{14}$$

where a is the average number of photons received from the background, $\Delta a + a$ is the average number of photons received from the background plus target, and k is a constant determined by the percentage of correct responses used to define threshold. The symbols in Figure 52.1 show the exact performance of the photon-noise-limited ideal observer for a number of background intensities expressed in units of quanta. The symbol at a background of 0.0 shows the absolute threshold of the ideal observer. The *straight line* is the approximation to the ideal observer given by equation 14, which is accurate for backgrounds above a few quanta. The *solid curve* is the approximation to the ideal observer obtained using the normality assumption described in the section "Optimal Discrimination Given Statistically Independent Sources of Information."

A number of human psychophysical studies have demonstrated that there is a substantial range of background intensities and target shapes where contrast detection follows the square-root relation both under rod-dominated (scotopic) conditions (e.g., Barlow, 1957; Blakemore and Rushton, 1965) and under cone-dominated (photopic) conditions (e.g., Banks et al., 1987; Barlow, 1958a; Kelly, 1972). For example, the symbols in Figure 52.2 show the contrast sensitivity (1 / contrast threshold) measured in the human fovea as a function of the spatial frequency of sine wave

targets for three background intensity levels. The spacing between the solid curves equals the value ($\sqrt{10}$) predicted by the square-root law.

There is additional evidence that photon noise may play a role in contrast detection. For example, in psychophysical studies, the shape of the receiver operating characteristic (ROC) for increment and decrement targets is often consistent with photon noise (Cohn and Lasley, 1986), and in electrophysiological studies, individual primate rods (Baylor et al., 1984) and cat ganglion cells (Barlow et al., 1971) produce reliable responses to the absorption of single photons. However, none of this evidence is definitive. Humans generally perform considerably worse than the photon-noise-limited ideal observer, making it quite possible that other factors (e.g., Poisson-like neural noise) are responsible for those visual performance characteristics that appear to be consistent with photon noise (Graham and Hood, 1992; Kortum and Geisler, 1995).

OPTICS AND PHOTORECEPTORS The optics of the eye, photoreceptor lattice, and absorption spectra of the photopigments are relatively well understood in humans and macaques, making possible a relatively rigorous ideal
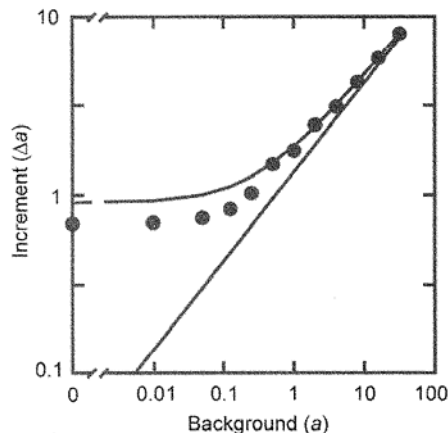

FIGURE 52.1. Performance of the photon-noise-limited ideal observer in a single-interval, two-alternative, forced-choice intensity discrimination task where the presentation probabilities for background and background plus target are equal and the criterion for threshold is 75% correct. The *solid symbols* show the exact predictions for a number of background energies (in units of quanta). The lumpiness of the predictions is due to the discrete nature of the Poisson probability density. The *solid straight line* shows the prediction of equation 14, which is reasonably accurate for background energies above a few quanta. The *solid curve* is based on the normal approximation that is commonly used in computing ideal observer predictions. It is quite accurate for background energies above 0.5 quanta, and for most purposes it is sufficiently accurate down to absolute threshold.
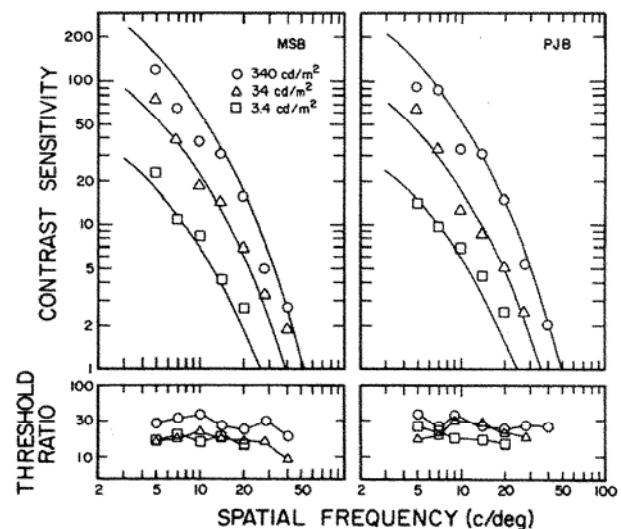

FIGURE 52.2. Contrast sensitivity functions measured in two human observers at three background intensity levels for sine wave targets having a fixed number of cycles. The targets were presented for 0.1 second in the center of the fovea. The solid curves are the performances of an ideal observer limited by photon noise, the optics of the eye, the photoreceptor lattice, and the absorption spectra of the cone photopigments. The quantum efficiency of the ideal observer has been adjusted so that the solid curves align optimally with the data. The spacing between the solid curves is due to photon noise and corresponds to the square-root law. The lower panels show the ratio of ideal threshold to real threshold, when there is no adjustment of quantum efficiency. (Adapted from Banks et al., 1987).

Point spread Ocular transmittance Receptor lattice Cone spectra

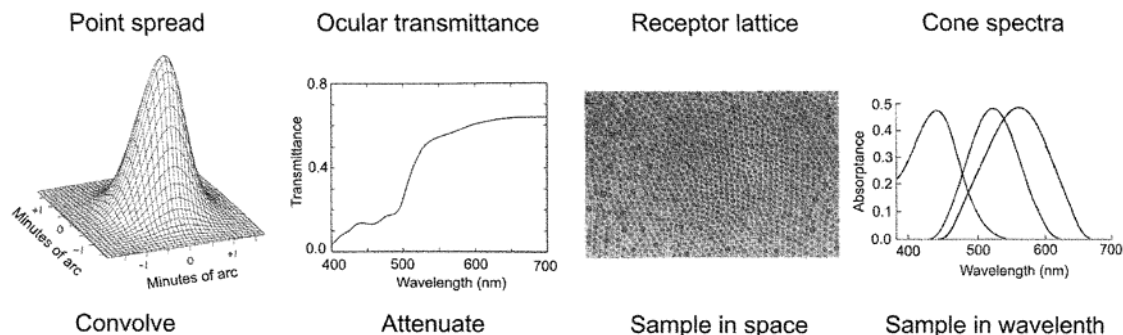### Convolve Attenuate Sample in space Sample in wavelenth

FIGURE 52.3. Computing the mean number of photons absorbed in each photoreceptor. First, the image is convolved with appropriate point spread functions. Second, the chromatic spectral energy distribution at each image location is attenuated by the ocular transmittance function. Third, the blurred and attenuated spectral energy distribution is summed over the aperture of each photoreceptor. Fourth, the spectral energy distribution entering each photoreceptor is multiplied by the effective absorption spectrum for that receptor, integrated, and converted to units of quanta.

observer analysis of how these factors limit detection and discrimination performance (Banks and Bennett, 1988; Banks et al., 1987; Beckmann and Legge, 2002; Geisler, 1984, 1989). The number of photons absorbed in a photoreceptor is described by a Poisson probability density (equation 8), and the number of photons absorbed in each receptor is statistically independent of the number absorbed in any other receptor (across presentations of the same stimulus). Thus, performance of the ideal observer can be easily calculated using the formulas in the section "Optimal Discrimination Given Statistically Independent Sources of Information" if the parameter (mean) of the Poisson distribution is known for each photoreceptor for the two stimuli being discriminated.

Typical steps for calculating the mean number of photons absorbed in each photoreceptor are illustrated in Figure 52.3: (1) convolve with the optical point spread function, (2) attenuate across wavelength using the ocular transmittance function, (3) sum the light across the aperture of each photoreceptor in the lattice, and (4) multiply by each photoreceptor's absorption spectrum and integrate across the wavelength. In terms of the general terminology of the section "Constrained Bayesian Ideal Observers," these processing steps constitute a constraint function $g_\theta(\mathbf{S})$ that maps the stimulus into an intermediate signal—the pattern of photon absorptions in the photoreceptors—although in this case without any free parameters.

Suppose that after applying these steps the mean number of photon absorptions in the ith photoreceptor is a_i for stimulus a and is b_i for stimulus b. Since the variance of the Poisson probability distribution is equal to its mean, the d-prime for the ideal observer using only the ith photoreceptor is

$$d_i' = \frac{|b_i - a_i|}{\sqrt{\dfrac{b_i + a_i}{2}}} \qquad (15)$$

and hence, by equation 11, d-prime for the full ideal observer is

$$d_{\text{ideal}}' = \sqrt{2 \sum_{i=1}^{n} \frac{(b_i - a_i)^2}{b_i + a_i}} \qquad (16)$$

To determine the discrimination threshold for the ideal observer, the difference between stimulus a and stimulus b is varied (along the dimension of interest) until the d_{ideal}' reaches a value corresponding to the chosen criterion level of accuracy (typically 75% correct).[7] The solid curve in Figure 52.1 shows the thresholds obtained with equation 16 for an intensity discrimination task where stimulus a and stimulus b are identical except for intensity (i.e., a scale factor). Ideal observer analyses at the level of the photoreceptors in humans have been carried out for a number of different detection, discrimination, and identification tasks.

One area that has received considerable attention is spatial vision. Banks et al. (1987) measured the high-frequency limb of the contrast sensitivity function in the fovea for briefly presented sine wave targets with a fixed number of spatial cycles. The symbols show the contrast sensitivities measured at three background intensities. The solid curves in Figure 52.2 show the performance of the ideal observer shifted vertically on the log contrast axis. (The vertical shifting, which is done to compare the shapes of the real and ideal threshold functions, corresponds to scaling the efficiency of the ideal observer down by a constant factor.[8]) The lower panels show the actual ratios of real to ideal

[7] For most stimulus conditions, these formulas for the ideal observer at the level of the photoreceptors are slightly less accurate than those described in Geisler (1989), but they are sufficiently accurate for most purposes, and are simpler and more intuitive.
[8] For a photon-noise-limited ideal observer, scaling the efficiency down by a constant factor is equivalent to placing a neutral density filter in front of the eye.

thresholds. The fact that the ratios are approximately constant as a function of spatial frequency and contrast suggests that much of the measured variation in human performance may be due to preneural factors. The fact that the ratios are high (approximately 20) implies that the overall efficiency of the neural processing subsequent to photon absorption in the receptors is relatively low (less than 1%).

Banks and Bennett (1988) performed a similar analysis of contrast sensitivity in human infants and found that a substantial fraction of the difference between adult and infant contrast sensitivity is consistent with the optical and photoreceptor immaturities that have been measured in the infant eye. Davila and Geisler (1991) showed that detection thresholds measured for spot targets as a function of target area vary in a similar fashion, for human and ideal observers, for targets up to several hundred square minutes of arc (Fig. 52.4). These results are consistent with the measurements of the contrast sensitivity function measured by Banks et al. (1987); cf. Figure 52.2.

Geisler (1984) and Geisler and Davila (1985) showed that some of the dramatic differences in performance observed across different types of acuity tasks are qualitatively consistent with the performance of an ideal observer operating at the level of the photoreceptors. In traditional acuity tasks, humans (with normal vision) can resolve changes of approximately 45 to 60 seconds of arc in the spatial position of two overlapping image features. This corresponds to a change in spacing of two foveal cone diameters. However, in *hyperacuity* tasks, humans can resolve changes in spatial position of a few seconds of arc if the features are slightly separated (e.g., Westheimer, 1979). The ideal observer shows similar differences in acuity because of differences in the nature of the discrimination information when features overlap on the retina as opposed to when they do not. This difference in the nature of the information also leads to rather different performance as a function of intensity: the performance of the ideal observer shows that the physical limit for resolving two overlapping features decreases with the fourth root of intensity, whereas the physical limit for resolving changes in the position of two spatially separated features decreases with the square root of intensity. This difference is also seen qualitatively in psychophysical studies (Geisler and Davila, 1985).

Another area that has received considerable attention is color vision. Once again, there are a number of examples in which human performance is qualitatively, and sometimes quantitatively, similar to that of an ideal observer at the level of the photoreceptors. Wavelength discrimination functions of humans are similar in shape to that of the ideal observer (Geisler, 1989; Vos and Walraven, 1972). The high spatial frequency limb of the chromatic contrast sensitivity function in humans is similar in shape to that of the ideal observer (Geisler, 1989; Sekiguchi et al., 1993). Some of the deficits
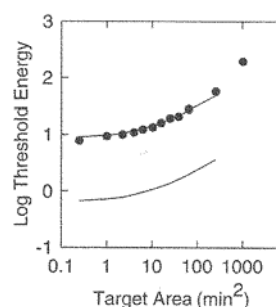

FIGURE 52.4. Detection threshold for spot targets in the central fovea as a function of target area. The data points are thresholds for a human observer, and the solid curves are the thresholds for an ideal observer operating at the level of photon absorptions in the receptors. The *lower solid curve* shows the thresholds of the ideal observer; the *upper solid curve* is the same curve translated vertically for the purpose of comparing shapes. (Adapted from Davila and Geisler, 1991.)

in color discrimination performance in human infants are consistent with the performance of an ideal observer that incorporates the known immaturities in the infant's optics and photoreceptors (Banks and Bennett, 1988).

The examples described above demonstrate that there are many tasks where human discrimination performance parallels the performance of an ideal observer limited by preneural factors. Thus, some of the variation in human performance in these tasks would seem to be explained by preneural factors in the sense that subsequent neural mechanisms must be extracting the available information from the photoreceptors with relatively constant efficiency. The importance of this observation is that it moves physiological research forward from the question of what mechanisms are responsible for the variations in discrimination performance for the stimuli entering the eye to the question of what mechanisms are responsible for achieving the nearly constant discrimination performance for the signals exiting the photoreceptors.

Of course, there are many tasks where human performance does not parallel the performance of the ideal observer limited by preneural factors. For example, even though the density of cones and the quality of the optics of the eye decline with eccentricity from the fovea, the decline in contrast sensitivity at high spatial frequencies is much more precipitous than predicted by the ideal observer (Banks et al., 1991). Similarly, the decline in letter identification performance drops much more quickly than that predicted by the ideal observer (Beckman and Legge, 2002). In these cases, most (or all) of the variation in performance must be due to neural mechanisms. These are also important results because they localize the relevant mechanisms beyond the photoreceptors.

NEURAL FACTORS IN THE RETINA AND PRIMARY VISUAL CORTEX The neural mechanisms in the early stages of the visual pathway are less well understood than the preneural factors (optics, receptor lattice, photopigments); nonetheless, there is sufficient knowledge to carry out limited ideal observer analyses, with the caveat that the analyses may change substantially as new anatomical and physiological knowledge accumulates. Banks et al. (1991) extended the earlier ideal observer analysis of Banks et al. (1987) to include a more complete description of the preneural factors, as well as a description the spatial pooling (summation) implied by the density of the ganglion cells and the sizes of their center mechanisms. They found that the performance of the ideal observer paralleled human performance in several detection and discrimination tasks measured as a function of retinal eccentricity. For example, Figure 52.5 shows grating acuity (*solid squares*) and vernier acuity (*solid circles*) as a function of retinal eccentricity. The *left panel* shows human performance reported by Westheimer (1982), and the *right panel* shows the performance of the ideal observer. The *open symbols* show the performance of an ideal observer that is limited only by preneural factors. As can be seen, the ideal observer that includes retinal spatial pooling parallels human performance considerably better than one that does not. Further, the retinal ideal observer displays the interesting property that vernier acuity declines more rapidly than grating acuity as a function of eccentricity. This is a counterintuitive prediction that further illustrates the value of ideal observer analysis for understanding the information processing consequences of the stimulus, task, and physiological/anatomical factors. As Banks et al. point out, there are aspects of the variation in human visual acuity with eccentricity unlikely to be explained by retinal factors; however, an ideal observer analysis (or something equivalent) is essential for determining those aspects.

Arnow and Geisler (1996) performed an ideal observer analysis similar to that of Banks et al., but they included some additional retinal factors based on the receptive field properties of ganglion cells in the macaque monkey (Croner and Kaplan, 1995; Croner et al., 1993). They found that ideal observer contrast sensitivity, at the level of the ganglion cell responses, parallels human contrast sensitivity fairly well as a function of target spatial frequency, eccentricity, and size (see also Geisler and Albrecht, 2000).

There are detection and discrimination tasks where human performance does not parallel the performance of the ideal observer at the level of the ganglion cell responses. These tasks include contrast discrimination and spatial frequency masking. Geisler and Albrecht (1997) performed an ideal observer analysis of contrast discrimination and spatial frequency discrimination at the level of the responses of neurons in primary visual cortex. However, this analysis is
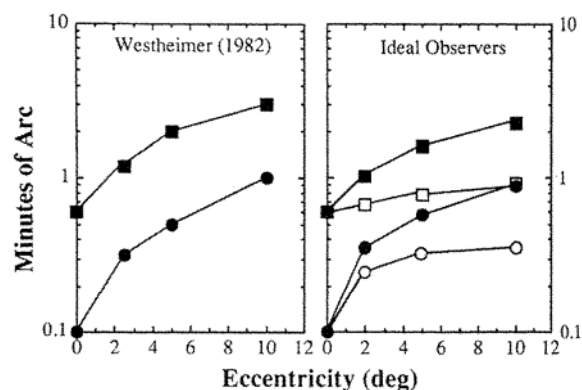

FIGURE 52.5. Grating acuity (*squares*) and vernier acuity (*circles*) as a function of retinal eccentricity. The *left panel* shows human performance (Westeimer, 1982). The *right panel* shows the performance of an ideal observer at the level of the photoreceptors (*open symbols*) and an ideal observer that also incorporates spatial summation consistent with the center sizes and spatial density of the retinal ganglion cells. (Adapted from Banks et al., 1991.)

more tentative because of the vast number of neurons in V1, their highly heterogeneous receptive field properties, and the potential for long-range interactions within V1 and from other cortical areas. The strategy was to measure (with single-unit electrophysiology) the spatial frequency tuning functions, contrast response functions, and noise characteristics of a large population of cortical neurons. From these response functions and noise measurements, Geisler and Albrecht determined the ideal contrast discrimination performance and spatial frequency discrimination performance of each cortical neuron in the population using equation 9. Finally, they determined the ideal performance for the whole population under the assumption of statistical independence equation 11. The solid curves in Figure 52.6 show the shapes of the contrast discrimination and spatial frequency discrimination functions from the ideal observer analysis. The open symbols show the data from several different studies in humans and monkeys (see figure caption). For comparison, note that the contrast discrimination function of an ideal observer at the level of the photoreceptors is flat; that is, contrast threshold is constant independent of the background contrast.

PIXEL NOISE, NEURAL NOISE, AND CENTRAL EFFICIENCY In the tasks considered so far, photon noise was the only source of stimulus noise (other than the random selection of the stimulus category on each trial). However, in natural and artificial environments, signals of interest are often embedded in complex, randomly varying background patterns which act as a source of stimulus noise. Thus, an important class of ideal observer analyses is those directed at tasks where pixel noise is added to the stimulus display. For
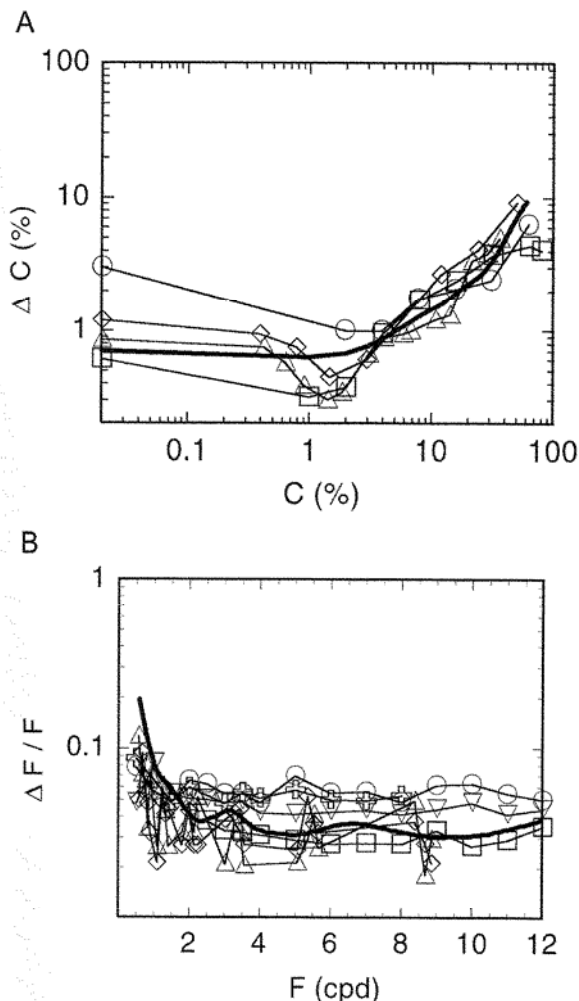
A

B

FIGURE 52.6.   *A*, Contrast discrimination threshold as function of base contrast. Open symbols are data from psychophysical studies (several in humans and one in monkey). *B*, Spatial frequency discrimination threshold as a function of base frequency. Open symbols are data from several psychophysical studies in humans. The solid curves show the relative performance of an ideal observer that combines the responses of a population of neurons whose contrast-response functions, spatial-frequency-tuning functions, and noise characteristics were measured one at a time in monkey V1. The efficiency of the ideal observer was reduced by the same factor in both plots to allow comparison of shapes.

example, Figure 52.7 shows human and ideal performance for amplitude discrimination of small targets (spots and grating patches) in white noise as a function of noise spectral power density (Burgess et al., 1981). The solid line of slope 1.0 shows the absolute performance of an ideal observer operating at the level of the cornea. Thus, the difference between real and ideal performance represents all losses of information within the eye, retina, and central visual pathways. For these conditions, efficiency ranges from about 20% to 70%, much higher than the efficiency for similar targets in uniform backgrounds (<1%). In other
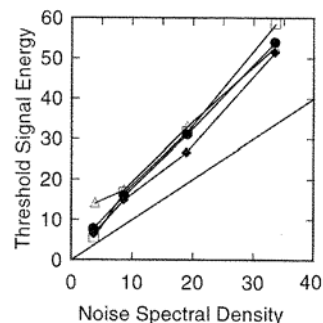
FIGURE 52.7.   Comparison of real and ideal performance for the discrimination of simple localized targets in white noise. The diagonal line of slope 1.0 shows the performance of the ideal observer; the symbols show the performances of human observers for several different spot and sine wave grating targets. (Adapted from Burgess et al., 1981.)

words, the human observers are performing much more closely to ideal in the pixel noise tasks. One possible explanation for this difference between tasks is that the pixel noise dominates the photon noise, and other sources of noise in the early levels of the visual system, effectively sidestepping inefficiencies of the peripheral visual system. Thus, the difference in performance between real and ideal observers in Figure 52.7 may reflect primarily the inefficiencies of central decision mechanisms (e.g., Barlow, 1978).

Pixel noise experiments, in conjunction with appropriate ideal observer analyses, have been used to isolate and measure central mechanisms that limit discrimination performance (Barlow, 1978; Burgess et al., 1981; Kersten, 1987; Pelli, 1990), to evaluate neural noise levels in the early visual system (Legge et al., 1987; Pelli, 1990), and to develop models relevant for understanding the detection and discrimination of targets in the noisy images created by radiological devices and image enhancers (Barrett et al., 1992; Burgess et al., 1982; Myers and Barrett, 1987; Myers et al., 1985).

Recently, new insights into the efficiency of central pooling and decision mechanisms have been gained by analyzing the samples of noise on each experimental trial, contingent upon the observer's response on that trial (Abbey et al., 1999; Ahumada, 1996; Beard and Ahumada, 1999; Gold et al., 2000). Specifically, suppose that there are two possible targets, *a* and *b*, and that on each trial one of the targets is randomly selected and then added to a random sample of white noise. By averaging together all samples of noise when the observer responded "*b*," averaging together all the samples of noise when the observer responded "*a*," and then subtracting these two average noise images, one obtains a *classification image*. If a sufficient number of such trials are run, the classification image provides a detailed map indicating the weight the observer placed on each pixel location in making the decision. This classification image

can be compared with the one generated by the ideal observer to give detailed information about inefficiencies in the real observer's central pooling and decision mechanisms. Measurements of classification images show that observers tend to make use of only a subset of the pixel locations that contain useful information, and sometimes use locations that contain no useful information (e.g., locations where target contours are occluded; Gold et al., 2000).

Natural scene statistics and natural selection

The introduction to this chapter began with the truism that the proper study of a visual system must include an analysis of those specific tasks that the system evolved to perform. However, as we have seen, ideal observer analyses have been largely confined to tasks involving relatively simple stimuli generated in the laboratory. Recently, measurements of statistical properties of natural environments have become available, allowing ideal observer analyses for tasks involving more naturalistic stimuli. This is an important direction for research because the results can speak more directly to the relationship between the statistics of natural environments and the design of perceptual systems.

One topic that has received considerable interest is color identification. For example, Regan et al. (1998, 2001) measured the wavelength distributions of primary food sources (fruits) of several New World monkeys and the wavelength distributions of the surrounding foliage. They then used an ideal observer analysis to determine optimal placement of the M and L cones for identifying food sources in the surrounding foliage. Interestingly, optimal placement corresponds fairly well with actual placement, although, as Regan et al. pointed out, other factors (such as minimizing chromatic aberration) may also contribute to actual placement. Osorio and Vorobyev (1996) have made a similar case for placement of the cone photopigments in Old World primates. Similarly, the two cone pigments in dichromatic mammals appear to be nearly optimally placed for discriminating between natural leaf spectra (Chiao et al., 2000; Lythgoe and Partridge, 1989).

Another topic that has received some attention is contour detection (Elder and Zucker, 1998; Geisler et al., 2001; Sigman et al., 2001). For example, Geisler et al. (2001) extracted edge elements from images of diverse natural scenes and then computed co-occurrence probabilities for different possible geometrical relationships between the edge elements. (Note that the geometrical relationship between a pair of edge elements is described by a vector in three dimensions: distance, direction, and orientation difference.) Two different co-occurrence probability distributions were measured: one for edge elements that belong to the same physical contour (i.e., the same surface boundary, shadow/lighting boundary, or surface-marking boundary)

and one for edge elements that belong to different physical contours. Using these two probability distributions, Geisler et al. derived an ideal observer for detecting contours embedded in complex backgrounds and compared its performance to human performance on the same tasks. Remarkably, the performance of the ideal observer based on natural image statistics was quite similar to human performance across all conditions—the correlation between human and ideal detection accuracy was approximately 0.9. This result suggests that there is a close relationship between contour grouping mechanisms in humans and the statistics of contours in natural images.

These two examples, and a number of the examples described earlier, demonstrate that there can be a close correspondence between real and ideal observers. Nonetheless, there are many reasons to expect real observers not to reach the performance of the ideal observer:

1. In the laboratory, the task (including the stimulus likelihoods and prior probabilities) is defined by the experimenter, and may not correspond well with the tasks that the organism has evolved or learned to perform in the natural environment.

2. Organisms evolve or learn to perform many different tasks, and hence there may be compromises in design that lead to nonideal performance in a given task.

3. There are limits to the range of materials that organisms can synthesize and exploit, and limits on the possible structure of organic molecules, but there are not similar limits on the ideal observer.

4. Perceptual systems (and the learning mechanisms that shape them) are designed through natural selection, and thus the intrinsic utility function is fitness (birth and death rates), which may imply an ideal observer rather different from the one implied by the utility function specified in a laboratory task.

5. Evolution through natural selection is an incremental process in which each change must produce an increase in fitness; thus, the real observer may correspond to a local maximum in the space of possible solutions, whereas the ideal observer corresponds to the global maximum in the space of possible solutions.

6. In general, evolution lags behind changes that occur in environmental likelihoods and prior probabilities; thus, a real observer may not even correspond to a local maximum in the space of possible solutions.

MAXIMUM FITNESS IDEAL OBSERVERS   One way to begin understanding the differences between real observers and ideal observers is to measure properties of the natural environment (e.g., natural scene statistics) and, from these, design more naturalistic laboratory tasks. The most appropriate ideal observers for analyzing such tasks are those where the measure of utility is fitness (birth and death rates), although

this may often not be practical. The fitness utility function can be represented as a growth factor function, $\gamma(\mathbf{r}, \omega)$, which equals 1 plus the birth rate minus the death rate for each possible response and state of the environment (Geisler and Diehl, 2002). Thus, given a particular stimulus $\mathbf{S}$, the maximum-fitness ideal observer will make the response that maximizes the growth factor averaged across all possible states of the environment. In other words, the maximum-fitness ideal observer will make the response

$$\mathbf{R} = \arg\max_r \left[\sum_\omega \gamma(\mathbf{r}, \omega) p(\mathbf{S}|\omega) p(\omega) \right] \quad (17)$$

This equation is identical to the standard Bayesian ideal observer equation 5, except that the utility function is the growth-factor function.

Geisler and Diehl (2002) have demonstrated that utility functions based on fitness can yield ideal observers that behave quite differently from those based on more traditional utility functions. For example, Figure 52.8 compares the decision criterion of an ideal observer that maximizes accuracy with one that maximizes fitness. In the hypothetical scenario they considered, a predator species is trying to detect prey (its only food source). All other things being equal, natural selection favors mutations that result in an increase in the birth rate of the predator. However, when birth rate increases, the result is a decline in the number of prey and hence a reduction in the prior probability that a prey is in the immediate vicinity. If the utility function of the ideal observer corresponds to maximizing prey detection accuracy (the typically utility function in laboratory experiments), then the optimal decision criterion equals the ratio of the prior probabilities, which is shown by the *dashed curve* in Figure 52.8. On the other hand, if the goal is to maximize fitness, then the optimal decision criterion is relatively invariant with birth rate (*solid curve*), because the decrease in target prior probability is balanced by the increase in payoff when a prey is captured. This example demonstrates the potential importance of considering fitness utility functions when evaluating the performance of real observers.

BAYESIAN NATURAL SELECTION A complementary approach to understanding the differences between real observers and ideal observers is to measure properties of the natural environment and incorporate those into a model that represents the process of natural selection, which, unlike the ideal observer, does not necessarily find the global optimum. It is possible to formulate a quantitative version of the theory of natural selection that incorporates the same terms—prior probability distributions, stimulus likelihood distributions, and utility functions—as those in a Bayesian ideal observer (Geisler and Diehl, 2002). This Bayesian formulation of natural selection provides a convenient conceptual frame-
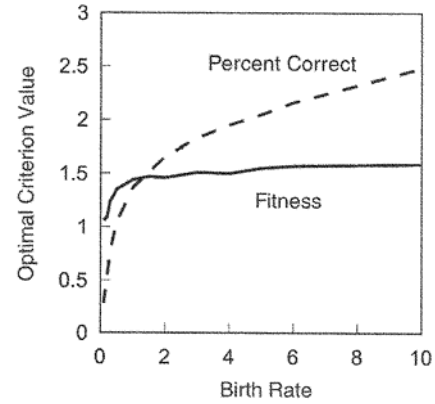

FIGURE 52.8. Comparison of an ideal observer having a utility function based on maximizing growth rate (1 + birth rate − death rate) versus a more standard ideal observer having a utility function based on maximizing detection accuracy. The *dashed curve* shows the decision criterion of an ideal predator that is maximizing prey detection accuracy as a function of the predators' birth rate (which causes a decrease in the prior probability of a prey being within the local vicinity). The *solid curve* shows the decision criterion of the ideal predator that is maximizing growth rate. The lumpiness of the solid curve is due to noise in the simulation process. (Adapted from Geisler and Diehl, 2002.)

work for understanding how natural scene statistics, and other properties of the environment, influence the evolution and hence design of a perceptual system.

Briefly, the fundamental equation of Bayesian natural selection shows how the expected number of organisms of a given species carrying a given vector of alleles $\mathbf{a}$ at time $t+1$ is related to the number of organisms carrying the same allele vector at time t, the prior probability of a state of the environment at time t, the likelihood of a stimulus given the state of the environment, the likelihood of a response given the stimulus, and the growth factor given the response and the state of the environment:

$$\overline{O}_a(t+1) = O_a(t) \sum_\omega p_a(\omega;t) \sum_\rho \gamma_a(\mathbf{r}, \omega) \sum_\sigma p_a(\mathbf{r}|\mathbf{s}) p_a(\mathbf{s}|\omega) \quad (18)$$

The actual number of organisms at time $t+1$ is a random number, $O_a(t+1)$, that is obtained by sampling from appropriate probability distributions for births, deaths, mutations, and sexual recombination. A separate fundamental equation is set up for each allele vector in each species under consideration, and the process of natural selection is represented by iteration of these fundamental equations in parallel over time. (Note that mutations and sexual recombination may create new allele vectors and, hence, new fundamental equations, and that extinction of alleles or species will remove fundamental equations.)

Using this formulation of natural selection and the corresponding maximum fitness ideal observer, Geisler and Diehl (2002) explored the coevolution of receptor spectral

sensitivities in a hypothetical predator species and camouflage in a hypothetical prey species. They found that there are many starting conditions where the spectral sensitivities (and decision criterion) of the predator species converge to that of the maximum-fitness ideal observer. Further, they found that there are different stable coevolutionary end states, and that these different end states are not equally good in terms of the average growth rate (fitness) for either the predator or the prey. This may seem contradictory (how can two different solutions both represent maximum fitness?), but in fact, it is expected because in realistic evolutionary scenarios the prior probability distributions change over time, and hence the maximum-fitness ideal observer changes over time. If the stable end-state prior probabilities are different, then the maximum-fitness ideal observer may be different. This example demonstrates the potential importance of measuring natural scene statistics, and of considering the incremental process of natural selection, when evaluating the performance of real observers.

Conclusion

This chapter outlined the basic concepts and formulas of ideal observer analysis as it has been applied in studies of biological vision, and then illustrated the application of ideal observer analysis in a number of cases involving detection, discrimination, and identification tasks. Although real visual systems are never ideal, deriving the performance of an ideal observer can be very useful for understanding the computational requirements of a task and the limits to performance imposed by specific anatomical and physiological factors. The range of situations where ideal observer analysis can be usefully applied has grown over the years, and it is likely to become an even more central tool as we attempt to understand natural scene statistics and their relationship to the design and evolution of visual systems.

REFERENCES

Abbey, C. K., M. P. Eckstein, and F. O. Bochud, 1999. Estimation of human-observer templates in two-alternative forced-choice experiments, *SPIE*, 3663:284–295.

Ahumada, A. J., 1996. Perceptual classification images from vernier acuity masked by noise, *Perception*, 25:18.

Arnow, T. L., and W. S. Geisler, 1996. Visual detection following retinal damage: predictions of an inhomgeneous retino-cortical model. *SPIE Proc. Hum. Vis. Electr. Imaging*, 2674:119–130.

Banks, M. S., and P. J. Bennett, 1988. Optical and photoreceptor immaturities limit the spatial and chromatic vision of human neonates. *J. Opt. Soc. Am. A*, 5:2059–2079.

Banks, M. S., W. S. Geisler, and P. J. Bennett, 1987. The physical limits of grating visibility, *Vis. Res.*, 27:1915–1924.

Banks, M. S., A. B. Sekuler, and S. J. Anderson, 1991. Peripheral spatial vision: limits imposed by optics, photoreceptors, and receptor pooling, *J. Opt. Soc. Am. A*, 8:1775–1787.

Barlow, H. B., 1957. Increment thresholds at low intensities considered as signal/noise discriminations, *J. Physiol. (Lond.)*, 136:469–488.

Barlow, H. B., 1958a. Intrinsic noise of cones, in *National Physical Laboratory Symposium on Visual Problems of Colour*, vol. 2, London: Her Majesty's Stationery Office, pp. 617–630.

Barlow, H. B., 1958b. Temporal and spatial summation in human vision at different background intensities, *J. Physiol. (Lond.)*, 141:337–350.

Barlow, H. B., 1977. Retinal and central factors in human vision limited by noise, in *Vertebrate Photoreception* (H. B. B. P. Fatt, ed.), London: Academic Press, pp. 337–358.

Barlow, H. B., 1978. The efficiency of detecting changes of density in random dot patterns, *Vis. Res.*, 18:637–650.

Barlow, H. B., W. R. Levick, and M. Yoon, 1971. Responses to single quanta of light in retinal ganglion cells of the cat, *Vis. Res.*, 11(suppl. 3):87–102.

Barrett, H. H., T. Gooley, K. Girodias, J. Rolland, T. White, and J. Yao, 1992. Linear discriminants and image quality, *Image Vis. Comput.*, 10:451–460.

Baylor, D. A., B. J. Nunn, and J. L. Schnapf, 1984. The photocurrent, noise and spectral sensitivity of rods of the monkey *Macaca fascicularis*, *J. Physiol. (Lond.)*, 357:575–607.

Beard, B. L., and A. J. Ahumada, 1999. Classification images for detection, *Invest. Ophthalmol. Vis. Sci.*, 40:3015.

Beckmann, P. J., and G. E. Legge, 2002. Preneural limitations to identification performance in central and peripheral vision, *J. Opt. Soc. Am. A*, 19:2349–2362.

Berger, J. O., 1985. *Statistical Decision Theory and Bayesian Analysis*, 2nd ed., New York: Springer-Verlag.

Blakemore, C. B., and W. A. H. Rushton, 1965. The rod increment threshold during dark adaptation in mornal and rod monochromat, *J. Physiol.*, 181:629–640.

Burgess, A. E., R. F. Wagner, and R. J. Jennings, 1982. Human signal detection performance for noisy medical images. In O. Nalcioglu, and M. S. Prewitt (eds.), *Proceedings of International Workshop on Physics and Engineering in Medical Images*: New York: IEEE.

Burgess, A. E., R. F. Wagner, R. J. Jennings, and H. B. Barlow, 1981. Efficiency of human visual signal discrimination, *Science*, 214:93–94.

Chiao, C., M. Vorobyev, T. W. Cronin, and D. Osorio, 2000. Spectral tuning of dichromats to natural scenes, *Vis. Res.*, 40:3257–3271.

Cohn, T. E., and D. J. Lasley, 1986. Visual sensitivity, *Annu. Rev. Psychol.*, 37:495–521.

Croner, L. J., and E. Kaplan, 1995. Receptive fields of P and M ganglion cells across the primate retina, *Vis. Res.*, 35:7–24.

Croner, L. J., K. Purpura, and E. Kaplan, 1993. Response variability in retinal ganglion cells of primates, *Proc. Natl. Acad. Sci. USA*, 90:8128–8130.

Davila, K. D., and W. S. Geisler, 1991. The relative contributions of pre-neural and neural factors to areal summation in the fovea, *Vis. Res.*, 31:1369–1380.

De Vries, H. L., 1943. The quantum character of light and its bearing upon threshold of vision, the differential sensitivity and visual acuity of the eye, *Physica*, 10(7):553–564.

Elder, J. H., and S. W. Zucker, 1998. Evidence for boundary-specific grouping, *Vis. Res.*, 38(1):143–152.

Geisler, W. S., 1984. Physical limits of acuity and hyperacuity, *J. Opt. Soc. Am. A*, 1:775–782.

Geisler, W. S., 1989. Sequential ideal-observer analysis of visual discriminations, *Psychol. Rev.*, 96:267–314.

Geisler, W. S., and D. G. Albrecht, 1997. Visual cortex neurons in monkeys and cats: detection, discrimination, and identification, *Visual Neurosci.*, 14:897–919.

Geisler, W. S., and D. G. Albrecht, 2000. Spatial vision. In *Seeing* (2nd ed.) (K. K. De Valois, ed.), New York: Academic Press, pp. 79–128.

Geisler, W. S., and K. D. Davila, 1985. Ideal discriminators in spatial vision: two-point stimuli, *J. Opt. Soc. Am. A*, 2:1483–1497.

Geisler, W. S., and R. Diehl, 2002. Bayesian natural selection and the evolution of perceptual systems, *Philos. Trans. R. Soc. Lond. B*, 357:419–448.

Geisler, W. S., J. S. Perry, B. J. Super, and D. P. Gallogly, 2001. Edge co-occurrence in natural images predicts contour grouping performance, *Vis. Res.*, 41:711–724.

Gold, J. M., R. F. Murray, P. J. Bennett, and A. B. Sekuler, 2000. Deriving behavioural receptive fields for visually completed contours, *Curr. Biol.*, 10:663–666.

Graham, N., and D. C. Hood, 1992. Quantal noise and decision rules in dynamic models of light adaptation, *Vis. Res.*, 32:779–787.

Green, D. M., and J. A. Swets, 1966. *Signal Detection Theory and Psychophysics*, New York: Wiley.

Hecht, S., S. Shlaer, and M. H. Pirenne, 1942. Energy, quanta, and vision, *J. Gen. Physiol.*, 25:819–840.

Kelly, D. H., 1972. Adaptation effects on spatio-temporal sine-wave thresholds, *Vis. Res.*, 12:89–101.

Kersten, D., 1987. Statistical efficiency for the detection of visual noise, *Vis. Res.*, 27:1029–1040.

Knill, D. C., and W. Richards, eds. 1996. *Perception as Bayesian Inference*, Cambridge: Cambridge University Press.

Kortum, P. T., and W. S. Geisler, 1995. Adaptation mechanisms in spatial vision II: flash thresholds and background adaptation, *Vis. Res.*, 35(11):1595–1609.

Legge, G. E., D. Kersten, and A. E. Burgess, 1987. Contrast discrimination in noise, *J. Opt. Soc. Am. A*, 4(2):391–404.

Lythgoe, J. N., and J. C. Partridge, 1989. Visual pigments and the acquisition of visual information, *J. Exp. Biol.*, 146:1–20.

Myers, K. J., and H. H. Barrett, 1987. Addition of a channel mechanism to the ideal-observer model, *J. Opt. Soc. Am. A*, 4:2447–2457.

Myers, K. J., H. H. Barrett, M. M. Borgstrom, D. D. Patton, and G. W. Seeley, 1985. Effect of noise correlation on detectability of disk signals in medical imaging, *J. Opt. Soc. Am. A*, 2:1752–1759.

Osorio, D., and M. Vorobyev, 1996. Colour vision as an adaptation to frugivory in primates, *Proc. R. Soc. Lond. B*, 263:593–599.

Pelli, D. G., 1990. The quantum efficiency of vision. in *Vision: Coding and Efficiency* (C. Blakemore, ed.). Cambridge: Cambridge University Press, pp. 3–24.

Regan, B. C., C. Julliot, B. Simmen, F. Vienot, P. Charles-Dominique, and J. D. Mollon, 1998. Frugivory and colour vision in *Alouatta seniculus*, a trichromatic platyrrine monkey, *Vis. Res.*, 38:3321–3327.

Regan, B. C., C. Julliot, B. Simmen, F. Vienot, P. Charles-Dominique, and J. D. Mollon, 2001. Fruits, foliage and the evolution of primate colour vision, *Philos. Trans. R. Soc. Lond. B*, 356:229–283.

Rose, A., 1948. The sensitivity performance of the human eye on an absolute scale, *J. Opt. Soc. Am.*, 38:196–208.

Schrater, P. R., and D. Kersten, 2001. Vision, psychophysics and Bayes, in *Statistical Theories of the Brain* (R. P. N. Rao, B. A. Olshausen, and M. S. Lewicki, eds.), Cambridge, MA: MIT Press.

Sekiguchi, N., D. R. Williams, and D. H. Brainard, 1993. Efficiency in detection of isoluminant and isochromatic interference fringes, *J. Opt. Soc. Am. A*, 10:2118–2133.

Sigman, M., G. A. Cecchi, C. D. Gilbert, and M. O. Magnasco, 2001. On a common circle: Natural scenes and Gestalt rules, *Proc. Natl. Acad. Sci. USA*, 98:1935–1940.

Simoncelli, E. P., and B. A. Olshausen, 2001. Natural image statistics and neural representation, *Annu. Rev. Neurosci.*, 24:1193–1215.

Tanner, W. P., Jr., and T. G. Birdsall, 1958. Definitions of d' and n as psychophysical measures, *J. Acoust. Soc. Am.*, 30:922–928.

Vos, J. J., and P. L. Walraven, 1972. An analytical description of the line element in the zone-fluctuation model of color vision: 1. Basic concepts, *Vis. Res.*, 12:1327–1344.

Westheimer, G., 1979. The spatial sense of the eye, *Invest. Ophthalmol. Vis. Sci.*, 18:893–912.

Westheimer, G., 1982. The spatial grain of the perifoveal visual field, *Vis. Res.*, 22:157–162.