

Bayesian natural selection and the evolution of perceptual systems

Wilson S. Geisler* and Randy L. Diehl

Department of Psychology and Center for Perceptual Systems, University of Texas at Austin TX 78712, Austin, USA

In recent years, there has been much interest in characterizing statistical properties of natural stimuli in order to better understand the design of perceptual systems. A fruitful approach has been to compare the processing of natural stimuli in real perceptual systems with that of ideal observers derived within the framework of Bayesian statistical decision theory. While this form of optimization theory has provided a deeper understanding of the information contained in natural stimuli as well as of the computational principles employed in perceptual systems, it does not directly consider the process of natural selection, which is ultimately responsible for design. Here we propose a formal framework for analysing how the statistics of natural stimuli and the process of natural selection interact to determine the design of perceptual systems. The framework consists of two complementary components. The first is a maximum fitness ideal observer, a standard Bayesian ideal observer with a utility function appropriate for natural selection. The second component is a formal version of natural selection based upon Bayesian statistical decision theory. Maximum fitness ideal observers and Bayesian natural selection are demonstrated in several examples. We suggest that the Bayesian approach is appropriate not only for the study of perceptual systems but also for the study of many other systems in biology.

1. BACKGROUND

The evolution of a species occurs as the result of interactions between its environment and its genome. Thus, achieving a rigorous account of evolution depends as much on understanding the properties of environments as it does on understanding the properties of nucleic acid and protein chemistry. Although many laws of nature are usefully described in deterministic terms, the complexity of evolution often requires a description of both environments and genetics in statistical/probabilistic terms. Research in the areas of population genetics and theoretical ecology has emphasized genetics and its relationship to natural selection. However, recent advances in measuring the statistical regularities of natural environments, especially in the area of perceptual systems, raise the possibility of developing more complete theories of evolution by combining precise statistical descriptions of the environment with precise statistical descriptions of genetics.

Here we show that a constrained form of Bayesian statistical decision theory provides an appropriate framework for exploring the formal link between the statistics of the environment and the evolving genome. The framework consists of two components. One is a Bayesian ideal observer with a utility function appropriate for natural selection. The other is a Bayesian formulation of natural selection that neatly divides natural selection into several factors that are measured individually and then combined to characterize the process as a whole. In the Bayesian

formulation, each allele vector (i.e. each instance of a polymorphism) in each species under consideration is represented by a fundamental equation, which describes how the number of organisms carrying that allele vector at time $t+1$ is related to: (i) the number of organisms carrying that allele vector at time t ; (ii) the prior probability of a state of the environment at time t ; (iii) the likelihood of a stimulus given the state of the environment; (iv) the likelihood of a response given the stimulus; and (v) the birth and death rates given the response and the state of the environment. The process of natural selection is represented by iterating these fundamental equations in parallel over time, while updating the allele vectors using appropriate probability distributions for mutation and sexual recombination.

Our proposal draws upon two important research traditions in sensation and perception: ideal observer theory (De Vries 1943; Rose 1948; Peterson *et al.* 1954; Barlow 1957; Green & Swets 1966) and probabilistic functionalism (Brunswik & Kamiya 1953; Brunswik 1956). After reviewing these two research traditions, we motivate and derive the basic formulae for maximum fitness ideal observers and Bayesian natural selection. We then demonstrate the Bayesian approach by simulating the evolution of camouflage in passive organisms and the evolution of two-receptor sensory systems in active organisms that search for prey. Although we describe Bayesian natural selection in the context of perceptual systems, the approach is quite general and should be appropriate for systems ranging from molecular mechanisms within cells to the behaviour of organisms.

*Author for correspondence (geisler@psy.utexas.edu).

(a) **Ideal observer theory in sensation and perception**

Ideal observer theory uses the concepts of Bayesian statistical decision theory to determine optimal performance in a task, given the physical properties of the stimuli. Organisms generally do not perform optimally in any given task, and thus the aim of ideal observer theory is often not to model the performance of the organism *per se*, but instead to provide a precise measure of the stimulus information available to perform the task, to provide a computational theory of how to perform the task, and to serve as an appropriate benchmark against which to compare the performance of the organism.

To illustrate the logic of ideal observer theory, consider a simple categorization task where there are n possible stimulus categories, c_1, c_2, \dots, c_n , and the observer's task is to correctly identify the category, given a particular stimulus \mathcal{S} arriving at the sensory organ. If there is substantial stimulus noise or overlapping of categories, then the task will be inherently probabilistic because errors are unavoidable. As one might guess intuitively, the ideal classifier achieves optimal performance by computing the probability of each category, given the stimulus, and then choosing the most probable category. In other words, the optimal decision rule is

If $p(c_i|\mathcal{S}) > p(c_j|\mathcal{S})$ for all $j \neq i$, then pick c_i .^{1,2}

These conditional probabilities are often computed by making use of Bayes's theorem:

$$p(c_i|\mathcal{S}) = \frac{p(\mathcal{S}|c_i)p(c_i)}{p(\mathcal{S})},$$

where $p(c_i|\mathcal{S})$ is the posterior probability, $p(\mathcal{S}|c_i)$ is the likelihood, and $p(c_i)$ is the prior probability.³ The probability in the denominator, $p(\mathcal{S})$, is a normalizing 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:

$$p(\mathcal{S}) = \sum_{j=1}^n p(\mathcal{S}|c_j)p(c_j).$$

Using Bayes's theorem, the optimal decision rule becomes

If $p(\mathcal{S}|c_i)p(c_i) > p(\mathcal{S}|c_j)p(c_j)$,
for all $j \neq i$, then pick c_i . (1.1)

In other words, Bayes's theorem implies that one can optimally categorize a given stimulus if one knows the probability of the different categories (the prior probabilities), and the probability of the stimulus given each of the possible categories (the likelihoods).

In perception research, the prior probabilities and the stimulus likelihoods are generally under experimental control and hence can be computed directly from the experimental design and from the irreducible physical noise properties of the stimuli (e.g. the Poisson noise property of light). Applications of ideal observer theory to simple categorization tasks have yielded important discoveries concerning, for example, the detection of signals in noise (De Vries 1943; Rose 1948; Peterson *et al.* 1954; Barlow 1957; Green & Swets 1966; Burgess *et al.* 1981; Watson *et*

al. 1983; Pelli 1990), the discrimination of colour (Vos & Walraven 1972), the discrimination of spatial patterns and spatial location (Geisler 1984, 1989), the discrimination of surface slant using texture cues (Knill 1998*a*), and the discrimination of 3D shape (Liu *et al.* 1995). These discoveries include optimal performance laws for various stimulus dimensions. To mention just a few, it has been shown that the physical limit for resolving differences in intensity increases with the square root of the average intensity, that the physical limit for spatially resolving two features decreases with the fourth root of stimulus intensity, and that the physical limit for detecting changes in the relative location of two spatially separated features decreases with the square root of intensity.

More generally, the ideal observer approach has led to the discovery of computational theories of how best to perform various tasks given the available information. In fact, an ideal observer amounts to the correct formalization of the intuitive notion of a computational theory as described by Marr (1982).⁴

As mentioned earlier, the performance of real organisms generally does not reach that of the ideal observer. Nonetheless, ideal observer theory serves as a useful starting point for developing realistic models of sensation and perception (e.g. Schrater & 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 (Green & Swets 1966; Barlow 1977; Pelli 1990), or known anatomical or physiological factors that would limit performance (Geisler 1989).

If the costs and benefits associated with different stimulus-response outcomes are more complex than maximizing accuracy, then a more general form of Bayesian decision theory is required (e.g. Berger 1985). Specifically, if the goal is to maximize the average utility, then the optimal decision rule becomes

If $u(c_i|\mathcal{S}) > u(c_j|\mathcal{S})$ for all $j \neq i$, then pick c_i

where the average utility associated with picking c_i is given by

$$u(c_i|\mathcal{S}) = \frac{1}{p(\mathcal{S})} \sum_{k=1}^n u(c_i, c_k) p(\mathcal{S}|c_k) p(c_k). \quad (1.2)$$

In this equation, $u(c_i, c_k)$ is the utility function,⁵ which specifies the benefit or loss associated with picking category i when the correct answer is category k . The remaining terms in the equation specify the probability of category k given that the stimulus is \mathcal{S} . (Once again, the value of $p(\mathcal{S})$ can be ignored because it has no effect on the decision.)

Note that if the utility is 1.0 for a correct classification and 0.0 otherwise, then this decision rule reduces to maximizing accuracy. The typical goal for categorization tasks in the laboratory is to maximize accuracy, and hence, in that case, the utility function takes on this simplified form. However, in natural categorization tasks there is often a

complex pattern of costs and benefits associated with different stimulus–response outcomes, and hence the utility function plays a more interesting role.

The utility function also plays a more interesting role in tasks that involve estimating physical properties of the environment. In general, estimates that are nearer the truth have greater utility than those that are wide of the mark. Thus, the proper utility functions tend to be smoothly varying (Yuille & Bülthoff 1996; Brainard & Freeman 1997). Ideal observers have been developed for a number of perceptual estimation tasks, such as estimating the direction and spectra of sound sources (Candy & Xiang 2001), estimating the reflectance of a surface independent of the spectral distribution of the illuminating light source (Brainard & Freeman 1997), and estimating the shape of a surface from shading (Freeman 1996) or the slants of surfaces from texture (Knill 1998*b*). A collection of discussions and applications of the Bayesian approach in perceptual estimation tasks is available in Knill & Richards (1996).

(b) *Probabilistic functionalism and the statistics of natural environments*

It is not surprising that humans fall short of ideal performance in both categorization and estimation tasks. Organisms evolve to perform many different tasks, and hence there may be compromises in design that lead to non-ideal performance in a given task. Also, there are limits to the range of materials that organisms can synthesize and exploit, and limits on the possible structure of organic molecules. Furthermore, perceptual systems are designed through natural selection, and thus the appropriate measure of utility is fitness (birth and death rates), which may lead to ideal observer predictions that are different from those obtained with other utility functions.

These factors will be taken up later, but first consider the obvious factor that the stimuli and tasks used in the laboratory generally do not correspond well with those that occur when an organism is operating in its natural environment. It is conceivable that an organism's performance may more closely approach that of the ideal observer in natural tasks performed with natural stimuli. This is especially true for those specific tasks and stimulus environments where natural selection has had an opportunity to operate over long periods of time or where the selection pressure is strong.

Ever since Darwin, the importance of characterizing natural environments for understanding biological design has been widely recognized. However, Brunswik & Kamiya (1953) and Brunswik (1956) were the first to fully appreciate the value of measuring directly the statistical properties of natural stimuli. For example, Brunswik & Kamiya (1953) measured the distances between parallel contours in visual images and showed that contours that are closer to each other are more likely to belong to the same physical object. Identifying those stimulus features that are likely to belong to the same object is critical for object recognition. Thus, Brunswik and Kamiya provided direct evidence of a selection pressure for perceptual mechanisms that perform grouping on the basis of 'proximity'. The Gestalt psychologists had demonstrated decades earlier (Wertheimer 1958) that the human visual

system tends to group visual features based on relative proximity, but Brunswik and Kamiya showed that this grouping rule has a physical basis in the statistics of natural environments.

In the 1950s, when Brunswik proposed measuring the statistical properties of natural stimuli, the technology required to make such measurements was largely non-existent. In recent decades, with the advent of more powerful computers, there has been growing interest in measuring statistics of natural stimuli. Among recent results for the visual domain are the following: the wavelength spectra of almost all natural light sources can be fitted with a single model having just two or three free parameters (Judd *et al.* 1964; Dixon 1978); the wavelength spectra of most natural surfaces can be described with a similar model having just two or three parameters (Buchsbaum & Gottschalk 1984; Maloney 1986); the distribution of contour orientations in natural scenes peaks in the vertical and horizontal directions (Switkes *et al.* 1978; Coppola *et al.* 1998); and spatial frequency spectra of natural scenes are not uniform, but instead amplitude varies inversely with spatial frequency (Burton & Moorehead 1987; Field 1987).

All these statistical facts about natural environments appear to be reflected in the design of the human visual system. The visual system is limited to three classes of cone photoreceptors, perhaps reflecting the highly constrained spectra of natural sources and surfaces (Maloney & Wandell 1986). The visual system is more sensitive to vertical and horizontal contours than to diagonal contours, perhaps reflecting the natural distribution of contour orientations (Annis & Frost 1973; Timney & Muir 1976). The visual system contains spatial frequency channels whose bandwidths are approximately constant on a logarithmic (octave) scale, perhaps compensating for the decline in amplitude as a function of spatial frequency (Field 1987).

The statistical properties of the environment may influence the design of perceptual systems either in a rigid genetically programmed fashion (a fixed adaptation) that is independent of the specific environment during the organism's lifespan or in a more flexible way (a facultative adaptation) that is dependent on the specific environment during the lifespan. Facultative adaptations include, for example, all mechanisms of learning and neural plasticity. The distinction between fixed and facultative adaptations is illustrated clearly in the example described by Williams (1966) of skin thickness: humans have a fixed adaptation of thicker skin in regions such as the soles of the feet and have a facultative adaptation to increase skin thickness in response to friction (i.e. to form a callus). This example also shows that a particular property of an organism can be dependent upon both fixed and facultative adaptations. With respect to the examples described above, the three classes of photoreceptors are surely a fixed adaptation, whereas the sensitivity to contour orientation and the bandwidths of spatial frequency channels may be at least partly the result of facultative adaptations (e.g. neural plasticity). Later, we return to the distinction between fixed and facultative adaptations. In the meantime, it is important to keep in mind that both kinds of adaptation result from natural selection (Williams 1966).

(c) ***Ideal observers and the statistics of natural environments***

Until recently, ideal observer theory was applied only to tasks involving relatively simple stimuli created in the laboratory. As measurements of statistical properties of natural environments have become available, it has become possible to apply the theory to tasks involving more natural stimuli. This is an important development because the results speak directly to the relationship between the statistics of natural environments and the design of perceptual systems.

The development of information theory (Shannon 1948) inspired the hypothesis that organisms may evolve perceptual systems that encode—or can learn to encode—natural environmental stimuli in an optimally efficient fashion (Attneave 1954; Barlow 1961). An optimally efficient code is one that represents as much of the input (sensory) information as possible, given certain assumed constraints (e.g. a fixed number of neurons with a fixed dynamic range). Coding is a task with well-defined goals, and hence optimal coding theory may be regarded as a particular form of ideal observer theory. The optimally efficient code in any given situation depends strongly on the particular statistical properties (i.e. the relevant probability distributions) of the stimuli being encoded. Thus, it is possible to test the efficient coding hypothesis by measuring response properties of neurons in a perceptual system (e.g. their receptive field properties) and then comparing these properties with the optimal code expected given the measured statistics of the relevant natural stimuli. A number of studies have demonstrated instances where the neural coding of natural stimuli appears to be nearly optimal. The contrast–response functions of certain neurons in the fly’s eye are closely matched to the statistical distribution of contrasts in the fly’s environment (Laughlin 1981). Similarly, spatial receptive fields in the fly’s eye (Srinivasan *et al.* 1982; van Hateren 1992), in the retina (Atick & Redlich 1992), and in the primary visual cortex of cats and monkeys (Bell & Sejnowski 1997; Olshausen & Field 1997; van Hateren & van der Schaaf 1998; Hyvarinen & Hoyer 2000) appear to be well matched to the spatial correlations in natural images. The correspondence between spatial coding in the primary visual cortex and natural image statistics appears to be even closer when one considers the nonlinear contrast response characteristics of cortical neurons (Simoncelli & Schwartz 1999). Finally, the spatio-temporal receptive field properties of neurons in the cat lateral geniculate nucleus (Dong & Atick 1995; Dan *et al.* 1996) and primary visual cortex (van Hateren & Ruderman 1998) appear to be well matched to the spatio-temporal correlational structure of natural images. (For a recent review of optimal coding theory and its applications towards understanding the relationship between natural image statistics and the neural representations in perceptual systems, see Simoncelli & Olshausen (2001).)

In addition to the numerous studies of encoding tasks, there have been several recent studies of categorization tasks. We will describe these examples in more detail because categorization tasks (particularly detection and discrimination tasks) will be emphasized here in illustrating the logic of Bayesian natural selection.

It has long been known that in trichromatic primates (including humans) the spectral sensitivity functions of the M (green) cones and the L (red) cones strongly overlap. This would seem to limit colour discrimination in the longer wavelength region of the spectrum. However, recent studies of natural image statistics have shown that the placement of the M and L spectral sensitivity functions along the wavelength dimension may in fact be nearly optimal for detecting important targets in background foliage. In a thorough study, Regan and colleagues measured the wavelength distributions of primary food sources (fruits) of several New World (platyrrhine) monkeys and the wavelength distributions of the surrounding foliage (Regan *et al.* 1998, 2001). They then used an ideal observer analysis to determine optimal placement of the M and L cones for detecting 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 & Vorobyev (1996) have made a similar case for placement of the cone photopigments in Old World (catarrhine) primates. Similarly, the two cone pigments in dichromatic mammals appear to be nearly optimally placed for discriminating between natural leaf spectra (Lythgoe & Partridge 1989; Chiao *et al.* 2000).

Spatial receptive field properties of neurons in the retina and primary visual cortex are often regarded as being optimized for detecting edges. However, this hypothesis has never been directly tested using the statistics of edges in natural images. Konishi *et al.* (2002) took a significant step in this direction by measuring the empirical probability distributions of theoretical receptive field responses at edges and away from edges in images of scenes containing a mixture of natural and human-made objects. They used existing presegmented image datasets, and thus the actual edge locations were known. For a range of different receptive field shapes, they determined the optimal decision rule (ideal observer) for classifying whether an image pixel is on or off an edge. They found that combining outputs from multiple sizes of receptive field (as found in the mammalian visual system) provides a substantial increase in performance in identifying edges. They also found that receptive field shapes similar to the first derivative of a Gaussian performed better for edge detection than those similar to the popular second derivative of a Gaussian (Marr 1982). It remains to be seen how well actual receptive field shapes found in the retina and primary visual cortex are suited for edge classification.

Several other recent studies have continued Brunswick’s original efforts to analyse the relationship between mechanisms of perceptual grouping and the statistics of natural images. Elder & Goldberg (1998), Geisler *et al.* (2001) and Sigman *et al.* (2001) measured natural image statistics that are relevant for contour grouping ‘good continuation’, and Malik *et al.* (2002) measured natural image statistics that are relevant for region grouping ‘similarity grouping’. Sigman *et al.* (2001) were primarily interested in comparing natural image statistics with long-range lateral connections in the primary visual cortex, and hence they did not measure the natural image statistics required for an ideal observer analysis of contour grouping. Further, Elder & Goldberg (1998), Sigman *et al.* (2001)

and Malik *et al.* (2002) did not make quantitative predictions for performance in categorization tasks. Hence, the Geisler *et al.* (2001) study is the most relevant example in the present context.

Geisler *et al.* (2001) extracted edge elements from images of diverse natural scenes⁶ and then computed co-occurrence probabilities for different possible geometrical relationships between the edge elements. The geometrical relationship between two arbitrary edge elements is described by three variables: distance between the elements (d), direction of one element relative to the other (ϕ), and orientation of one element relative to the other (θ). Two kinds of co-occurrence probabilities were measured—absolute probabilities and conditional (Bayesian) probabilities. The absolute probabilities were obtained simply by counting the relative number of times each possible geometrical relationship occurred in a large number of natural images. The result is a three-dimensional (3D) probability distribution, $p(d, \phi, \theta)$, which was found to be similar from one natural image to the next. This absolute probability distribution has a complex structure, revealing two properties of contours in natural images: nearby contour elements tend to be parallel to each other, and nearby contour elements tend to be co-circular (tangent to the same imaginary circle). It is not implausible that high co-occurrence probabilities are associated with belonging to the same physical contour, and thus these two properties may reflect natural selection pressures for grouping on the basis of ‘orientation similarity’ and ‘good continuation’, respectively. However, the only way to establish this definitively is to measure (as did Brunswik & Kamiya 1953) co-occurrence probability distributions, given that a pair of edge elements belongs to the same physical contour and given that the elements belong to different contours.

Geisler *et al.* (2001) measured these conditional probability distributions for contours using an image tracing technique, where each edge element in each natural image was assigned by hand to a unique physical contour with the aid of specially designed software. With this assignment information, they were then able to determine the probability of each possible geometrical relationship between two edge elements given that they belong to the same physical contour, $p(d, \phi, \theta | c)$, or given that they belong to different physical contours, $p(d, \phi, \theta | \sim c)$. These conditional probability distributions clearly demonstrate the existence of a selection pressure for grouping on the basis of ‘good continuation’, in the sense that nearby edge elements that are approximately co-circular are likely to belong to the same physical contour, and hence should be grouped perceptually if the organism is to represent the physical environment correctly.

These distributions can also be used to determine the optimal decision rule for contour grouping (given only the geometrical relationships between pairs of edge elements). Using equation (1.1), we see that the ideal observer should group elements (i.e. categorize a pair of edge elements as belonging to the same physical contour) according to the following decision rule:

$$\text{If } \frac{p(\mathbf{S} | c)}{p(\mathbf{S} | \sim c)} > \frac{p(\sim c)}{p(c)}, \text{ then group,} \quad (1.3)$$

where \mathbf{S} is the geometrical relationship between the pair

of edge elements, i.e. the particular value of the vector (d, ϕ, θ) . The ratio of the stimulus probabilities on the left is the likelihood ratio, and the ratio of the prior probabilities on the right is a criterion whose value is independent of the geometrical relationship between edge elements.

Using this decision rule and the measured conditional probability distributions for natural images, Geisler *et al.* (2001) generated predictions for contour detection performance and compared those predictions with the ability of humans to detect contours embedded in complex backgrounds. Human contour detection performance was measured in an extensive parametric study, testing a wide range of natural contour shapes. Remarkably, the predictions from the natural image statistics were quite accurate across all conditions, with only one free parameter—the value of the criterion. (The correlation between observed and predicted detection accuracy was approximately 0.9.) This result demonstrates that there is a close relationship between contour grouping mechanisms in humans and the statistics of contours in natural images. Further, the result suggests that natural selection may have created contour grouping mechanisms (probably both fixed and facultative adaptations) that are approximately ideal for detecting natural contours.⁷

(d) *Maximum fitness ideal observers*

Perceptual systems evolve through natural selection, and thus biologically appropriate ideal observers are those where the measure of utility is fitness (birth and death rates). Specifically, natural selection picks genes that maximize the number of organisms carrying those genes. We can represent the fitness utility function as a growth-factor function $\gamma(\mathbf{r}, \omega)$. The greater the growth factor, the greater the increase in the number of organisms carrying a given gene. Obviously, the growth factor is a function of the response \mathbf{r} made by the organism in each particular state of the environment ω . Thus, given a particular stimulus \mathbf{S} , the maximum fitness ideal observer will make the response $\mathbf{r}_{\text{opt}}(\mathbf{S})$ that maximizes the growth factor averaged across all possible states of the environment. In other words, the ideal observer will make the response that maximizes the quantity

$$\gamma(\mathbf{r} | \mathbf{S}) = \sum_{\omega} \gamma(\mathbf{r}, \omega) p(\mathbf{S} | \omega) p(\omega). \quad (1.4)$$

This equation is identical to the standard Bayesian formula (equation (1.2)), except that the utility function is the growth-factor function, and we have dropped the term $p(\mathbf{S})$ because (as mentioned earlier) it has no effect on the optimal response. Note that \mathbf{r} is a vector that can contain any number of discrete or continuous elements, and hence equation (1.4) is appropriate for a wide range of tasks including categorization, estimation, and coding. As in other applications of ideal observer analysis, one can incorporate physiological and anatomical constraints into a maximum fitness ideal observer, and hence determine the optimal responses given those constraints (see Appendix A).

Considering ideal observer theory from the standpoint of maximum fitness leads immediately to several conclusions. First, in order to determine the appropriate ideal observer it is necessary to measure (or know) the growth-

factor function. Second, it is obvious that the ideal observer will vary for different organisms depending upon their particular growth-factor function. In particular, optimal perceptual design will depend upon the organism's nominal birth and death rates, as well as upon the consequences of different responses and states of environment on those rates. Third, natural selection can produce changes in nominal birth and death rates (e.g. via changes in reproductive and ageing mechanisms), and hence the ideal observer may change depending upon the genes (alleles) controlling nominal birth and death rates.

If fitness is the appropriate utility function when considering the design of perceptual systems in biological organisms, then why do there appear to be so many close parallels between actual perceptual performance and ideal observers derived with other utility functions (as reviewed in the previous section)? The most probable answer is that in these cases the assumed utility function is approximately monotonic with the appropriate growth-factor function, and thus a design feature that is optimal for the assumed utility function is also optimal for fitness. This is not to say that the assumed utility function always has little effect on optimal design. Later, we show that utility functions based on fitness can yield ideal observers that are quite different from those obtained with more traditional utility functions, even for relatively simple tasks.

Even though the maximum fitness ideal observer is based upon the utility function defined by natural selection, there are many reasons to expect that natural selection will often fail to achieve ideal performance. Nonetheless, the maximum fitness ideal observer has an important role to play in providing the appropriate computational theory for natural tasks, and in providing the appropriate benchmark against which to evaluate both the performance of the organism and the process of natural selection. For example, the maximum fitness ideal observer allows one to determine how closely natural selection approaches optimal performance. Also, as we will see, the maximum fitness ideal observer is very useful for interpreting and validating simulations of natural selection.

2. BAYESIAN NATURAL SELECTION

As discussed earlier, there is sometimes a close correspondence between the statistics of natural environments and the design of perceptual systems, in that performance of a perceptual system approaches that of an ideal observer informed or limited by the relevant environmental statistics. However, even for natural tasks there must be many situations where a perceptual system falls well short of ideal. Some of the reasons were mentioned earlier: organisms evolve to perform many different tasks leading inevitably to compromises in design that result in non-ideal performance in some tasks; there are limitations on the possible structure of organic molecules; and the assumed utility function may not match the organism's intrinsic utility function (fitness). In addition, even without these factors one would not generally expect to evolve ideal performance because of inherent limits in the process of natural selection. By definition, an ideal observer is obtained by considering the entire space of possible solutions for a task and picking the best according to its utility function.




Figure 1. Natural selection involves an interaction between the environment, the responses of the organism to the environment, and the survival/reproduction rate of the organism. In general, natural selection results in changes in: (i) the responses of the organism; (ii) the survival/reproduction rate of the organism; and (iii) the structure of the environment within which the organism exists. Our definition of the environment includes the organism's internal state.

Natural selection, on the other hand, does not look across the landscape of possibilities and pick the best. Instead, natural selection must always move in small steps, where each step must produce an increase in fitness. Natural selection cannot produce a temporary decrease in fitness in order to later reach a higher point in the landscape of possible solutions. In mathematical jargon, natural selection generally creates a perceptual system that corresponds to a local maximum in the space of possible solutions, not the ideal system that corresponds to the global maximum. The small step size in natural selection also implies that there will be a time-lag in reaching the local maximum.

It is clear then that the design of a perceptual system is limited not only by the task and the relevant statistics of natural stimuli, but by a variety of other factors including inherent limits set by the process of natural selection. How can we assess the contributions of these other factors? There is a growing consensus in the perception community that Bayesian statistical decision theory (through development of ideal observers) provides the appropriate formal framework for understanding how the task and the statistical properties of environments contribute to the design of perceptual systems. Here we propose that Bayesian natural selection, a constrained form of Bayesian statistical decision theory, provides the appropriate formal framework for understanding how all factors contribute to the design of perceptual systems; including the task, relevant statistics of the natural stimuli, compromises among tasks, limits on materials and organic molecules, and the incremental character of natural selection. As we will see, the Bayesian formalization of natural selection may be useful in formulating and testing hypotheses both about the design of existing perceptual systems and about the process of natural selection itself.

To motivate the Bayesian formalization of natural selection, we note first that natural selection involves a complex interaction between the environment, behaviour, and reproduction (see figure 1). At a particular time t , there is some prior probability distribution over the possible states of the environment (ω). Given a particular state of the environment, the organism will respond in some fashion, which may include a passive response such as reflecting light. This response (r) is, in general, probabilis-




Figure 2. 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 alleles at time t . The actual number of organisms at time $t + 1$ is a random quantity that results from sampling appropriate probability density functions. The vector ω represents a particular state of the environment. In general, it is a vector that represents aspects of the environment at time t that are relevant to the evolution of the organism, such as the numbers and kinds of nearby predator and prey species and the background environment. In Bayesian terminology, the probability density function for ω is called the ‘prior probability’. The vector \mathbf{s} represents a particular stimulus arriving at the organism. In Bayesian terminology, the probability density function for \mathbf{s} given ω is called the ‘stimulus likelihood’. The vector \mathbf{r} represents the response of the organism. The response of the organism is probabilistic and its response probability function depends upon the organism’s species, its particular allele vector, and the particular stimulus. In Bayesian terminology, the probability density function for \mathbf{r} given \mathbf{s} is called the ‘response likelihood’. Finally, the growth factor consists of one plus the birth rate minus the death rate, which both depend upon the response, the state of the environment, the organism’s species and its particular allele vector. Each different vector of alleles in each species under consideration is represented by a separate fundamental equation; all of the equations are evaluated and iterated in parallel. The effects of mutation and sexual reproduction are described by other equations (see text).

tic and is determined by a vector of alleles (\mathbf{a}) carried in the organism, although we must keep in mind that the response may reflect both fixed and facultative adaptations. The response of the organism in a particular environment will have consequences for survival and reproduction. If the growth factor (γ)—one plus the birth rate minus the death rate—is greater than one, then the number of organisms containing the set of alleles will on average increase. However, success for a set of alleles (a growth factor greater than one) has an inevitable, and eventually powerful, feedback effect on the environment. Specifically, growth in the number of organisms containing any given set of alleles can go on for only so long. Eventually, equilibrium with the environment must be approached (i.e. the growth factor must converge towards one or become less than one). Thus, the prior probability distribution over the possible states of the environment must change over time. Under some circumstances (discussed later), the prior probability and the growth factor may depend on the allele vector.

This description of natural selection translates directly into the following Bayesian formula:

$$\bar{O}_a(t + 1) = O_a(t) \sum_{\omega} p_a(\omega;t) \sum_r \gamma_a(r,\omega) p_a(r|\omega). \quad (2.1)$$

In this equation $\bar{O}_a(t + 1)$ represents the average

(expected) number of organisms of a given species at time $t + 1$ carrying a particular vector of alleles \mathbf{a} , and $O_a(t)$ represents the actual number of organisms at time t . The rest of the terms on the right of the equation give the total average growth factor for an individual with allele vector \mathbf{a} . This total average is obtained by summing specific growth factors over all the possible states of the environment and over all the possible responses of the organism. Specifically, $p_a(\omega;t)$ is the prior probability distribution over the possible states of the environment parametrized by t , $p_a(\mathbf{r}|\omega)$ is the likelihood distribution over the possible responses given the state of the environment, and $\gamma_a(\mathbf{r},\omega)$ is the growth factor (utility function) associated with each possible response in each possible state of the environment. Equation (2.1) gives the average number of organisms expected at time $t + 1$. The actual number at time $t + 1$ is a random number, $O_a(t + 1)$, that can be represented by sampling from appropriate probability distributions for births, deaths, mutations, and sexual recombination (see later).

To explicitly represent the statistical properties of stimuli reaching the organism, we can expand the last term on the right of equation (2.1) using the definition of conditional probability:

$$\bar{O}_a(t + 1) = O_a(t) \sum_{\omega} p_a(\omega;t) \sum_r \gamma_a(\mathbf{r},\omega) \sum_S p_a(\mathbf{r}|\mathbf{s}) p_a(\mathbf{s}|\omega). \quad (2.2)$$

We call this the fundamental equation of Bayesian natural selection.⁸

The fundamental equation describes natural selection for one particular allele vector in a given species. Thus, the full description of natural selection requires a separate fundamental equation for each allele vector in each species under consideration, with all the equations iterating in parallel over time. At this general level of description, the unit of time is unspecified; depending on the application, it could range from units that are a small fraction of the average lifespan to units that are many lifespans. Also, it is important to note that the relevant factors (components) defining the environment vector (ω), stimulus vector (\mathbf{s}), and response vector (\mathbf{r}) will generally be different for each species.

In the Bayesian framework, fitness is the value of the expression to the right of $O_a(t)$ in equation (2.2), that is, the growth factor averaged over all possible states of the environment, all possible stimuli, and all possible responses. This definition of fitness turns out to be consistent with Hamilton’s concept of inclusive fitness (Hamilton 1964a,b). Specifically, the fundamental equation describes the number of organisms carrying a given allele vector irrespective of lines of descent. For example, the fitness can include the birth and death rates of collateral as well as lineal descendants.

Given that fitness is defined with respect to the allele vector, it follows that the possible states of the environment can include not only the environment external to the organism carrying the allele vector, but also the environment internal to the organism (e.g. properties of the organism that depend on age). A related point is that the logic of Bayesian natural selection does not require that ‘stimulus’ and ‘response’ be interpreted only in their usual psychological sense. Rather, they can be interpreted

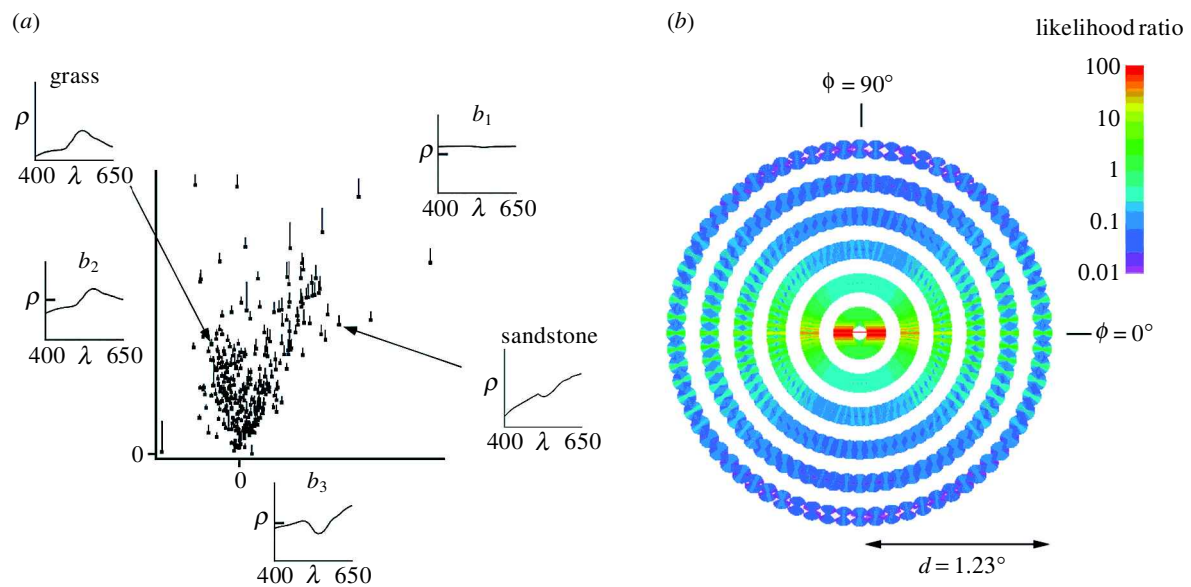


Figure 3. Examples of the complex statistical structure of natural stimuli. (a) Spectral reflectance distributions of natural surfaces measured by Krinov (1947). The Krinov spectra are adequately described by the weighted sum of the three basis functions labelled b_1 , b_2 and b_3 (obtained by a principal components analysis). Each dot represents a particular spectral distribution; the location of a dot indicates the weights on the second and third basis functions, and the length of the line segment attached to a dot indicates the weight on the first basis function. Also shown are the spectral distributions for grass and red sandstone; their representation in terms of the basis functions is given by the dots located at the arrowheads. (b) Likelihoods of different geometrical relationships between contour elements in natural images. The central line segment represents an arbitrary contour element and each of the other 7776 line segments represents a possible geometrical relationship with that element. The colour of a line segment indicates the likelihood ratio: the probability of observing the given geometrical relationship when the two elements belong to the same physical contour divided by the probability of observing the geometrical relationship when they belong to different physical contours. These data show that contour elements that have a co-circular relationship are more likely to belong to the same physical contour. (Adapted from Geisler *et al.* (2001)).

broadly to stand for the input and output of any biological system or subsystem, from the molecular to the behavioural level.

The value of the Bayesian formulation of natural selection is that it neatly divides natural selection into its key components, which are identified in figure 2. However, it is important to emphasize at the outset that each component of the fundamental equation is highly complex, easily being a whole science unto itself. Further, the equation does not constitute a theory or hypothesis, but is simply a general formalization of the principle of natural selection, which we assume to be true. Theories or hypotheses are represented by assumptions concerning the components of the fundamental equation(s). Failures of prediction are interpreted as rejection of these assumptions. Later, we demonstrate with a few simple examples how the Bayesian formulation of natural selection might be used to formulate and test specific hypotheses. We now consider the components of the fundamental equation one at a time.

(a) *Prior probability*

The prior probability distribution specifies the probability of the possible states of the environment relevant for the evolution of the species under consideration. The state of the environment is described by a vector, $\omega = (\omega_1, \omega_2, \dots)$, which could represent any relevant environmental factors either external or internal to the organism. The environment vector could represent categorical states of the external environment such as the sub-

population of conspecifics (deme), whether a predator or prey is nearby, whether the nearby predator or prey is moving, or the kind of background foliage or background sound. It could also represent more continuous states of the external environment such as the geographically dependent characteristics of conspecifics (cline), the 3D location of a nearby predator or prey, level of local illumination, direction of the primary light source, temperature, or intensity and direction of the wind. Similarly, the environment vector could represent internal states such as the sex or age of the organism.

Natural selection must ultimately produce substantial changes in the prior probability distribution, and thus the prior probability, $p_a(\omega; t)$, must be dependent on time, as indicated by the parameter t . However, many (perhaps most) relevant environmental factors will not be affected by the evolution of the species under consideration. Thus, in practice it may be possible to simplify the Bayesian analysis by separating the prior probability distribution into those factors that vary during the course of evolution (ω_r) and those that do not (ω_v):

$$p_a(\omega; t) = p_a(\omega_r | \omega_v; t) p(\omega_v). \quad (2.3)$$

Factors likely to vary during the course of evolution might be, for example, spatial distributions of the species' predators and prey, camouflage of the species' predators and prey, or the age distribution of organisms in the species; factors less likely to vary on an evolutionary time-scale might be distributions of background foliage, background

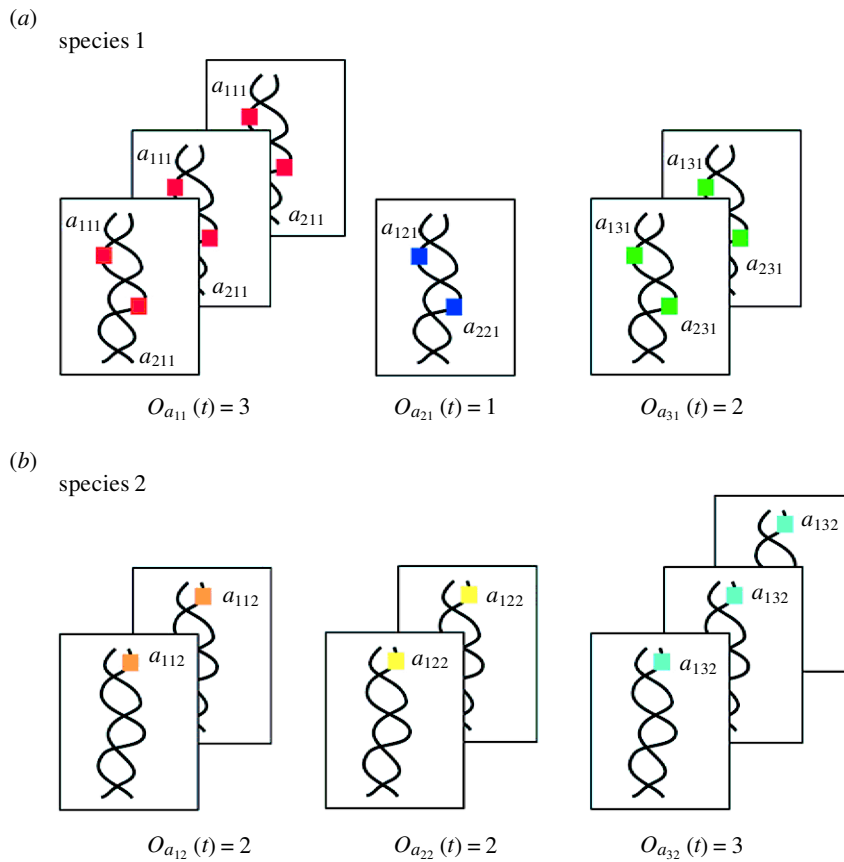


Figure 4. Illustration of how alleles, species, and organisms are represented (in haploid species). The alleles carried by an organism are represented as a vector \mathbf{a}_{jk} , where k indexes the species and j indexes the particular allele vector in that species. (a) In species 1 there are three different allele vectors for the same two gene locations, as indicated by the differently coloured pairs of boxes (the vector is of length 2). (b) In species 2 there are three different allele vectors for the same single gene location (the vector is of length 1). The total number of organisms carrying allele vector \mathbf{a}_{jk} at time t is given by $O_{\mathbf{a}_{jk}}(t)$. Note that each box represents a separate organism.

sound, temperature, or light level. Of course, there may also be long-term variations that have nothing to do with the evolution of the species under consideration. For example, some aspect of the auditory or visual background could change over time owing to the evolution of some other species that has no direct interaction with the species under consideration, or there could be infrequent or gradual changes in climate such as an extended drought or an ice age.

Dynamic change in the prior probability distribution due to feedback between behaviour and environment is an important difference between Bayesian natural selection and other applications of Bayesian statistical decision theory in perception, which generally assume that there is no such feedback.

Identifying relevant environmental factors and measuring their prior probability distributions is a difficult task, generally requiring a combination of extensive physical measurements, computational analyses, and hypothesis testing. Nonetheless, as the numerous studies measuring natural scene statistics have demonstrated, it is a task on which genuine progress is being made.

(b) Stimulus likelihood

The stimulus likelihood, $p_a(\mathbf{s}|\boldsymbol{\omega})$, specifies the probability of each possible stimulus arriving at the organism, given the current state of the environment. The stimulus,

like the state of the environment, is described by a vector, $\mathbf{s} = (s_1, s_2, \dots)$, which can represent any sort of sensory information. The possible descriptions of stimuli are more constrained than the possible descriptions of the state of the environment: for visual stimuli, the most general description is light intensity as a function of space, time, and wavelength; for auditory stimuli, it is sound pressure as a function of time and frequency; for olfactory stimuli, it is concentration as a function of time and type of molecule; and so on.

Nonetheless, stimulus likelihood distributions in the natural world often have a complex structure. For example, if the states of the environment ($\boldsymbol{\omega}$) are the presence or absence of a particular kind of material (e.g. a particular kind of grass), there will still be a great deal of variation in the stimulus (\mathbf{s}) reaching the organism owing to variation in the spectral reflectance functions of the target and background materials, the spectral illumination function, and the lighting geometry. Some of this variation is illustrated in figure 3a, which shows the spectral reflectance functions of a large sample of materials measured by Krinov (1947). Maloney (1986) showed that the Krinov spectra (which are representative of natural environments) can be described accurately by the weighted sum of three basis functions that were derived via principal components analysis. The three basis functions are shown in the inset plots labelled, b_1 , b_2 and b_3 . The main figure shows the

weights that describe the reflectance function of each material; the location of a datapoint indicates the weights on b_2 and b_3 , and the length of the line segment attached to a datapoint indicates the weight on b_1 . The remaining two inset plots show the weighted sums (i.e. the reflectance functions) for a sample of grass and for a sample of red sandstone. As can be seen, the distribution of natural reflectance spectra has a complex V-like structure, with a higher concentration near the base of the V, and a tendency for greater weight on b_1 in the upper right of the plot.

Another example of the complex structure of stimulus likelihood distributions is the edge co-occurrence measurements of Geisler *et al.* (2001), described in § 1c. In that example, the two states of the environment (ω) are whether or not a given pair of edge elements belongs to the same physical contour, and the stimulus vector (\mathbf{s}) is the geometrical relationship between the elements. Figure 3b plots the ratio of the stimulus likelihoods for the two states of the environment, for all possible geometrical relationships between contour elements. The centre horizontal line segment in the plot represents an arbitrary contour element (taken as the reference) and every other line segment represents a second contour element in one of the possible geometrical relationships with that reference. The six concentric circles represent increasing distance between the element and the reference, the location of an element around the circle represents the direction of the element from the reference, and the orientation of an element at a given location indicates the orientation of the element with respect to the reference. The colour of an element indicates the likelihood ratio. Note that for each distance and direction there are 36 orientations represented (one every 5°); hence, it is difficult to resolve individual elements except in cases where the likelihood ratio is high. As can be seen, edge elements that are co-circular (consistent with a smooth continuous contour) are more likely to belong to the same physical contour.

Analysis of stimulus likelihood has long been a major focus in perception research, and thus techniques for measuring and computing stimulus likelihood are relatively well developed. As we have seen, stimulus likelihood distributions can be measured directly from stimuli captured in the environment (e.g. from natural images). It is also possible to combine more limited direct measurements with known facts of physics, such as the laws of sound propagation and interference, the geometry of perspective projection, the laws of light reflectance, transmittance and refraction, the laws of diffusion, the Poisson noise characteristics of light, and the Brownian noise characteristics of sound. This is a common approach in the laboratory and in modelling, and it could be extended to the measurement of stimulus likelihood in natural environments.

At the level of abstraction represented by the fundamental equation, few constraints are placed on what might constitute a stimulus. The stimulus could refer to events that occur instantaneously or events that occur over some period of time.

It is important to note that there is considerable flexibility in how the environmental statistics are split into prior probabilities and stimulus likelihoods. If the states of the environment are defined in terms of a few general categories or variables, then the prior probability distributions are relatively simple and the stimulus likelihood

distributions are relatively complex; just the reverse holds if the states of the environment are defined in terms of many categories or variables. The same knowledge of the world can be represented either way, and thus the choice is largely pragmatic and part of the art of Bayesian analysis.

(c) Alleles

The alleles carried by an organism are represented by a list or vector $\mathbf{a} = (a_1, a_2, \dots)$. This vector could represent a single allele or any set of alleles up to and including the entire genome. Across organisms within a given species, there are generally variations (i.e. polymorphisms) in the particular allele vector for the same given set of genes. To indicate the particular allele vector and the particular species we can add subscripts j and k , respectively. Thus, the general notation for a set of alleles is $\mathbf{a}_{jk} = (a_{1jk}, a_{2jk}, \dots)$. The total number of organisms at time t carrying the set of alleles is given by $O_{\mathbf{a}_{jk}}(t)$ (see figure 4). The total number of organisms of a given species is the sum across all allele vectors:

$$O_k(t) = \sum_{j=1}^{n_k} O_{\mathbf{a}_{jk}}(t), \quad (2.4)$$

where n_k is the number of allele vectors for species k .

In application, a separate fundamental equation is set up for each relevant allele vector within each species under consideration (i.e. one equation for each \mathbf{a}_{jk}), and all of these equations are evaluated and iterated in parallel. To evaluate the equations it is necessary to specify the initial set of genes, the starting allele vectors for those genes, and the initial number of organisms carrying each allele vector. Also note that when $O_{\mathbf{a}_{jk}}(t)$ becomes zero the allele vector \mathbf{a}_{jk} is extinct and when $O_k(t)$ becomes zero then species k is extinct.

(d) Response likelihood

The response likelihood, $p_{\mathbf{a}}(\mathbf{r}|\mathbf{s})$, specifies the probability of each possible response of the organism given the organism's set of alleles and the stimuli arriving at the organism. An organism's contact with the external environment is entirely through proximal stimuli, and so the response does not depend directly on the value of the environment vector ω .⁹

The response is described by a list or vector, $\mathbf{r} = (r_1, r_2, \dots)$, which could represent either fixed or facultative responses. The number of possible responses is of course enormous. The simplest fixed responses include the passive responses of the surface of the organism, such as the spatio-chromatic distribution of light reflected from the organism, the sound pattern reflected from the organism, and the types and concentrations of molecules passively released from the organism into the surrounding medium. For example, the stimulus vector might be the spatio-chromatic distribution of light falling on the organism, and the response vector might be the spatio-chromatic distribution of light reflected from the organism. Passive surface responses often determine the signature stimulus information made available to other organisms (e.g. they determine the level of camouflage or attractiveness). Other relatively simple fixed responses might be those of the sensory systems, including the optical responses of the eye, the acoustical and mechanical

responses of the ear, the turbulence responses of the nasal cavity, the responses of the sensory receptors, and the responses of various sensory neural circuits. Surface and sensory responses are often the simplest to study because the mechanisms are more accessible and hence easier to measure and analyse. A few other types of response (which may be either fixed or facultative, depending on the specific response and species) might include stimulus-dependent changes in the organism's surface properties, sensory organs, metabolism, hormone levels, neural representations, neural algorithms, and locomotor activities.

As with the stimuli, few constraints are placed on what might constitute a response. The response could refer to events that occur instantaneously, or even events that occur over the lifetime of the organism. The only events excluded are evolutionary changes in the species due to natural selection.

(e) Growth factor

The growth factor is one plus the average birth rate minus the average death rate:

$$\gamma_a(\mathbf{r}, \boldsymbol{\omega}) = 1 + \beta_a(\mathbf{r}, \boldsymbol{\omega}) - \chi_a(\mathbf{r}, \boldsymbol{\omega}). \tag{2.5}$$

The birth and death rate each depend on the response (\mathbf{r}) made within the existing state of the environment ($\boldsymbol{\omega}$) and possibly on the set of alleles under consideration (\mathbf{a}). We allow dependence on the alleles because the alleles could affect the internal reproductive and ageing mechanisms.

Population geneticists often use a growth index based on age-specific birth and death rates. On the one hand, such an index can readily be incorporated into equation (2.5) by allowing the state of the environment ($\boldsymbol{\omega}$) to include the age distribution of the population. On the other hand, if the individual life cycles run largely asynchronously (which will often be the case), then the accuracy of the fundamental equation is little improved by this added complexity. Specifically, given a large population of organisms carrying a given allele vector, the average growth factor (over the scale of a lifetime) will be constant even if reproductive efficiency and mortality vary over the lifespan. In the simulations described later, we assume a simple average birth and death rate per unit time.

The growth factor corresponds to the utility function in Bayesian statistical decision theory. In previous applications of Bayesian statistical decision theory to natural tasks, there has been considerable uncertainty about what form the utility function should take. An important feature of Bayesian natural selection is that there can be no doubt that the growth factor is the proper utility function. Of course, to evaluate the fundamental equation it is necessary to specify the initial conditions for the growth factor, but this specification is less open-ended than in most other applications of Bayesian statistical decision theory to natural tasks.

As mentioned earlier, the actual number of organisms at time $t + 1$ is a random number, which can be obtained by sampling from appropriate probability distributions. To describe this probabilistic process, we first substitute equation (2.5) into the fundamental equation and obtain

$$\bar{O}_a(t + 1) = O_a(t) + \bar{B}_a(t + 1) - \bar{D}_a(t + 1), \tag{2.6}$$

where $\bar{B}_a(t + 1)$ is the average number of births and

$\bar{D}_a(t + 1)$ is the average number of deaths in the t th time-step:

$$\bar{B}_a(t + 1) = O_a(t) \sum_{\boldsymbol{\omega}} p_a(\boldsymbol{\omega}, t) \sum_{\mathbf{r}} \beta_a(\mathbf{r}, \boldsymbol{\omega}) \sum_{\mathbf{s}} p_a(\mathbf{r} | \mathbf{s}) p_a(\mathbf{s} | \boldsymbol{\omega}), \tag{2.7}$$

$$\bar{D}_a(t + 1) = O_a(t) \sum_{\boldsymbol{\omega}} p_a(\boldsymbol{\omega}, t) \sum_{\mathbf{r}} \chi_a(\mathbf{r}, \boldsymbol{\omega}) \sum_{\mathbf{s}} p_a(\mathbf{r} | \mathbf{s}) p_a(\mathbf{s} | \boldsymbol{\omega}). \tag{2.8}$$

A plausible hypothesis is that the actual number of births, $B_a(t + 1)$, is a random sample from a Poisson distribution with a mean of $\bar{B}_a(t + 1)$, and that the actual number of deaths, $D_a(t + 1)$, is a random sample from a binomial probability distribution with parameters $\bar{D}_a(t + 1)/O_a(t)$ and $O_a(t)$. There are other possibilities, but whatever the sampling distributions, the number of organisms at time $t + 1$ is given by

$$O_a(t + 1) = O_a(t) + B_a(t + 1) - D_a(t + 1). \tag{2.9}$$

Note also that when the value of $O_a(t + 1)$ becomes zero, the allele vector is extinct and its fundamental equation disappears.

(f) Mutation and sexual reproduction

In order for natural selection to proceed there must be mechanisms for changing, rearranging, activating, deactivating or exchanging alleles. Mutation and sexual reproduction are the primary mechanisms.

Relevant mutations are those affecting the reproductive cells, and hence the average number of mutations is proportional to the number of births:

$$\bar{M}_a(t + 1) = m_a B_a(t + 1), \tag{2.10}$$

where m_a is a proportionality constant (the mutation rate) that may depend upon the particular allele vector. The actual number of mutations, $M_a(t + 1)$, is plausibly modelled as a random sample from a binomial probability distribution with parameters m_a and $B_a(t + 1)$. For each individual mutation there will be a new allele vector \mathbf{a}' created. The structure of this new allele vector will also be probabilistic and is described by another sampling distribution $p_a(\mathbf{a}')$. In general, the different possible mutations are not equally likely. The most likely would be a single nucleotide mutation in just one of the alleles. Furthermore, the mutation probability will probably vary across alleles within the allele vector and across sites within a given allele. For each new allele vector we subtract one from $O_a(t + 1)$ and add one to $O_{a'}(t + 1)$. Obviously, when a new allele vector is created, it is represented by a new fundamental equation. To describe the effects of mutation, one must specify the mutation rates and the allele sampling distributions.

To characterize sexual recombination it is necessary to explicitly represent the diploid structure of the genome. Thus, each allele vector is a list of pairs of alleles, one pair for each gene location being considered. Like mutation, sexual recombination is also directly linked to the number of births. Half of the alleles from each parent are combined in the offspring. Different more or less sophisticated descriptions of this process are possible. One simple version is to select, for each birth associated with allele vector \mathbf{a} , an allele vector $\hat{\mathbf{a}}$ for the mate, where the allele vector

for the mate is randomly sampled from a probability distribution $p_a(\hat{a}|\omega)$ that may depend on the allele vector \mathbf{a} and on the current state of the environment (e.g. the deme to which the organism belongs). The allele vector of an offspring \mathbf{a}' is then obtained by randomly selecting for each gene location one allele from \mathbf{a} and one from $\hat{\mathbf{a}}$.

(g) *Initial conditions*

In order to evaluate the equations of Bayesian natural selection, one must specify the starting points (initial conditions) for all the terms in the equations. There is no general prescription for doing this.

A given species is likely to be near equilibrium with its environment, and thus the current environment, described by the prior probabilities and stimulus likelihoods, cannot be the same as when the species began to diverge from its ancestors. One tactic for setting the initial conditions of the environment is to identify through measurement, experiment and common sense the factors likely to have changed during recent natural selection for the given species (i.e. the factors in ω_r of equation (2.3)). The prior probability and stimulus likelihood distributions for the other factors (i.e. the factors in ω_s) could be set to their current values, and then the effect of different distributions for the dynamic factors could be explored in simulations. The results could provide evidence concerning the probable starting states, possible ending states, and the adequacy of the assumed static and dynamic factors. Another tactic might be to start the simulations with the current environment and generate predictions for the future. This could provide evidence concerning whether the species is in a nearly optimal design state or, depending on the range of mutations and allele interactions allowed, whether evolution of the species is trapped at a local maximum. The tactic of predicting the future might be particularly useful for species that evolve rapidly and for species where the environment, mutation rate and reproduction can be manipulated.

It is even more certain that the response likelihood (the stimulus–response term) for a given species is different from when the species began to diverge from its ancestors. Much of perception research has been directed at measuring the anatomy, physiology and biochemistry of the stimulus–response mechanisms within organisms, as well as the stimulus–response behaviour of organisms as a whole. In addition, there is information about the evolution of perceptual systems from comparative studies at the biochemical, physiological–anatomical and behavioural levels. All of this information can be brought to bear in specifying plausible initial conditions for the response likelihood distribution. One tactic is to use current knowledge to form a working model of the perceptual system under consideration. The initial conditions for the response likelihood might then be created from this model by deleting some relevant feature or mechanism. Simulations from such a starting point should provide insight into the evolution of the deleted mechanism and its relationship to the statistics of the environment. For example, one might start from the assumption that the earliest Old World primates were dichromatic (like other mammals). Starting with an appropriate characterization of the primate retina, the relevant statistics of the natural environment, the possible photopigment alleles, and the effect of

successful foraging and mating on the growth factor, one could, in principle, simulate the evolution of trichromatic vision. Another tactic would be to start with the working model and then generate predictions for the future.

This completes a general description of the Bayesian framework. We now demonstrate with some examples how the concepts of Bayesian natural selection might be used to formulate and test specific hypotheses via simulation.

3. SIMULATION OF BAYESIAN NATURAL SELECTION

We have seen how the Bayesian framework divides natural selection into its key components. Thus, the framework offers a potential simplification because each of the components can be measured and characterized individually and then combined, via the fundamental equation, to describe the evolutionary process as a whole. However, as noted earlier, each component is highly complex, easily being a whole science unto itself. In the face of this complexity, how does one begin translating the Bayesian formalization of natural selection into concrete testable models of the evolutionary process in particular situations? Formulating testable models will be easiest for simple organisms living in simple environments, but even then a number of detailed technical issues must be considered for each term of the fundamental equation(s).

Here we describe simulations for three common evolutionary scenarios: (i) evolution of a species whose initial polymorphism is transient; (ii) coevolution of two species whose initial polymorphisms are transient; and (iii) evolution of a species that maintains a stable polymorphism. In each case, there are at least two interacting species living in a hypothetical environment (see figure 5). One of the species is actively searching for one or more other species, which are its only food source. The other species are assumed to be passive, like plants. Although the responses of the active species are simple, its perceptual system is reasonably complex and the search task fairly realistic. For convenience we assumed that both species reproduce asexually. These are admittedly simple cases, but they suffice to illustrate the flexibility of the Bayesian approach.

The present simulations are generic, in the sense that we deliberately avoid making detailed assumptions about the stimulus dimensions, stimulus likelihoods and response likelihoods. As it turns out, there are some interesting general conclusions that can be drawn from these generic simulations. Nonetheless, our real aim here is not to formulate or test hypotheses concerning evolution but to demonstrate the Bayesian approach, and to describe three simulation templates into which specific measurements or assumptions could be substituted in order to test specific hypotheses.

Before proceeding, a few words should be said about the value and feasibility of simulation in the study of evolution. Its value arises largely because evolution is so complex. In any natural environment there are a vast number of factors that could affect the course of evolution. It will never be possible to measure all of these factors, and even if they could all be measured, their simultaneous effect on evolution could never be evaluated experimentally. Simul-




Figure 5. Hypothetical environment and interacting species for the simulations of Bayesian natural selection. In the simulations of transient polymorphism and coevolution, there were two species, an active predator species with a sensory system (species 1) and a passive prey species with some surface characteristic (species 2). In the simulation of stable polymorphism, there was a second prey species with a different surface characteristic (species 3). The surface characteristic of the background is not illustrated in this diagram. The maximum number of organisms for a species was defined by the minimum space it requires for reproduction (the dashed circles). Species 2 and 3 were assumed to compete for the same space. Species 1 was assumed to have a sensory system with a limited range indicated by the solid circle; the probability of detecting a prey outside of this range was zero.

ations can be used to determine which factors are likely to be important in a particular situation, and hence where one should concentrate effort in making further physical measurements and in conducting experiments. Once some of the important factors have been identified, simulations are necessary to understand how they work together. For example, simulations can be used to identify those factors that are likely to influence evolution relatively independently and those likely to influence evolution in a highly interactive fashion. Simulations can also be used to explore conditions under which natural selection produces optimal or nearly optimal solutions (in the ideal observer sense). Finally, for well-controlled laboratory environments, simulations can generate quantitative predictions for the course of evolution, thereby providing a basis for rigorous hypothesis testing.

Obviously, a Bayesian model will be useful only if it is computationally tractable. There are reasons to be cautiously optimistic. Bayesian models have been successfully developed and evaluated in a number of complex domains including perception, signal processing, computational neuroscience, and economics. Many of the mathematical techniques and theorems used to make these models trac-

table should transfer to the domain of Bayesian natural selection. There are some aspects of Bayesian natural selection that may appear at first glance to raise computational difficulties. One is that a separate fundamental equation must be evaluated for each allele vector. However, the total number of separate fundamental equations that must be evaluated at each time-step of the simulation tends to be relatively small because mutation rates are limited and unfit alleles soon become extinct. Also, adding new allele vectors has only a linear effect on computation time. Another apparent complicating factor is the potential length of allele vectors. In practice, however, hypotheses will tend to be focused on those limited subsets of genes assumed to be most relevant for the adaptation under study.

(At this point, readers not interested in the details of the simulations may wish to skip to the simulation results (§ 4).)

(a) *Transient polymorphism*

In this example, we consider evolutionary transitions in which multiple allele vectors in a population are eventually replaced by a single allele vector or a cluster of nearly equivalent allele vectors. This situation might arise, for example, when a polymorphic species moves into a new environment that favours just one allele vector. In this first example, we assume that the environment contains just two interacting species, where species 2 (prey) is a passive organism such as a plant, and species 1 (predator) is an active organism that moves around the environment in search of species 2 (see figure 5). We assume that species 2 is not evolving or is doing so on a time-scale that is large relative to species 1. We further assume that there is no significant polymorphism in species 2; thus we can let $\mathbf{a}_{12} = (a_{12})$ represent a single allele vector (of length one) that determines the stimulus signature of species 2 relevant for detection by species 1. In this case there is just one fundamental equation for species 2.

Species 1 is assumed to be evolving rapidly relative to species 2 and to have some initial polymorphism in those genes that are relevant for detecting and responding to species 1. This polymorphism is represented by a collection of different allele vectors, $\mathbf{a}_{11}, \mathbf{a}_{21}, \mathbf{a}_{31}, \dots$. There is a separate fundamental equation for each of these allele vectors.

(i) *Prior probability*

To specify the prior probability distributions relevant for a species, one must first decide how to represent the possible states of the environment. We suppose here that two general states of the environment are important for the survival and reproduction of species 2: (i) whether species 2 is within range of possible detection by an individual of species 1 and (ii) whether species 2 has the opportunity to reproduce. Thus, the environment vector for species 2 is in two dimensions (2D), $\omega_2 = (\omega_{12}, \omega_{22})$. The first dimension has $n_2 + 1$ possible values: not within detection range of species 1 ($\omega_{12} = 0$), within detection range of an individual of species 1 having the allele vector \mathbf{a}_{j1} ($\omega_{12} = j$). The second dimension has two possible values: having space to reproduce ($\omega_{22} = 1$), or not having space to reproduce ($\omega_{22} = 0$). Similarly, we suppose that the same two general states of the environment are

important for the survival and reproduction of species 1: (i) whether species 1 is within range to possibly detect species 2 and (ii) whether species 1 has the opportunity to reproduce. Thus, the environmental vector for species 1 is also in 2D, $\omega_1 = (\omega_{11}, \omega_{21})$. For both species we assume that the prior probabilities do not depend on the value of the allele vector.

Assuming that the probability of being within detection range is independent of having space to reproduce, the prior probability distribution for the different states of the environment is the product of the probability distributions for the two components:

$$p_a(\omega_1; t) = p(\omega_{11}; t)p(\omega_{21}; t),$$

$$p_a(\omega_2; t) = p(\omega_{12}; t)p(\omega_{22}; t).$$

First, consider the probability of having space to reproduce. Even under the best of circumstances, a species can only reach a certain maximum density that is allowed by the available space and properties of the species. In figure 5, the dotted circle around each species indicates the minimum amount of space (the reproduction space) required for each individual of that species. Dividing each of these areas into the total available area determines the maximum number of organisms for each species, o_{\max_1} and o_{\max_2} . We assume that the probability of not having a space to reproduce is equal to the fraction of the total space occupied at time t , which is the fraction of the maximum possible number of organisms:

$$p(\omega_{21} = 0; t) = \frac{O_1(t)}{o_{\max_1}}, \quad (3.1)$$

$$p(\omega_{22} = 0; t) = \frac{O_2(t)}{o_{\max_2}}. \quad (3.2)$$

The probability of having space to reproduce is just one minus the probability of not having space to reproduce. Notice that the prior probabilities change over time in direct proportion to the population of the species.

Now consider the probability of being within detection range. The perceptual system of a species collects information over a limited range. The solid circle in figure 5 indicates this range for the perceptual system of species 1. If species 2 were to reach its maximum density, then the probability of a given individual of species 1 being within range to detect any individual of species 2 would reach its maximum. If we let this maximum probability be p_{range_2} , then the probability of an individual of species 1 being within range to detect an individual of species 2 at time t is

$$p(\omega_{11} = 1; t) = p_{\text{range}_2} \frac{O_{a_{112}}(t)}{o_{\max_2}}, \quad (3.3)$$

and the probability of not being within range, $p(\omega_{11} = 0; t)$, is one minus the probability of being within range.

Conversely, if species 1 were to reach its maximum density, then the probability that a given individual of species 2 is within range of possible detection by any individual of species 1 would also reach its maximum. If we let this maximum probability be p_{range_1} , then the probability of an individual of species 2 being within range of an individual of species 1 with allele vector a_{j1} at time t is

$$p(\omega_{12} = j; t) = p_{\text{range}_1} \frac{O_{a_{j1}}(t)}{o_{\max_1}}, \quad (3.4)$$

and the probability of not being in range, $p(\omega_{12} = 0; t)$, is one minus the sum of all the probabilities of being within range.

This completes the specification of the prior probability distributions. However, there are a couple of final points to make. First, these formulae implicitly assume that the probability of two or more individuals of species 1 being within range of the same individual of species 2 is negligible. This assumption holds if p_{range_1} is sufficiently small. Second, it may appear from equations (3.1)–(3.4) that the prior probability distribution is defined by four parameters: o_{\max_1} , o_{\max_2} , p_{range_1} and p_{range_2} . However, there are in fact only three independent parameters because of a constraint that follows from consideration of figure 5; thus, p_{range_2} is completely determined from the other three parameters.

(ii) Stimulus likelihood

To specify the stimulus likelihood one must identify the relevant stimulus dimensions. Here we assume that the stimulus dimensions for the two species consist of energy distributions (see left panels in figures 6 and 7). The stimulus might, for example, be the wavelength spectrum of the light reaching the organism. Because of variation in the environment there will be variation in the energy distribution; three examples are shown on the left in figures 6 and 7. For the present generic simulations, we do not make particular assumptions about the form of these stimulus probability distributions. However, in general, the precise form of these distributions will be critical in specific applications.

(iii) Response likelihood

To specify the response likelihood one must identify the relevant properties of the organism and how they depend upon the stimuli and the allele vector. For species 2, the relevant property is the characteristic of the organism's surface, for example its spectral reflectance function. As illustrated in figure 6, this surface characteristic is determined by the allele a_{112} . The response from the surface is a joint function (e.g. the product) of the stimulus energy distribution and the surface characteristic. Because of variation in the stimulus (and possibly in the surface characteristic), there will be individual variation in the surface response. Again, for the present generic simulations, we do not make particular assumptions about the form of the surface response characteristic of species 2, although the precise form will be critical in specific applications.

For species 1 the relevant properties of the organism are the characteristics of the receptors, the form of the receptor response integration, and the form of the behavioural decision process. We assume that there are two receptors, a fixed receptor a , and a variable receptor b . The sensitivity function of the fixed receptor is shown as the dashed curve in the middle panel of figure 7, and the variable receptor as the solid curve. The sensitivity function for receptor b is controlled by the first allele of an allele vector $a_{j1} = (a_{1j1}, a_{2j1})$. The activation level of each receptor is a scalar quantity, such as the integral of the product of the stimulus energy distribution or the receptor sensitivity




Figure 6. Hypothetical stimuli, surface characteristic, and surface responses for the passive prey species (species 2). The stimuli for species 2 are energy distributions along some arbitrary wavelength dimension; the possible energy distributions and their probabilities are described by a stimulus likelihood distribution $p(s_2|\omega_2)$. The surface characteristic describes how the surface of the organism responds to the stimuli; although it is determined by some allele a_{112} , there may still be random variation from one individual to another. The surface responses are also energy distributions that are a joint function of the stimulus and the surface characteristic (e.g. the product). The possible surface responses and their probabilities are described by a response likelihood distribution $p_{a_{112}}(r_2|s_2)$, which depends on the stimulus and the allele.

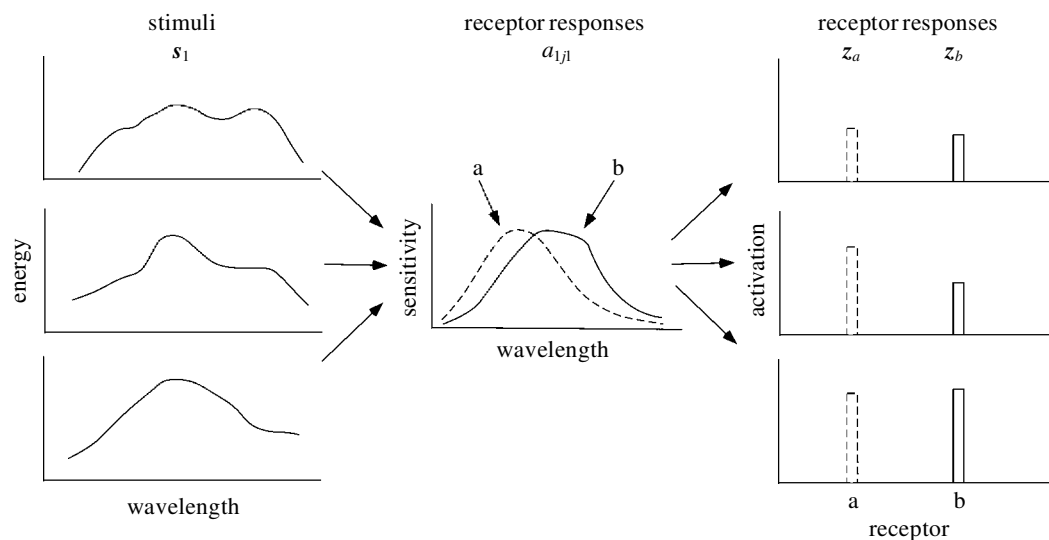


Figure 7. Hypothetical stimuli, receptor sensitivity functions, and receptor activations for the active predator species (species 1). The stimuli for species 1 are energy distributions along some arbitrary wavelength dimension; the possible energy distributions and their probabilities are described by a stimulus likelihood distribution $p(s_1|\omega_1)$. There are two receptors, a fixed receptor a that does not evolve (dashed curve), and a variable receptor b that does evolve (solid curve); the receptor sensitivity function of receptor b is determined by the polymorphic allele a_{1j1} . The activation level of each receptor class is a function of the particular stimulus and the receptor sensitivity function (e.g. the integral of the product).

function. Thus, receptor activation can be described by a two-valued vector (z_a, z_b) , as shown in the right panel of figure 7.

As illustrated in figure 8a, we assume that z_a and z_b represent the combined (e.g. summed) responses of the two classes of receptor within a cluster. For simplicity, we assume that z_a and z_b are then combined into a single response, z (e.g. their difference), which represents the output of a cluster.

On the one hand, if a receptor cluster is receiving a stimulus from the surface of species 2, then the stimulus

will be a sample of the surface response from species 2 (e.g. one of the functions on the left in figure 6). On the other hand, if a cluster of receptors is receiving a stimulus from the background environment, then the stimulus will be a sample from one of the possible surfaces in the background.¹⁰ Thus, there will be a different probability distribution for z depending upon whether the cluster of receptors is receiving a stimulus from the surface of species 2 or from the background. When an individual of species 2 is not within range ($\omega_{11} = 0$), then all of the receptor clusters will receive a stimulus from the background.

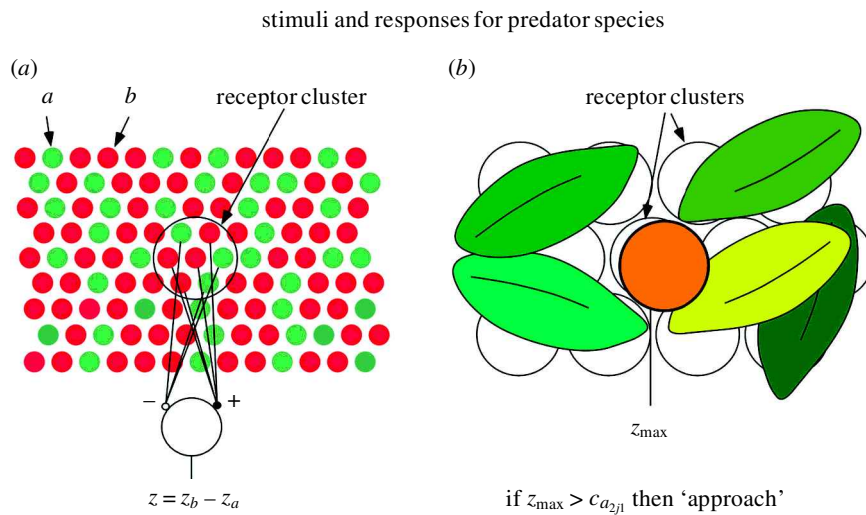


Figure 8. Example of a sensory circuit and its stimulation for the active predator species (species 1). (a) Receptor outputs are processed in clusters (one of which is circled) that may consist of anything from a single receptor to the whole receptor array. Within each cluster, the activities of the *a* receptors are combined into a signal z_a and those of the *b* receptors into a signal z_b ; these two signals are then combined (e.g. subtracted) to obtain the cluster response z . The maximum of the z responses across all clusters is a value z_{\max} . The predator approaches a potential target if the value z_{\max} exceeds some criterion determined by the polymorphic allele a_{2j1} . (b) One example is receptor clusters consisting of photoreceptors and stimuli consisting either of background foliage alone or background foliage plus a prey (as in this illustration). If receptor *b* is well matched to the prey and receptor *a* to the background foliage, then (as shown) z_{\max} might occur in the cluster stimulated by the prey. Note that high values of z_{\max} might occur by chance even when no prey is present.

When an individual of species 2 is within detection range ($\omega_{11} = 1$), then some of the receptor clusters will receive a stimulus from the surface of species 2 and the remainder will receive a stimulus from the background. For example, figure 8*b* illustrates the situation assuming that the background is foliage and the receptors are photoreceptors: the central receptor cluster is receiving a stimulus from the surface of species 2, and the rest are receiving a stimulus from the surfaces of the background foliage.

We assume two possible behavioural responses of species 1, represented by a response vector of length one, $\mathbf{r}_1 = (r_{11})$: ‘approach’ ($r_{11} = 1$) and ‘avoid’ ($r_{11} = 0$). The behavioural decision process is assumed to be a threshold criterion on the value of z from each cluster of receptors. If the maximum of all the z ’s (z_{\max}) exceeds this criterion, then species 1 approaches an object within its range; otherwise, it avoids all the objects within its range and moves to another location. We assume that the value of the criterion is controlled by the second component of the allele vector $\mathbf{a}_{j1} = (a_{1j1}, a_{2j1})$.

As shown in figure 9, the behaviour of species 1 is characterized ultimately by two probability functions, $f_{a_{1j1}}(z_{\max} | \omega_{11} = 0)$ and $f_{a_{1j1}, a_{112}}(z_{\max} | \omega_{11} = 1)$, and a criterion, $c_{a_{2j1}}$. The function on the left of figure 9 shows the probability of each possible value of z_{\max} given that no individual of species 2 is within detection range. The function on the right of the figure shows the probability of each possible value of z_{\max} given that some individual of species 2 is within detection range. The function on the right depends on a_{112} because when species 2 is within detection range its surface response is part of the stimulus.

From figure 9 it follows that the probability of an avoidance response given that no individual of species 2 is within range (a correct rejection), $p_{a_{j1}}(r_{11} = 0 | \omega_{11} = 0)$, is the area under the left function to the left of the criterion. The probability of an approach response given that no

individual of species 2 is within range (a false alarm), $p_{a_{j1}}(r_{11} = 1 | \omega_{11} = 0)$, is the area under the left function to the right of the criterion. The probability of an avoidance response given that an individual of species 2 is within range (a miss), $p_{a_{j1}}(r_{11} = 0 | \omega_{11} = 1)$, is the area under the right function to the left of the criterion. Finally, the probability of an approach response given that an individual of species 2 is within range (a hit), $p_{a_{j1}}(r_{11} = 1 | \omega_{11} = 1)$, is the area under the right function to the right of the criterion. These four probabilities describe the behaviour of species 1 given any state of the environment. In the fundamental equation (see figure 2), they specify all the possible values of

$$\sum_{s_1} p_{a_{j1}}(\mathbf{r}_1 | \mathbf{s}_1) p(\mathbf{s}_1 | \omega_1).$$

In a specific application, the two probability distributions in figure 9 would be determined by calculations from statistics of the surface characteristics of the background and of species 2, and from appropriate assumptions about the receptor sensitivity functions and sensory circuits of species 1. In our generic simulations we simply took the two probability functions to be equal-variance normal density functions whose means shifted in the intuitively appropriate fashion with changes in the alleles of the two species.

(iv) *Growth factor*

To specify the growth factor, one must describe how it depends upon the state of the environment and the organism’s response. First consider species 2. If species 1 is not within range for detecting species 2, then the average growth factor for species 2 is some nominal value defined by a fixed death rate, χ_2 , and birth rate, β_2 . If species 1 is within range and makes an approach response, then the individual of species 2 dies and does not give birth in that




Figure 9. Decision process and response likelihoods for the active predator species (species 1). The predator is assumed to have two possible responses: approach and avoid. The response is to approach if the maximum of the cluster responses, z_{\max} , exceeds a criterion value, $c_{a_{2j1}}$, which is determined by the polymorphic allele a_{2j1} . When the state of the environment is such that no prey is present, then the probability function for z_{\max} depends only on the allele controlling the predator’s receptors. When the state of the environment is such that a prey is present, then the probability function for z_{\max} depends on the allele controlling the predator’s receptors and on the allele controlling the prey’s surface characteristic. For each state of the environment, the probability of an approach response is the area under the corresponding probability function to the right of the criterion, and the probability of an avoidance response is the area to the left of the criterion.

time-step. If species 1 is within range but makes an avoidance response, then the birth and death rates for species 2 are their nominal values.

Next consider species 1. On the assumption that reproduction requires consumption of individuals from species 2, the birth rate is assumed to be zero unless there are hits. When there is a hit the birth rate has the value β_{1h} . The death rate will in general be different for each different combination of state of the environment and response; thus we allow a different death rate for hits (χ_{1h}), correct rejections (χ_{1cr}), misses (χ_{1m}) and false alarms (χ_{1fa}). (In the present simulations, we assumed that the death rates were the same for correct rejections and misses.)

(v) *Mutation*

In this example, we assume that mutations can occur only in species 1, and that reproduction is asexual. We further assume that the mutation rate is the same for all alleles of species 1. In our generic simulation, allele vectors are hypothetical and are represented by a value on some dimension (e.g. the location of the peak of the variable receptor’s sensitivity function or the location of the decision criterion). When mutations occur, the value of the new allele is obtained by sampling uniformly from a small range around the value of the allele before the mutation occurred ($\pm 10\%$ of the full range possible for that allele). Thus, every mutation results in a phenotypic change (i.e. there are no ‘silent’ mutations).

(vi) *Parameters*

The subsections above define all aspects of Bayesian natural selection needed for the simulations, except the starting values of the parameters.

The first set of parameters concerns the allele vectors. Specifically, we must select the starting number of allele vectors (n_1) for species 1, and the values of those vectors (\mathbf{a}_{j1}). In the present simulations, each vector component represents a specific phenotypic property of the organism.

The first component (a_{1j1}) represents the location of the sensitivity function of the variable receptor. The second component (a_{2j1}) represents the location of the decision criterion.

The second set of parameters concerns the numbers of organisms. We must specify the numbers of organisms carrying each allele vector at the beginning of the simulation ($t=0$) for species 1 ($O_{a_{j1}}(0)$) and for species 2 ($O_{a_{112}}(0)$). We must also specify the maximum number of organisms in species 1 ($o_{\max 1}$) and species 2 ($o_{\max 2}$) that can exist in the environment.

The third set of parameters concerns the growth rates. We must specify the birth and death rates for the active species 1 ($\beta_{1h}, \chi_{1h}, \chi_{1fa}, \chi_{1m}, \chi_{1cr}$) and for the passive species 2 (β_{2}, χ_{2}).

The fourth set of parameters concerns mutation. We must specify the mutation rate for species 1 (m_1) and the range of possible mutations for each of the allele vector components ($\Delta_{a_{11}}, \Delta_{a_{21}}$). We assume that the mutation rate is very small and that mutations occur independently in each allele.

The final parameter is the maximum probability that a given individual of species 2 is within detection range of any individual of species 1 ($p_{\text{range } 1}$).

(b) *Coevolution*

In this example, we generalize the previous example to allow for simultaneous evolution of species 2 (prey). Specifically, we now assume that there is an initial polymorphism in the allele vector that determines the stimulus signature of species 2 relevant for detection by species 1 (predator). This polymorphism is represented by a collection of different alleles, $a_{112}, a_{122}, a_{132}, \dots$, and thus we now have a different fundamental equation for each of these alleles. We also allow mutations in the allele vectors of both species 1 and 2.

With respect to the prior probabilities, the state of the environment for species 1 now depends not only on

whether an individual of species 2 is within detection range but also on what allele it carries. Thus, there are now $n_2 + 1$ possible values along the first dimension of the prior probability distribution for species 1: $\omega_{11} = 0$ (no individual of species 2 within detection range), $\omega_{11} = i$ (an individual of species 2 with allele a_{1i2} within detection range). Thus, the prior probability distribution for ω_{11} becomes

$$p(\omega_{11} = i; t) = p_{\text{range}_2} \frac{O_{a_{1i2}}(t)}{o_{\text{max}_2}}; \tag{3.5}$$

cf. equation (3.3).

With respect to the response likelihoods for species 1, recall that z is the integrated response from a cluster of fixed and variable receptors, and that z_{max} represents the maximum value of z across all clusters of receptors. Also recall that the probability function for z_{max} , when an individual of species 2 is within detection range, depends on the allele of that individual. Because there is now a different probability function, $f_{a_{1j1}, a_{1i2}}(z_{\text{max}} | \omega_{11} = i)$, for each different allele a_{1i2} in species 2, the probabilities of hits and misses now vary with a_{1i2} . Note that the probabilities of correct rejections and false alarms do not vary because in these cases no individual of species 2 is within detection range.

Specification of the initial conditions is nearly the same as for the case of transient polymorphism. The only difference is that now we must also specify the alleles and mutation parameters for species 2.

(c) Stable polymorphism

Here we generalize the first example to three species (see figure 5), one active predator (species 1) and two passive prey (species 2 and 3). We assume that both species 2 and 3 are food sources for species 1, and that species 2 and 3 are competing for space to reproduce. As in the first example, we assume that the passive species are not evolving or are doing so on a time-scale that is large relative to species 1. We further assume that there is no significant polymorphism in either species 2 or 3 that is, we can let a_{112} and a_{113} represent the alleles that determine the surface characteristic of species 2 and species 3, respectively. Thus, there is just one fundamental equation for species 2 and one for species 3. Different alleles in species 2 and 3 can lead to a stable polymorphism in species 1.

The prior probability distributions for species 2 and 3 are essentially the same as for species 2 in the first example. The environment vector for each species is two-dimensional, $\omega_2 = (\omega_{12}, \omega_{22})$ and $\omega_3 = (\omega_{13}, \omega_{23})$. The first dimension represents whether or not an individual of species 1 is within range and, if so, which allele vector it carries. The second dimension represents whether or not there is space to reproduce. The prior probability for species 1 is slightly different from the first example. The environment vector is still two-dimensional, but the first dimension represents whether a food source is within range and, if so, which food source: no food source ($\omega_{11} = 0$), species 2 ($\omega_{11} = 2$) or species 3 ($\omega_{11} = 3$).

First, consider the probability of having space to reproduce. Species 2 and 3, although passive, are in competition for space. We assume that the space required for

reproduction is the same for the two species, and thus the maximum number of organisms is the same, $o_{\text{max}_{2+3}}$. As before, the probability of not having space to reproduce is equal to the fraction of the total space occupied at time t :

$$p(\omega_{22} = 0; t) = p(\omega_{23} = 0; t) = \frac{O_2(t) + O_3(t)}{o_{\text{max}_{2+3}}}. \tag{3.6}$$

Of course, the probability of having space to reproduce is just one minus the probability of not having space to reproduce. For species 1 the probability of having space to reproduce is the same as in the first example.

Next, consider the probability of being within range of possible detection. For species 1, the probabilities of an individual of species 2 and 3 being within detection range are

$$p(\omega_{11} = 2; t) = p_{\text{range}_{2+3}} \frac{O_2(t)}{o_{\text{max}_{2+3}}}, \tag{3.7}$$

$$p(\omega_{11} = 3; t) = p_{\text{range}_{2+3}} \frac{O_3(t)}{o_{\text{max}_{2+3}}}. \tag{3.8}$$

For species 2 and 3, the probability of being within range of possible detection by an individual of species 1 with allele vector α_{j1} is given by

$$p(\omega_{12} = j; t) = p(\omega_{13} = j; t) = p_{\text{range}_1} \frac{O_{\alpha_{j1}}(t)}{o_{\text{max}_1}}, \tag{3.9}$$

cf. equation (3.4), and the probability of not being within range is one minus the sum of all the probabilities being within range.

With respect to the response likelihoods for species 1, we note that the probability function for z_{max} , when a food source is within detection range, depends on the species of the food source. Hence, the probabilities of hits and misses will also differ depending on the species of the food source.

Specification of the initial conditions is nearly the same as for the case of transient polymorphism. The only difference is that now we must also specify the allele, the initial number of organisms, and the growth parameters for species 3.

4. SIMULATION RESULTS

As described above, we carried out simulations of Bayesian natural selection for three cases: transient polymorphism, coevolution and stable polymorphism. In all three cases, we assumed that there was an active predator species (species 1) foraging for a passive prey species (species 2 or species 2 and 3). Species 1 was assumed to have a two-receptor sensory system, where the location of one of the receptors along the stimulus axis and the location of the response decision criterion were free to evolve. We also assumed that individuals within each species were competing for space to reproduce. See figure 5 for a schematic of the hypothetical environment.

In the first simulation (transient polymorphism) there were just two species, and species 2 was assumed to have a fixed surface characteristic that was not free to evolve. In the second simulation (coevolution) there were also two species, but now species 2 was assumed to have a surface characteristic that was free to evolve. In the third simulation (stable polymorphism) there were three species, and species 2 and 3 were assumed to have fixed (but different) surface characteristics that were not free to evolve.

The specific assumptions for each of the three cases are described in the previous section. The computational flow was the same in each case.

- (i) Set the initial conditions; create a separate fundamental equation for each allele vector in each species; set $t = 0$.
- (ii) Evaluate the fundamental equation for each allele vector to determine the expected number of births and deaths at time $t + 1$.
- (iii) For each allele vector, randomly sample from a Poisson probability distribution to obtain the actual number of births and from a binomial probability distribution to obtain the actual number of deaths.
- (iv) For each allele vector, randomly sample from the actual number of births to obtain the number of mutations; for each mutation randomly sample from the mutation range to obtain a new allele vector; for each new allele add a new fundamental equation.
- (v) For each allele vector, update the count of the number of organisms; for each allele vector with a count of zero, eliminate its fundamental equation.
- (vi) Set $t = t + 1$; go to step (ii).

These simulation steps are straightforward except for step (ii), which is described in more detail in Appendix A.

For each simulation we also computed the maximum fitness ideal observer for species 1, under the assumption that all properties of the organism are fixed, except for the location of the variable receptor and the location of the decision criterion (see Appendix A). In other words, we determined the locations of the variable receptor and decision criterion that maximized fitness (average growth factor). The maximum fitness ideal observer is useful for interpreting the results of simulations and provides a check on their validity. It is important to note that maximum fitness ideal observers typically change during the course of evolution because of changes in the prior probability distributions.

(a) *Transient polymorphism*

Representative results from the simulation of transient polymorphism are displayed as needle diagrams in figure 10. Figure 10*a* shows the states of species 1 (lower panel) and species 2 (upper panel) after one step of the simulation, figure 10*d* shows the asymptotic (equilibrium) states, and figure 10*b,c* show intermediate steps. Each small square in the plots indicates a particular allele vector, and the length of the vertical line attached to a square indicates the relative number of organisms carrying that allele vector. For species 1 (predator), the horizontal axis represents the location of the variable receptor's sensitivity function, and the vertical axis represents the location of the decision criterion. For species 2 (prey), the horizontal

axis represents the location of the surface characteristic. Stars show the optimal allele vector as given by the maximum fitness ideal observer.

We chose an initial state in which species 1 is a relatively small population containing a diverse set of allele vectors (chosen at random), and species 2 is a relatively large population containing a single allele vector. As can be seen, the allele vectors for species 1 converge over time to the values predicted by the maximum fitness ideal observer. There is a fast initial pruning of non-optimal allele vectors followed by a slow convergence limited by the rate and range of mutations. As species 1 evolves, there is a decline in the number of organisms in species 2, especially in earlier stages.

It is important to note that the optimal allele vector changes continuously during the course of evolution. These changes occur only in the optimal location of the decision criterion; the optimal location for the receptor's sensitivity function does not change. Two factors are responsible for the changes in the optimal location of the criterion. First, as the number of organisms in species 2 declines, the prior probability that an individual of species 1 will be within detection range of an individual of species 2 also declines. Second, as the number of organisms in species 1 increases, the prior probability of finding space to reproduce decreases. This simple example illustrates the dynamic character of prior probability functions in Bayesian natural selection.

The pattern of results in figure 10 is representative of those we obtained using a variety of different initial parameter values. We found that under most conditions, the simulations converge to the optimum given by the maximum fitness ideal observer. The only exceptions involved parameter values that led to extinction of one species or the other. As expected, varying the location of the surface characteristic of species 2 led to equivalent changes in the equilibrium location of the receptor sensitivity function of species 1. Varying the other relevant parameters—detection range of species 1, birth and death rates of species 1 and 2, and maximum numbers of the two species—led to changes in the equilibrium location of the decision criterion. Surprisingly, some of these changes were small relative to what would be expected given the magnitude of changes in the prior probability distributions.

For example, figure 11 shows the effect on the optimal decision criterion of varying the birth rate parameter (β_{1h}) of species 1. As birth rate increases, the population of species 2 decreases, and hence the prior probability that an individual of species 1 will have a target within range decreases (see figure 11*a*). According to classical ideal observer theory, as the prior probability of the target decreases, the decision criterion that maximizes detection accuracy increases (see the dashed curve in figure 11*b*). Thus, when the birth rate is 10 the target prior probability is 0.05 and optimal criterion is 2.5, whereas when the birth rate is near zero the target prior probability is 0.8 and the optimal criterion is an order of magnitude smaller. By contrast, when fitness rather than accuracy is maximized, the optimal criterion value changes only slightly (see the solid curve in figure 11*b*). This result occurs because the increase in birth rate and the resulting decrease in target prior probability have opposite effects on optimal criterion placement. The higher birth rate increases the




Figure 10. Results from simulation of transient polymorphism. These needle diagrams show the alleles and the number of organisms carrying those alleles at four times during the course of evolution. Each allele is represented by a small square; the length of the line attached to the square indicates the number of organisms carrying that allele. Stars indicate the optimal allele vector according to the maximum fitness ideal observer. (a) After the first step ($t = 1$), with randomly chosen alleles for species 1. (b,c) Intermediate steps ($t = 100$ and $t = 1000$). (d) Asymptotic/equilibrium state ($t = 20\,000$).

reproductive payoff for finding a target, pushing the criterion down, while the simultaneous lowering of the target prior probability decreases the probability of reproductive payoff, pushing the criterion up.

The relatively invariant decision criteria (at equilibrium) revealed in this simulation might reflect a general principle. Species having similar receptor systems—but different reproductive efficiencies—might be expected to have almost the same sensory decision criteria. Similarly, evolution of the genes controlling reproductive efficiency within a species might be expected to require little or no coevolution of the genes controlling sensory decision criteria.

(b) *Coevolution*

The simulation for coevolution was the same as in the previous case, except that we allowed evolution in the sur-

face characteristic of species 2 (prey). The population of species 2 was relatively large and contained an initial polymorphism. The mutation rate and maximum mutation step size for species 2 were the same as for species 1 (predator). We assumed that the surface characteristic for species 2 could approach but not reach the location of the surface characteristic of the background, which was at the centre of the horizontal axis (i.e. we did not allow perfect camouflage).

Representative results from the simulation are shown in figure 12. Stars indicate the optimal allele values according to the maximum fitness ideal observer. As can be seen, the allele populations for both species evolved over time. The alleles for species 2 migrated to be as close to the background surface characteristic as possible, and the alleles for species 1 followed along. As in the first simu-

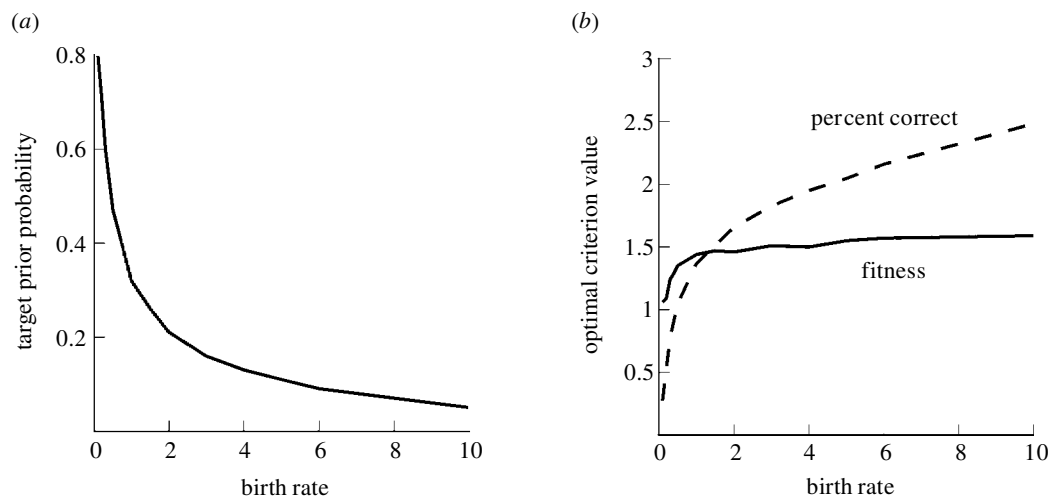


Figure 11. Effect of variations in reproductive efficiency on the equilibrium values of the decision criterion in the predator species (species 1). (a) This plot shows the effect on the prior probability (at equilibrium) of a prey (a target) being within detection range of an individual predator, as a function of the birth rate parameter of the predator. (b) On the one hand, if the goal is to maximize prey detection accuracy, then the optimal decision criterion should increase by more than an order of magnitude over the range of birth rates (dashed curve), because of the decline in the target prior probability. On the other hand, if the goal is to maximize fitness, then the decision criterion should be 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. The solid curve is predicted both by the maximum fitness ideal observer and by the simulation of natural selection.

ation, the alleles for the predator species generally migrated to the optimal values. Also, the optimal values defined by the maximum fitness ideal observer again changed during the course of evolution but even more dramatically than in the previous case.

The coevolutionary trajectories sampled in figure 12 are represented in a more continuous fashion in figure 13*a,b*. Figure 13*a* plots the mean value of the allele determining the surface characteristic of species 2 on the vertical axis and the mean value of the allele determining the location of the receptor sensitivity function of species 1 on the horizontal axis. (Note that the mean value of the allele determining the location of the decision criterion is not represented in the figure.) Figure 13*b* plots the number of organisms in species 2 on the vertical axis and the number of organisms in species 1 on the horizontal axis. In both plots, the arrowheads indicate the direction of the coevolutionary trajectory. The four arrowheads have been placed in figure 13*a,b* so that they represent corresponding points in time in the two figure parts. The black circles indicate the stable asymptotic endpoint of the coevolution. Initially, the receptor location in species 1 evolves rapidly to match the wavelength location of the surface characteristic of species 2. During this period, the population of species 1 increases while that of species 2 decreases. Once the population of species 2 declines sufficiently, its surface characteristic begins to evolve towards an optimal match with the background environment. During this period, the population of species 2 increases while that of species 1 decreases. Finally, after the surface characteristic of species 2 has reached the optimal match with the background, the receptor location of species 1 continues to evolve towards its optimum. During this period the population of species 2 decreases while that of species 1 increases.

Figures 12 and 13*a,b* illustrate one pattern of coevolutionary trajectory. There are two other stable patterns, which are illustrated in figure 13*c-f*. Figure 13*c,d* show a

typical coevolutionary trajectory if the initial population of alleles in species 2 is located further to the right than in figure 12*a*. In this case, the surface characteristic for species 2 initially shifts away from an optimal match with the background environment. Given the relatively small mutation step size, shifts towards the background increase the probability of detection by species 1 (i.e. decrease the fitness of species 2). However, once the receptor location of species 1 evolves to match the surface characteristic of species 2, mutations towards the background in species 2 produce increases in fitness. Thus, the surface characteristic of species 2 and the receptor location of species 1 eventually reach the same endpoints as in the first coevolutionary simulation (figure 13*a,b*). Interestingly, this does not occur if the range of possible receptor locations in species 1 is substantially less than the range of possible surface characteristics in species 2 (see figure 13*e,f*). In this case, species 2 becomes trapped with a surface characteristic that is far from an optimal match with the background, and its population size is permanently depressed relative to the other coevolutionary trajectories.

(c) *Stable polymorphism*

The simulation for stable polymorphism was the same as for transient polymorphism, except that there were two prey species with different surface characteristics. All other aspects of the two prey species were initially the same. As mentioned earlier, the two prey species competed for the same reproduction space. Representative results from the simulation are shown in figure 14. As before, stars show the optimal allele vectors for species 1 according to the maximum fitness ideal observer. As can be seen, the allele vectors for species 1 evolve towards the two locally optimal vectors, resulting in a stable polymorphism. However, notice that throughout the course of evolution, even in the equilibrium state, there are oscillations in the numbers of organisms in species 2 and 3, and corresponding oscillations




Figure 12. Results from simulation of coevolution. Each allele vector is represented by a small square; the length of the line attached to the square indicates the number of organisms carrying that allele vector. Stars indicate the optimal allele vector according to the maximum fitness ideal observer. (a) After the first step ($t = 1$). (b,c) Intermediate steps ($t = 100$ and $t = 1000$). (d) Asymptotic/equilibrium state ($t = 20\,000$).

lations in the numbers of organisms in species 1 that carry the alleles matched to these two species. This can be understood as follows. When, by chance, the individuals of species 1 carrying an allele vector matched to one of the prey species, say species 2, is particularly successful in foraging and/or reproduction, the competition for reproduction space drives down the number of individuals with the allele vector matched to species 3. All things being equal, fewer individuals means fewer offspring, and hence the disparity in numbers between individuals carrying the two alleles grows larger. However, when this happens, the number of individuals in species 2 is driven down, increasing the reproduction space available for species 3 and hence increasing their number. This inevitably leads to recovery in the number of individuals of species 1 whose alleles are matched to species 3. This alternation continues

indefinitely, although the periods and amplitudes of the successive cycles vary randomly.

Corresponding to the alternations in numbers of individuals among the three species, there are alternations in the globally optimal allele vector. When the prior probability of being within detection range is higher for species 2 than for species 3, then the globally optimal allele vector is the one matched to species 2, and vice versa.

5. DISCUSSION

There has been a great deal of interest recently in measuring the statistical properties of natural stimuli. This interest has been motivated in part by the intuition that statistical properties of the environment drive the design of perceptual systems through evolution and learning. As

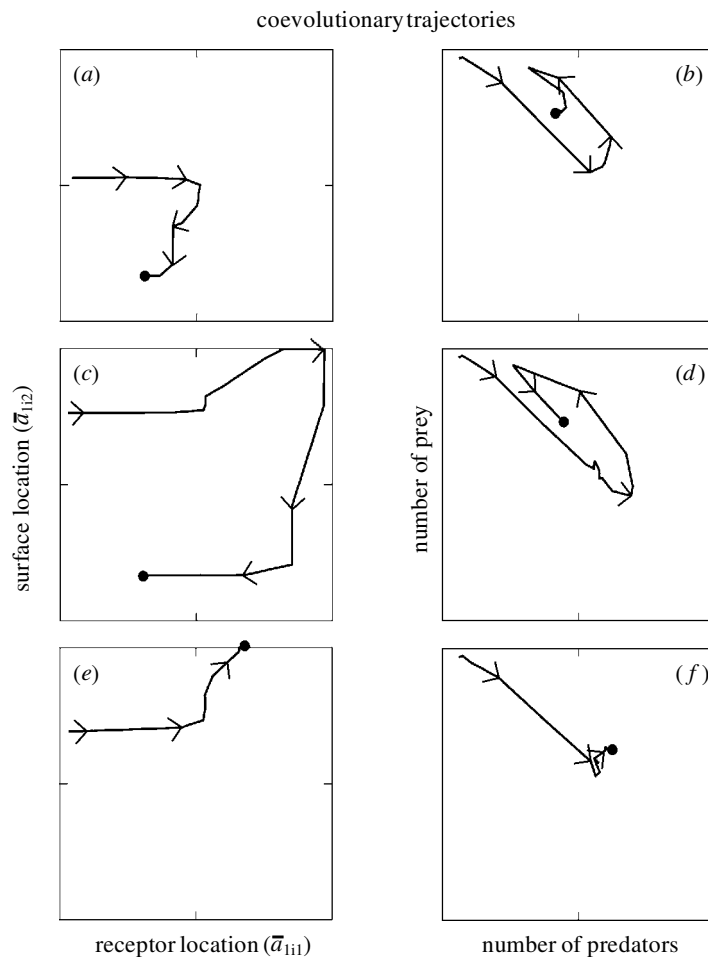


Figure 13. Three patterns of coevolutionary trajectory. In the left panels, the horizontal axis gives the average location of the wavelength peak of the receptor in the predator species, and the vertical axis gives the average location of the wavelength peak in the surface characteristic of the prey species. In the right panels, the horizontal and vertical axes give the population sizes of the predator and prey species. (a,b) Prey species evolves towards an optimal match with the background. (c,d) Prey species initially evolves away from an optimal match with the background but later evolves towards it after the receptor of the predator species evolves to match the surface characteristic of the prey. (e,f) Prey species evolves away from an optimal match with the background and becomes trapped because of a limited range in the evolution of the predator's receptor location.

described in § 1, some evidence in support of this intuition has been gained by comparing the processing of natural stimuli by real perceptual systems with that of ideal observers derived within the framework of Bayesian statistical decision theory. These ideal observer analyses have not only provided evidence for the importance of natural stimulus statistics in determining the design of perceptual systems, they have also provided a deeper understanding of the information contained in natural stimuli as well as of the computational principles employed in perceptual systems.

Formulating a Bayesian ideal observer requires selecting a utility function. In the past, the utility function has been selected by intuition or from the goals of some laboratory task (e.g. maximizing categorization or estimation accuracy). However, when ideal observers are considered from the viewpoint of natural selection, it is clear that the appropriate utility function is fitness (i.e. the birth and death rates for the species). Therefore, we propose that a maximum fitness ideal observer is generally more appropriate when considering the design and performance of perceptual systems.

Although appropriate ideal observer analyses can provide insight into how natural tasks and stimuli contribute to the design of a perceptual system, there are many other factors, including the incremental process of natural selection, which also make important contributions. To represent all of these factors in a coherent theoretical framework, we have proposed a formal version of natural selection based upon Bayesian statistical decision theory. The heart of our proposal is that each allele vector in each species under consideration is represented by a fundamental equation, which describes how the number of organisms carrying the allele vector at time $t + 1$ is related to the number of organisms carrying that 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. The process of natural selection is represented by iterating these fundamental equations in parallel over time, whilst updating the allele vectors using appropriate probability distributions for mutation and sexual recombination.

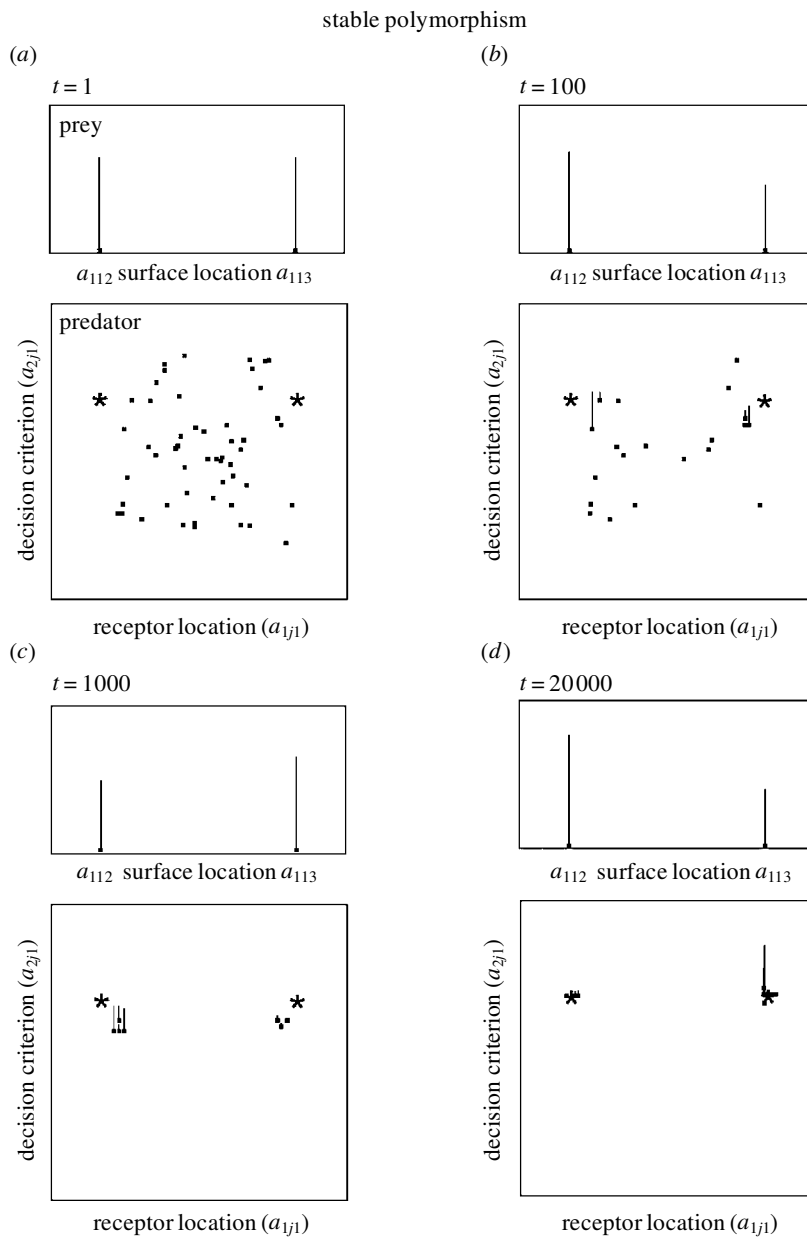


Figure 14. Results from the simulation of stable polymorphism. Each allele is represented by a small square; the length of the line attached to the square indicates the number of organisms carrying that allele. Stars indicate the two locally optimal allele vectors according to the maximum fitness ideal observer. (a) After the first step ($t = 1$). (b,c) Intermediate steps ($t = 100$ and $t = 1000$). (d) Asymptotic/equilibrium state ($t = 20\,000$).

Following the general description of Bayesian natural selection, we demonstrated how it might be applied in specific cases by developing generic simulations of transient polymorphism, coevolution and stable polymorphism. The primary purpose of these simulations was to demonstrate some of the flexibility of the Bayesian approach and to provide a few simulation templates into which specific measurements or assumptions might later be substituted.

Even our generic simulations provided some new insights into the design and evolution of perceptual systems. One interesting result is that for all three of our simulations, under all initial conditions tested, the alleles of the predator species converged to those corresponding to the maximum fitness ideal observer. Apparently, optimal perceptual mechanisms should evolve in situations where the assumptions in our generic simulations are satisfied.

Another insight from the simulations concerns the effect of reproductive efficiency on sensory decision criteria. As expected, we found that as the birth rate of a predator species increases, there is a decrease in the population of the prey species. According to classical ideal observer theory, where the utility function involves maximizing detection accuracy, this decrease in prey population should lead to a more conservative decision criterion for expending the energy required to make an approach response. However, our simulations suggest that, in fact, natural selection should pick a decision criterion that is relatively invariant across changes in reproductive efficiency of the predator species. This result, which is predicted by the maximum fitness ideal observer, occurs because the reduced probability of detecting prey is largely offset by the increased reproductive payoff of finding prey. Although in these trade-off situations we found a relative

invariance of the decision criterion, this will not always be the case. In general, decision criteria (e.g. categorization boundaries and estimation biases) are strongly dependent on the prior probability functions.

Most efforts to understand the evolution of perceptual systems have been directed at physical and physiological properties of the sensory organs. Relatively little effort has been directed at perceptual decision criteria, even though they play a critical role in determining perceptual performance. One reason may be the high likelihood that most properties of sensory organs are fixed adaptations, which are easier to study, whereas most properties of decision mechanisms are facultative adaptations, which are harder to study. Certainly in humans, and probably in most mammals, many perceptual decision criteria are adjusted via neural learning mechanisms to match the particular prior probabilities and stimuli that occur during the organism's lifetime. On the other hand, in lower organisms many decision criteria are fixed adaptations. For example, in reptiles, amphibians and birds, decisions as fundamental as which direction to strike, move or peck are often fixed adaptations, in that even newborns never learn to compensate for changes in visual direction produced by rotations or displacements of the visual image (Sperry 1951, 1956; Hess 1956). Thus, in lower organisms it may often be reasonable to regard both sensory organs and decision mechanisms as fixed adaptations, as we have done in our generic simulations.

In general, it is not obvious how to apply the equations of Bayesian natural selection to the study of facultative adaptations in perception. However, one possible place to start would be the evolution of habituation. Even in our simple generic simulations, it is possible to envision how habituation mechanisms might evolve. For example, prey targets that are within detection range of a predator may not be equally accessible; access could be blocked by elevation off the ground or by some intervening object such as a body of water. If the prey target is highly detectable, it may take a long time before a chance movement of the predator would result in a failure to detect the target, allowing the predator to move on to another location. This scenario could easily create a strong selection pressure for a habituation response in the decision mechanism. In other words, fitness could be substantially improved by temporarily raising the decision criterion, if the criterion has been exceeded for some extended period of time without acquisition of the prey.

Further insights from the simulations concern the effect of starting point and range of possible allele values on the pattern of coevolution between a predator's perceptual system and a prey's camouflage. If the predator's perceptual system is initially tuned to detect a surface characteristic more dissimilar from the background than the prey's surface characteristic, then the prey's surface characteristic will evolve towards a match with the background and the predator's perceptual system will evolve towards a match with the prey's surface characteristic. If the predator's perceptual system is initially tuned to detect a surface characteristic more similar to the background than the prey's surface characteristic, then the prey's surface characteristic will evolve away from a match to the background. However, as long as the evolvable range of the predator's perceptual system is greater than the evolvable range of the

prey's surface characteristic, then the prey's surface characteristic will eventually begin evolving towards the background after the predator's perceptual system becomes well matched to the prey's surface characteristic. On the other hand, if the evolvable range of the predator's perceptual system is less than that of the prey's surface characteristic, then the prey's surface characteristic becomes trapped in a poor match to the background. In general, there is no reason to think that the evolvable ranges for perceptual systems and surface characteristics should be equal and thus such evolutionary traps should not be uncommon in nature. From knowledge of the background stimulus statistics and the range of possible alleles in the predator and prey species, it should be possible to predict the pattern of coevolutionary trajectory.

It is useful to consider the relationship between the present proposal and other formal approaches to evolutionary biology and population genetics. Optimization theory has played an important role in evolutionary biology over the past three decades (see, for example, Maynard Smith & Price 1973; Alexander 1982; Maynard Smith 1982; Parker & Maynard Smith 1990). The aim of optimization theory has been to calculate physically, physiologically or behaviourally optimal solutions to adaptive problems that may serve as a benchmark in evaluating actual biological adaptations. Thus, ideal observer theory can be regarded as a special case of optimization theory that is well suited to analysing perceptual systems. However, there are some significant differences. Optimization theory has typically been expressed in the form of deterministic equations. By contrast, ideal observer theory emphasizes the statistical properties of the signals and system being optimized. These statistical properties are indispensable for determining the correct optimization in most perceptual tasks. A second difference is that ideal observer theory makes explicit the prior probabilities, thus forcing one to consider explicitly the context within which the optimization occurs.

Both classical ideal observer theory and optimization theory make use of indirect measures (proxies) for fitness, such as categorization accuracy or energy efficiency, in determining optimal design. Using a proxy is valid if it is monotonically related to fitness (Parker & Maynard Smith 1990). The difficulty is knowing whether monotonicity holds in particular cases. Intuition would suggest, for example, that in foraging animals, fitness is monotonically related to accuracy of detecting food sources. However, all three of our simulations demonstrated that this proxy leads to incorrect predictions for optimal criterion placement. As can be inferred from figure 11, detection accuracy and fitness are monotonically related at only one birth rate (the point of intersection of the solid and dashed curves). Thus, it is safest to use a maximum fitness ideal observer, whenever possible, in assessing optimal design.

Theories in population genetics have been concerned with modelling actual evolutionary trajectories rather than describing optimal solutions (see, for example, Dobzhansky 1970; Lewontin 1974; Crow 1986; Maynard Smith 1989; Hartl & Clark 1997),¹¹ and hence Bayesian natural selection can be regarded as a particular approach to population genetics. In population genetics, the emphasis is on the statistics of genetic variation within and across species and how those statistics change over time. Although

Bayesian natural selection also focuses on the statistics of genes, its unique contributions are to place a corresponding emphasis on the statistics of natural environments and to provide a formal representation of the environment–genome interaction. The approach is potentially valuable because of the critical role that complex statistical regularities of natural environments must play in natural selection. The sophistication and diversity of the organisms created through natural selection would surely be impossible without highly complex environments interacting with a highly complex organic chemistry.

Focusing on environmental statistics is useful, for example, in predicting when adaptations are likely to be fixed versus facultative. All things being equal, facultative adaptations generally require more resources than fixed adaptations (Williams 1966). Therefore, if the relevant statistical properties of a species's environment are such that a facultative adaptation would not perform any better than a fixed adaptation, then the adaptation is likely to be fixed. By incorporating measured environmental statistics into simulations of Bayesian natural selection, it may be possible to determine when variation in the environment is small enough to allow fixed adaptations and when it is large enough to promote facultative adaptations.

As might be expected, our Bayesian formulation of natural selection is compatible with many of the standard formulae in population genetics. Consider, for example, the fundamental equations for the simple case of two alleles competing for reproduction space in the world of our simulations, but otherwise having no effect on the prior probabilities in the environment. To simplify the notation, let x be the mean number of organisms carrying one allele and let y be the mean number carrying the other allele. The fundamental equations are probabilistic difference equations that can be approximated by deterministic differential equations. For this example, these differential equations take the form

$$\frac{dx}{dt} = \beta_x x (1 - (x + y)/o_{\max_{x+y}}) - \chi_x x$$

$$\frac{dy}{dt} = \beta_y y (1 - (x + y)/o_{\max_{x+y}}) - \chi_y y,$$

where β_x , β_y , χ_x and χ_y are the birth and death rates averaged over responses and stimuli. Inspection of these equations shows that population y will eventually drive population x to extinction if $\beta_y/\chi_y > \beta_x/\chi_x$. These equations are consistent with the so-called logistic equation, which is widely used in theoretical ecology (see, for example, Maynard Smith 1989). Specifically, in the case of a single population with a single allele, the size of the population is described by a logistic growth curve with an 'intrinsic rate of growth' of $\beta_x - \chi_x$ and a 'carrying capacity' of $(1 - \beta_x/\chi_x) o_{\max_x}$. As Maynard Smith (1989) notes, in population genetics the logistic equation is derived in an *ad hoc* fashion. By contrast, our version follows from first principles.

A shortcoming of existing approaches to evolutionary biology has been the lack of an explicit link between optimization theory and the formulae describing natural selection. In the Bayesian approach proposed here, such a link becomes feasible because the optimization theory (maximum fitness ideal observers) and the representation

of evolution (Bayesian natural selection) are cast in exactly the same terms—the same prior probability, stimulus likelihood and growth factor functions. Thus, the Bayesian approach should make it easier to grasp the relationship between optimal and actual design.

The aim of this paper has been to present a Bayesian framework for analysing the evolution of perceptual systems. Although the framework provides some conceptual coherence in thinking about perceptual systems, the real payoff will come if it proves useful in specific applications. Bayesian natural selection is most likely to be valuable in applications where measurements are available for many of the factors in the fundamental equations, or where there is a reasonable prospect of measuring many of the factors. For example, it may be possible to test hypotheses about the evolution (or coevolution) of colour vision in some species by incorporating actual measurements of relevant surface characteristics, surface pigment alleles, photoreceptor response characteristics, photopigment alleles, nominal birth and death rates, and the consequences of detection on these nominal rates. As reviewed by Regan *et al.* (2001), colour vision is an area where many of these measurements are being made in a rigorous fashion, and hence it may be an area where simulations based on Bayesian natural selection could usefully be applied in the near future. Other areas of potential application in the near future might involve perceptual systems in certain lower organisms that have a preponderance of fixed adaptations, well-characterized genetics and physiology, and measurable or controllable environments (e.g. *Drosophila*, zebra fish, *Aplysia*, and various bacteria).

It is important to note that focusing on fixed adaptations in lower organisms does not preclude the study of sophisticated perceptual mechanisms. For example, Gestalt perceptual grouping rules may well involve facultative adaptations in humans and other primates, but they are so important for object identification that similar rules must have evolved as fixed adaptations in lower organisms. Thus, it seems likely that the co-occurrence statistics of contour elements in natural images (like those in figure 3b) have led to the evolution of fixed contour grouping rules in the visual systems of many lower organisms.

Finally, we note that the Bayesian approach as outlined here is quite flexible. In the simulation examples, we left out many of the complexities that would be required in particular applications: the stimulus likelihood distributions, realistic anatomical and physiological constraints, realistic probability distributions for possible mutations and for sexual inheritance of alleles, and so on. However, the framework allows direct incorporation of these complexities once estimates of the relevant information are available. Also, the Bayesian approach is not restricted to perceptual systems. It applies to any biological system that can be characterized in terms of an input and an output. For example, species 2 in our simulations did not possess a perceptual system but a surface characteristic whose input (stimulus) was an energy distribution from the environment and whose output (response) was another energy distribution emitted back into the environment. Thus, the stimulus and response in Bayesian natural selection could represent just about anything from the input and output of a transcription enzyme, to the input and

output of a cell membrane, to the input and output of an organ, to the input and output of an organism.

We thank John Loehlin, Robert Sekuler, David Buss and Dennis McFadden for many helpful comments and suggestions. Supported by NIH grants EY02688 and EY11747 to W.S.G. and by NIH grant DC00427 to R.L.D.

ENDNOTES

¹In the case of ties at the highest probability, one can pick arbitrarily from those tied categories.

²Throughout this paper we adopt the standard convention that capital letters refer to random quantities, and boldface letters refer to vector quantities. Throughout this paper, ‘vector’ refers to an ordered list of properties (generally, integer- or real-valued quantities).

³Bayes theorem follows directly from the definition of conditional probability:

$$p(c_i \& S) = p(c_i | S)p(S) = p(S | c_i)p(c_i).$$

⁴Marr characterized the notion of a computational theory in the context of information processing. A computational theory defines an information processing problem, the constraints that apply in solving the problem, and the formal solution to the problem given the constraints. The solution is specified abstractly (i.e. independently of a particular algorithm or implementation).

⁵In many standard treatments of Bayesian statistical decision theory, the utility function is referred to as the ‘loss function’, and the expected utility as the ‘risk’. We have chosen the terminology that is more common in economics and psychology in order to avoid representing benefits as negative numbers.

⁶The method of extracting edge elements was different from any tested by Konishi *et al.* (2002), but it accurately determined location and orientation of the edge elements.

⁷The examples in this section are cases in which certain aspects of system design (e.g. photoreceptor placement) approach optimality. However, overall performance in a perceptual task will typically fall short of ideal because of neural noise or inefficiency in other aspects of the system.

⁸Equation (2.1) is actually the more general version of the fundamental equation. Equation (2.2) is appropriate for cases in which the relationship between environment states and responses is mediated by stimulus vectors.

⁹If the environment vector includes some internal factors, then the response likelihood is $p_a(\mathbf{r} | \mathbf{s}, \boldsymbol{\omega}_I)$, where $\boldsymbol{\omega}_I$ is a vector representing the internal factors.

¹⁰Here we do not make any assumptions about the surface characteristics of the background other than to assume that the fixed receptor a is well matched to the average background stimulus.

¹¹This includes game theoretic models of evolution. Some applications of game theory can be regarded as an optimization theory because they describe the rational (optimal) strategy given the rules of a particular game. However, most applications in evolutionary biology include the hill-climbing constraints of natural selection.

APPENDIX A

(a) Constrained maximum fitness ideal observers

To be useful, maximum fitness ideal observers must often incorporate some biological constraints. For example, to be of value in interpreting the present simulations, the ideal observer was constrained to include only two classes of receptor. Here, we briefly describe the formal structure of constrained maximum fitness ideal observers.

Constraints can be represented by a function g that maps the stimulus \mathbf{S} into an intermediate signal \mathbf{Z} . This intermediate signal could represent the output of any subsystem of the organism. Further, the properties of the subsystem that are allowed to vary in optimizing fitness can be represented by a parameter vector $\boldsymbol{\theta}$. Thus, we have

$$\mathbf{Z} = g_{\boldsymbol{\theta}}(\mathbf{S}).$$

Given these subsystem constraints, the average growth factor across all states of the environment is

$$\gamma_{\boldsymbol{\theta}}(\mathbf{r} | \mathbf{Z}) = \sum_{\boldsymbol{\omega}} \gamma(\mathbf{r}, \boldsymbol{\omega}) p_{\boldsymbol{\theta}}(\mathbf{Z} | \boldsymbol{\omega}) p(\boldsymbol{\omega}),$$

where the signal likelihood, $p_{\boldsymbol{\theta}}(\mathbf{Z} | \boldsymbol{\omega})$, is calculated from the subsystem function, $g_{\boldsymbol{\theta}}(\mathbf{S})$, and the stimulus likelihood, $p(\mathbf{S} | \boldsymbol{\omega})$. This equation is the same as equation (1.4) in the text, except that the stimulus \mathbf{S} is replaced by the intermediate signal \mathbf{Z} and the likelihood depends on $\boldsymbol{\theta}$.

Now, the optimum response for particular values of \mathbf{Z} and $\boldsymbol{\theta}$ is the one that maximizes the average growth factor across all states of the environment:

$$\mathbf{r}_{\text{opt}}(\mathbf{Z}, \boldsymbol{\theta}) = \arg \max[\gamma_{\boldsymbol{\theta}}(\mathbf{r} | \mathbf{Z})],$$

where $\arg \max$ is an operator that yields the argument corresponding to the maximum value of a function. For any particular value of the parameter vector, the average growth factor of the ideal observer across all states of the environment and all stimuli is

$$\gamma(\boldsymbol{\theta}) = \sum_{\boldsymbol{\omega}} p(\boldsymbol{\omega}) \sum_{\mathbf{z}} \gamma(\mathbf{r}_{\text{opt}}(\mathbf{z}, \boldsymbol{\theta}), \boldsymbol{\omega}) p_{\boldsymbol{\theta}}(\mathbf{z} | \boldsymbol{\omega}),$$

or equivalently,

$$\gamma(\boldsymbol{\theta}) = \sum_{\mathbf{z}} \max_{\mathbf{r}} [\sum_{\boldsymbol{\omega}} \gamma(\mathbf{r}, \boldsymbol{\omega}) p_{\boldsymbol{\theta}}(\mathbf{z} | \boldsymbol{\omega}) p(\boldsymbol{\omega})].$$

The optimum value of the parameter vector is the one that maximizes this average growth factor:

$$\boldsymbol{\theta}_{\text{opt}} = \arg \max[\gamma(\boldsymbol{\theta})].$$

We see then that the constrained maximum fitness ideal observer will pick the response that satisfies the equation

$$\mathbf{r}_{\text{opt}}(\mathbf{Z}) = \arg \max[\gamma_{\boldsymbol{\theta}_{\text{opt}}}(\mathbf{r} | \mathbf{Z})].$$

In other words, natural selection will have achieved the optimal solution (given the assumed subsystem constraints) if it converges simultaneously on the optimal parameter vector for the subsystem and on the optimal processing of the intermediate signals.

(b) Evaluation of fundamental equations

Here we describe in more detail how the fundamental equations were evaluated in the simulations. We describe only the calculations for the case of coevolution, but the others are very similar.

First, consider the calculations for species 1. The fundamental equation for a particular allele vector of species 1 is as follows:

$$\begin{aligned} \bar{O}_{a_{j1}}(t+1) &= O_{a_{j1}}(t) \sum_{\boldsymbol{\omega}_{11}, \boldsymbol{\omega}_{21}} p(\boldsymbol{\omega}_{11}, \boldsymbol{\omega}_{21}; t) \times \\ &\quad \sum_{r_{11}} \gamma(r_{11}, \boldsymbol{\omega}_{11}, \boldsymbol{\omega}_{21}) \sum_{s_1} p_{a_{j1}}(r_{11} | \mathbf{s}_1) p(\mathbf{s}_1 | \boldsymbol{\omega}_{11}). \end{aligned}$$

To evaluate the summation over the possible stimuli, recall that, by definition,

$$p_{a_{j1}}(r_{11} | \boldsymbol{\omega}_{11}) = \sum_{s_1} p_{a_{j1}}(r_{11} | \mathbf{s}_1) p(\mathbf{s}_1 | \boldsymbol{\omega}_{11}).$$

There are just two possible responses that species 1 can make, approach ($r_{11} = 1$) or avoid ($r_{11} = 0$). Also, the poss-

ible states of the environment relevant to the summation over possible stimuli are: no individual of species 2 is within detection range ($\omega_{11} = 0$), or an individual of species 2 with allele vector i is within detection range ($\omega_{11} = i$). As described in the text (see figure 9), the summation across the possible stimuli reduces to an integral of a probability distribution for the maximum response (z_{\max}) from the receptor clusters in the sensory system. Thus, the probabilities of correct rejection, false alarm, miss and hit are, respectively,

$$p_{a_{j1}}(r_{11} = 0 | \omega_{11} = 0) = g_{a_{j1}}(c_{a_{2j1}}) = \int_{-\infty}^{c_{a_{2j1}}} f_{a_{j1}}(z_{\max}) dz_{\max}$$

$$p_{a_{j1}}(r_{11} = 1 | \omega_{11} = 0) = 1 - g_{a_{j1}}(c_{a_{2j1}}),$$

$$p_{a_{j1}}(r_{11} = 0 | \omega_{11} = i) = g_{a_{j1}, a_{1i2}}(c_{a_{2j1}}) = \int_{-\infty}^{c_{a_{2j1}}} f_{a_{j1}, a_{1i2}}(z_{\max}) dz_{\max}$$

$$p_{a_{j1}}(r_{11} = 1 | \omega_{11} = i) = 1 - g_{a_{j1}, a_{1i2}}(c_{a_{2j1}}).$$

Because we assumed that the probability distributions are normal, each of these integrals could be evaluated using the standard cumulative normal integral function.

To compute the sum over the possible responses, note that the average growth factor for each possible state of the environment is

$$\gamma_{a_{j1}}(\omega_{11}, \omega_{21}) = \sum_{r_{11}} \gamma(r_{11}, \omega_{11}, \omega_{21}) p_{a_{j1}}(r_{11} | \omega_{11}).$$

Using the definitions of the birth and death rate parameters for species 1, we have

$$\gamma_{a_{j1}}(\omega_{11} = 0, \omega_{21} = 0) = 1 - \chi_{1cr} g_{a_{j1}}(c_{a_{2j1}}) - \chi_{1fa} [1 - g_{a_{j1}}(c_{a_{2j1}})],$$

$$\gamma_{a_{j1}}(\omega_{11} = i, \omega_{21} = 0) = 1 - \chi_{1m} g_{a_{j1}, a_{1i2}}(c_{a_{2j1}}) - \chi_{1h} [1 - g_{a_{j1}, a_{1i2}}(c_{a_{2j1}})],$$

$$\gamma_{a_{j1}}(\omega_{11} = 0, \omega_{21} = 1) = 1 - \chi_{1cr} g_{a_{j1}}(c_{a_{2j1}}) - \chi_{1fa} [1 - g_{a_{j1}}(c_{a_{2j1}})],$$

$$\gamma_{a_{j1}}(\omega_{11} = i, \omega_{21} = 1) = 1 - \chi_{1m} g_{a_{j1}, a_{1i2}}(c_{a_{2j1}}) + [\beta_{1h} - \chi_{1h}] [1 - g_{a_{j1}, a_{1i2}}(c_{a_{2j1}})].$$

Finally, to compute the sum over the possible states of the environment, we note that the average growth factor is

$$\gamma_{a_{j1}} = \sum_{\omega_{11}, \omega_{21}} p(\omega_{11}, \omega_{21}; t) \gamma_{a_{j1}}(\omega_{11}, \omega_{21}).$$

From our assumptions about the prior probabilities,

$$p(\omega_{11} = i, \omega_{21} = 0; t) = p_{\text{range}_2} \frac{O_{a_{1i2}}(t)}{o_{\max_2}} \frac{O_1(t)}{o_{\max_1}},$$

$$p(\omega_{11} = i, \omega_{21} = 1; t) = p_{\text{range}_2} \frac{O_{a_{1i2}}(t)}{o_{\max_2}} \left[1 - \frac{O_1(t)}{o_{\max_1}} \right],$$

$$p(\omega_{11} = 0, \omega_{21} = 0; t) = \frac{O_1(t)}{o_{\max_1}} \left[1 - \sum_{i=1}^{n_2} p_{\text{range}_2} \frac{O_{a_{1i2}}(t)}{o_{\max_2}} \right],$$

$$p(\omega_{11} = 0, \omega_{21} = 1; t) = \left[1 - \frac{O_1(t)}{o_{\max_1}} \right] \left[1 - \sum_{i=1}^{n_2} p_{\text{range}_2} \frac{O_{a_{1i2}}(t)}{o_{\max_2}} \right].$$

Next, consider the calculations for species 2. Because all relevant aspects of the responses of species 2 are represented in the behaviour of species 1, we can begin with the calculation of the average growth factor of species 2 for each possible state of the environment. If species 1 is not within range for detecting species 2 ($\omega_{12} = 0$), then the average growth factors for species 2 are the nominal values:

$$\gamma_{a_{1i2}}(\omega_{12} = 0, \omega_{22} = 0) = 1 - \chi_2,$$

$$\gamma_{a_{1i2}}(\omega_{12} = 0, \omega_{22} = 1) = 1 - \chi_2 + \beta_2.$$

If species 1 is within range, then the growth factors are scaled by the probability that species 1 does not approach:

$$\gamma_{a_{1i2}}(\omega_{12} = j, \omega_{22} = 0) = (1 - \chi_2) g_{a_{j1}, a_{1i2}}(c_{a_{2j1}}),$$

$$\gamma_{a_{1i2}}(\omega_{12} = j, \omega_{22} = 1) = (1 - \chi_2 + \beta_2) g_{a_{j1}, a_{1i2}}(c_{a_{2j1}}).$$

In other words, we assume that if species 1 approaches, then the growth factors are zero; otherwise, they are the nominal values.

To compute the sum over the possible states of the environment, we note that the average growth factor is

$$\gamma_{a_{1i2}} = \sum_{\omega_{12}, \omega_{22}} p(\omega_{12}, \omega_{22}; t) \gamma_{a_{1i2}}(\omega_{12}, \omega_{22}).$$

From our assumptions about the prior probabilities,

$$p(\omega_{12} = j, \omega_{22} = 0; t) = p_{\text{range}_1} \frac{O_{a_{j1}}(t)}{o_{\max_1}} \frac{O_2(t)}{o_{\max_2}},$$

$$p(\omega_{12} = j, \omega_{22} = 1; t) = p_{\text{range}_1} \frac{O_{a_{j1}}(t)}{o_{\max_1}} \left[1 - \frac{O_2(t)}{o_{\max_2}} \right],$$

$$p(\omega_{12} = 0, \omega_{22} = 0; t) = \frac{O_2(t)}{o_{\max_2}} \left[1 - \sum_{i=1}^{n_1} p_{\text{range}_1} \frac{O_{a_{i1}}(t)}{o_{\max_1}} \right],$$

$$p(\omega_{12} = 0, \omega_{22} = 1; t) = \left[1 - \frac{O_2(t)}{o_{\max_2}} \right] \left[1 - \sum_{i=1}^{n_1} p_{\text{range}_1} \frac{O_{a_{i1}}(t)}{o_{\max_1}} \right].$$

REFERENCES

Alexander, R. M. 1982 *Optima for animals*. London: Arnold.
 Annis, R. C. & Frost, B. 1973 Human visual ecology and orientation anisotropies in acuity. *Science* **182**, 729–731.
 Atick, J. J. & Redlich, A. N. 1992 What does the retina know about natural scenes? *Neural Comput.* **4**, 196–210.
 Attneave, F. 1954 Some informational aspects of visual perception. *Psychol. Rev.* **61**, 183–193.
 Barlow, H. B. 1957 Increment thresholds at low intensities considered as signal/noise discriminations. *J. Physiol. Lond.* **136**, 469–488.
 Barlow, H. B. 1961 Possible principles underlying the transformations of sensory messages. In *Sensory communication* (ed. W. A. Rosenblith), pp. 217–234. Cambridge, MA: MIT Press.
 Barlow, H. B. 1977 Retinal and central factors in human vision limited by noise. In *Vertebrate photoreception* (ed. H. B. P. Fatt), pp. 337–358. London: Academic.
 Bell, A. J. & Sejnowski, T. J. 1997 The ‘independent components’ of natural scenes are edge filters. *Vision Res.* **37**, 3327–3338.

- Berger, J. O. 1985 *Statistical decision theory and Bayesian analysis*. New York: Springer.
- Brainard, D. H. & Freeman, W. T. 1997 Bayesian color constancy. *J. Optical Soc. Am. A* **14**, 1393–1411.
- Brunswik, E. 1956 *Perception and the representative design of psychological experiments*. Berkeley, CA: University of California Press.
- Brunswik, E. & Kamiya, J. 1953 Ecological cue-validity of 'proximity' and other Gestalt factors. *Am. J. Psychol.* **66**, 20–32.
- Buchsbaum, G. & Gottschalk, A. 1984 Chromaticity coordinates of frequency-limited functions. *J. Optical Soc. Am. A* **1**, 885–887.
- Burgess, A. E., Wagner, R. F., Jennings, R. J. & Barlow, H. B. 1981 Efficiency of human visual signal discrimination. *Science* **214**, 93–94.
- Burton, G. J. & Moorehead, I. R. 1987 Color and spatial structure in natural scenes. *Appl. Optics* **26**, 157–170.
- Candy, J. V. & Xiang, N. (eds) 2001 *Signal processing acoustics and underwater acoustics: Bayesian signal processing approach in acoustics*. *J. Acoustical Soc. Am.* **109**, 2382–2384.
- Chiao, C., Vorobyev, M., Cronin, T. W. & Osorio, D. 2000 Spectral tuning of dichromats to natural scenes. *Vision Res.* **40**, 3257–3271.
- Coppola, D. M., Purves, H. R., McCoy, A. N. & Purves, D. 1998 The distribution of oriented contours in the real world. *Proc. Natl Acad. Sci. USA* **95**, 4002–4006.
- Crow, J. F. 1986 *Basic concepts in population, quantitative and evolutionary genetics*. New York: Freeman.
- Dan, Y., Atick, J. J. & Reid, R. C. 1996 Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *J. Neurosci.* **16**, 3351–3362.
- 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**, 553–564.
- Dixon, R. E. 1978 Spectral distribution of Australian daylight. *J. Optical Soc. Am.* **68**, 437–450.
- Dobzhansky, T. 1970 *Genetics of the evolutionary process*. New York: Columbia University Press.
- Dong, D. W. & Atick, J. J. 1995 Temporal decorrelation: a theory of lagged and nonlagged responses in the lateral geniculate nucleus. *Network: Comput. Neural Syst.* **6**, 159–178.
- Elder, J. H. & Goldberg, R. M. 1998 Statistics of natural image contours. In *IEEE Computer Society Workshop on Perceptual Organization in Computer Vision*. Santa Barbara, CA (available from marathon.csee.usf.edu/~sarkar/pocv_program.html).
- Field, D. J. 1987 Relations between the statistics of natural images and the response properties of cortical cells. *J. Optical Soc. Am. A* **4**, 2379–2394.
- Freeman, W. T. 1996 The generic viewpoint assumption in a Bayesian framework. In *Perception as Bayesian inference* (ed. D. C. Knill & W. Richards), pp. 365–389. Cambridge University Press.
- Geisler, W. S. 1984 Physical limits of acuity and hyperacuity. *J. Optical 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., Perry, J. S., Super, B. J. & Gallogly, D. P. 2001 Edge co-occurrence in natural images predicts contour grouping performance. *Vision Res.* **41**, 711–724.
- Green, D. M. & Swets, J. A. 1966 *Signal detection theory and psychophysics*. New York: Wiley.
- Hamilton, W. D. 1964a The genetical evolution of social behaviour, I. *J. Theor. Biol.* **7**, 1–16.
- Hamilton, W. D. 1964b The genetical evolution of social behaviour, II. *J. Theor. Biol.* **7**, 17–52.
- Hartl, D. L. & Clark, A. G. 1997 *Principles of population genetics*, 3rd edn. Sunderland, MA: Sinauer.
- Hess, E. H. 1956 Space perception in the chick. *Sci. Am.* **195** (7), 71–80.
- Hyvarinen, A. & Hoyer, P. 2000 Emergence of topography and complex cell properties from natural images using extensions of ICA. In *Advances in neural information processing systems*, vol. 12 (ed. S. A. Solla, T. K. Leen & K.-R. Muller), pp. 827–833. Cambridge, MA: MIT Press.
- Judd, D. B., MacAdam, D. L. & Wyszecki, G. W. 1964 Spectral distribution of typical daylight as a function of correlated color temperature. *J. Optical Soc. Am.* **54**, 1031–1040.
- Knill, D. C. 1998a Discrimination of planar surface slant from texture: human and ideal observers compared. *Vision Res.* **38**, 1683–1711.
- Knill, D. C. 1998b Surface orientation from texture: ideal observers, generic observers and the information content of texture cues. *Vision Res.* **38**, 1655–1682.
- Knill, D. C. & Richards, W. (eds) 1996 *Perception as Bayesian inference*. Cambridge University Press.
- Konishi, S., Yuille, A. L., Coughlan, J. M. & Zhu, S. C. 2002 Fundamental bounds on edge detection: learning and evaluating edge cues. *Pattern Anal. Machine Intell.* (Submitted.)
- Krinov, E. 1947 *Spectral reflectance properties of natural formations*. Ottawa, Ontario: National Research Council of Canada.
- Laughlin, S. B. 1981 A simple coding procedure enhances a neuron's information capacity. *Zeitschrift für Naturforschung.* **36c**, 910–912.
- Lewontin, R. C. 1974 *The genetic basis of evolutionary change*. New York: Columbia University Press.
- Liu, Z., Knill, D. & Kersten, D. 1995 Object classification for human and ideal observers. *Vision Res.* **35**, 549–568.
- Lythgoe, J. N. & Partridge, J. C. 1989 Visual pigments and the acquisition of visual information. *J. Exp. Biol.* **146**, 1–20.
- Malik, J., Martin, D., Fowlkes, C. & Tal, D. 2002 A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. *International Conference on Computer Vision* (Submitted.)
- Maloney, L. T. 1986 Evaluation of linear models of surface reflectance with small numbers of parameters. *J. Optical Soc. Am. A* **3**, 1673–1683.
- Maloney, L. T. & Wandell, B. A. 1986 Color constancy: a method for recovering surface spectral reflectance. *J. Optical Soc. Am. A* **3**, 29–33.
- Marr, D. 1982 *Vision: a computational investigation into the human representation and processing of visual information*. San Francisco, CA: Freeman.
- Maynard Smith, J. 1982 *Evolution and the theory of games*. Cambridge University Press.
- Maynard Smith, J. 1989 *Evolutionary genetics*. Oxford University Press.
- Maynard Smith, J. & Price, G. R. 1973 The logic of animal conflict. *Nature* **246**, 15–18.
- Olshausen, B. A. & Field, D. J. 1997 Sparse coding with an overcomplete basis set: a strategy by V1? *Vision Res.* **37**, 3311–3325.
- Osorio, D. & Vorobyev, M. 1996 Colour vision as an adaptation to frugivory in primates. *Proc. R. Soc. Lond. B* **263**, 593–599.
- Parker, G. A. & Maynard Smith, J. 1990 Optimality theory in evolutionary biology. *Nature* **348**, 27–33.
- Pelli, D. G. 1990 The quantum efficiency of vision. In *Vision: coding and efficiency* (ed. C. Blakemore), pp. 3–24. Cambridge University Press.
- Peterson, W. W., Birdsall, T. G. & Fox, W. C. 1954 The theory of signal detectability. *Trans. Inst. Radio Eng. Profess. Group Inform. Theory* **4**, 171–212.

- Regan, B. C., Julliot, C., Simmen, B., Vienot, F., Charles-Dominique, P. & Mollon, J. D. 1998 Frugivory and colour vision in *Alouatta seniculus*, a trichromatic platyrrhine monkey. *Vision Res.* **38**, 3321–3327.
- Regan, B. C., Julliot, C., Simmen, B., Vienot, F., Charles-Dominique, P. & Mollon, J. D. 2001 Fruits, foliage and the evolution of primate colour vision. *Phil. Trans. R. Soc. Lond. B* **356**, 229–283.
- Rose, A. 1948 The sensitivity performance of the human eye on an absolute scale. *J. Optical Soc. Am.* **38**, 196–208.
- Schrater, P. & Kersten, D. 2001 Vision, psychophysics and Bayes. In *Statistical theories of the brain* (ed. R. P. N. Rao, B. A. Olshausen & M. S. Lewicki), pp. 39–64. Cambridge, MA: MIT Press.
- Shannon, C. 1948 A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 379–423.
- Sigman, M., Cecchi, G. A., Gilbert, C. D. & Magnasco, M. O. 2001 On a common circle: natural scenes and Gestalt rules. *Proc. Natl Acad. Sci. USA* **98**, 1935–1940.
- Simoncelli, E. P. & Olshausen, B. A. 2001 Natural image statistics and neural representation. *A. Rev. Neurosci.* **24**, 1193–1215.
- Simoncelli, E. P. & Schwartz, O. 1999 Image statistics and cortical normalization models. In *Advances in neural information processing systems*, vol. 11 (ed. M. S. Kearns, S. A. Solla & D. A. Cohn), pp. 153–159. Cambridge, MA: MIT Press.
- Sperry, R. W. 1951 Mechanisms of neural maturation. In *Handbook of experimental psychology* (ed. S. S. Stevens), pp. 236–280. New York: Wiley.
- Sperry, R. W. 1956 The eye and the brain. *Sci. Am.* **194**, 48–52.
- Srinivasan, M. V., Laughlin, S. B. & Dubs, A. 1982 Predictive coding: a fresh view of inhibition in the retina. *Proc. R. Soc. Lond. B* **216**, 427–459.
- Switkes, E., Mayer, M. J. & Sloan, J. A. 1978 Spatial frequency analysis of the visual environment: anisotropy and the carpentered environment hypothesis. *Vision Res.* **18**, 1393–1399.
- Timney, B. & Muir, D. W. 1976 Orientation anisotropy: incidence and magnitude in Caucasian and Chinese subjects. *Science* **193**, 699–700.
- van Hateren, J. H. 1992 A theory of maximizing sensory information. *Biol. Cybernetics* **68**, 23–29.
- van Hateren, J. H. & Ruderman, D. L. 1998 Independent component analysis of natural image sequences yields spatio-temporal filters similar to simple cells in primary visual cortex. *Proc. R. Soc. Lond. B* **265**, 2315–2320.
- van Hateren, J. H. & van der Schaaf, A. 1998 Independent component filters of natural images compared with simple cells in primary visual cortex. *Proc. R. Soc. Lond. B* **265**, 359–366.
- Vos, J. J. & Walraven, P. L. 1972 An analytical description of the line element in the zone-fluctuation model of color vision: 1. Basic concepts. *Vision Res.* **12**, 1327–1344.
- Watson, A. B., Barlow, H. B. & Robson, J. G. 1983 What does the eye see best? *Nature* **302**, 419–422.
- Wertheimer, M. 1958 Principles of perceptual organization. In *Readings in perception* (ed. D. C. Beardslee & M. Wertheimer), pp. 115–135. Princeton, NJ: Van Nostrand.
- Williams, G. C. 1966 *Adaptation and natural selection*. Princeton University Press.
- Yuille, A. L. & Bülthoff, H. H. 1996 Bayesian decision theory and psychophysics. In *Perception as Bayesian inference* (ed. D. C. Knill & W. Richards), pp. 123–161. Cambridge University Press.