PERSPECTIVE

How biological vision succeeds in the physical world

Dale Purves[a,b,c,1], Brian B. Monson[a], Janani Sundararajan[a], and William T. Wojtach[c]

[a]Neuroscience and Behavioral Disorders Program, Duke-NUS Graduate Medical School, Republic of Singapore 169857; [b]Department of Neurobiology, Duke University Medical Center, Durham, NC 27710; and [c]Duke Institute for Brain Sciences, Duke University, Durham, NC 27708

Biological visual systems cannot measure the properties that define the physical world. Nonetheless, visually guided behaviors of humans and other animals are routinely successful. The purpose of this article is to consider how this feat is accomplished. Most concepts of vision propose, explicitly or implicitly, that visual behavior depends on recovering the sources of stimulus features either directly or by a process of statistical inference. Here we argue that, given the inability of the visual system to access the properties of the world, these conceptual frameworks cannot account for the behavioral success of biological vision. The alternative we present is that the visual system links the frequency of occurrence of biologically determined stimuli to useful perceptual and behavioral responses without recovering real-world properties. The evidence for this interpretation of vision is that the frequency of occurrence of stimulus patterns predicts many basic aspects of what we actually see. This strategy provides a different way of conceiving the relationship between objective reality and subjective experience, and offers a way to understand the operating principles of visual circuitry without invoking feature detection, representation, or probabilistic inference.

visual stimuli | luminance | lightness | empirical ranking | Bayes' theorem

In the 1960s and for the following few decades, it seemed all but certain that the rapidly growing body of information about the electrophysiological and anatomical properties of neurons in the primary visual pathway of experimental animals would reveal how the brain uses retinal stimuli to generate perceptions and appropriate visually guided behaviors (1). However, despite the passage of 50 years, this expectation has not been met. In retrospect, the missing piece is understanding how stimuli that cannot specify the properties of physical sources can nevertheless give rise to generally successful perceptions and behaviors.

The problematic relationship between visual stimuli and the physical world was recognized by Ptolemy in the 2nd century, Alhazen in the 11th century, Berkeley in the 18th century, Helmholtz in the 19th century, and many others since (2–12). To explain how accurate perceptions and behaviors could arise from stimuli that cannot specify their sources, Helmholtz, arguably the most influential figure over this history, proposed that observers augmented the information in retinal stimuli by making "unconscious inferences" about the world based on past experience. The idea of vision as inference has been revived in the last two decades using Bayesian decision theory, which posits that the uncertain provenance of retinal images illustrated in Fig. 1 is resolved by making use of the probabilistic relationship between image features and their possible physical sources (13–16).

The different concept of vision we consider here is based on a more radical reading of the challenge of responding to stimuli that cannot specify the metrics of the environment (17–20). The central point is that because there is no biologically feasible way to solve this problem by mapping retinal image features onto real-world properties, visual systems like ours circumvent it by generating perceptions and behaviors that depend on the frequency of occurrence of biologically determined stimuli that are tied to reproductive success. In what follows, we describe how this strategy of vision operates, how it explains the anomalous way we experience the physical world, and what it implies about visual system circuitry.

Vision in Empirical Terms

Although it is often assumed that the purpose of the evolved properties of the eye and early-level visual processing is to present stimulus features to the brain so that neural computations can recreate a representation of the environment, there is overwhelming evidence that we do not see the physical world for what it is (17, 18, 20–24). Whatever else this evidence may suggest, it indicates that to be useful, perceptions need not accord with measured reality. Indeed, generating veridical perceptions seems impossible given the uncertain significance of information conveyed by retinal stimuli (Fig. 1), even when the constraints of physics that define the world are taken into account (10–12).

In terms of neo-Darwinian evolution, however, a visual strategy that can circumvent

the inverse optics problem and explain why perceptions differ from the measured properties of the world is straightforward. Random changes in the structure and function of visual systems in ancestral forms would be favored by natural selection according to how well the ensuing percepts guided behaviors that promoted reproductive success. Any configuration of an eye and/or neural circuitry that strengthened the empirical link between visual stimuli and useful behavior would tend to increase in the population, whereas less beneficial ocular properties and circuit configurations would not. As a result, both perceptions and, ultimately, behaviors would depend on previously instantiated neural circuitry that promoted reproductive success; consequently, the recovery or representation of the actual properties of the world would be unnecessary.

Stimulus Biogenesis

The key to understanding how and why this general strategy explains the anomalous way we perceive the world when the properties of objects cannot be directly determined is recognizing that visual stimuli are not the passive result of physics or the statistics of physical properties in the environment, but are actively
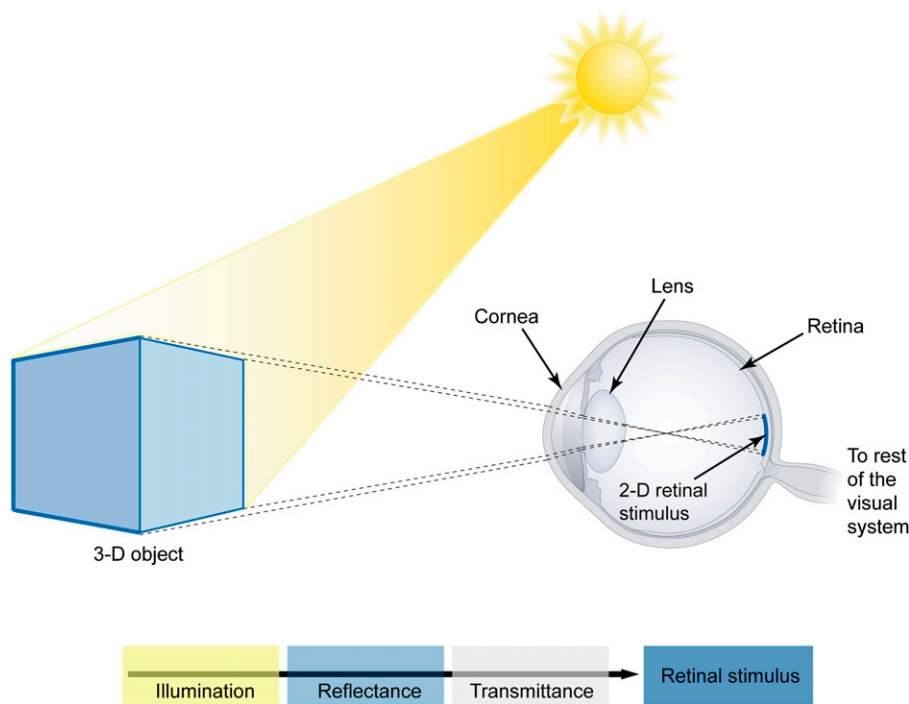
Fig. 1. The uncertain provenance of retinal stimuli. Images formed on the retina cannot specify physical properties such as illumination, surface reflectance, atmospheric transmittance, and the many other factors that determine the luminance values in visual stimuli. The same conflation of physical information holds for geometrical, spectral (color), and sequential (motion) stimulus properties. Thus, the behavioral significance of any visual stimulus is uncertain. Understanding how the image formation process might be inverted to recover properties of the environment under these circumstances is referred to as the inverse optics problem.

created according to their influence on reproductive success.

In contrast to the intuition that vision begins with a retinal image that is then processed and eventually represented in the visual brain according to a series of more-or-less logical steps, in the present argument the retinal image is just one of a series of stages in the biological transformation of disordered photon energy that begins at the corneal surface and continues in the processing carried out by the retina, thalamus, and cortex. In this framework, the "visual stimulus" is defined by the transformation of information by a recurrent network of ascending and descending connections, where the instrumental goal of generating perceptions and behaviors that work is met despite the absence of information about the actual properties of the world in which the animal must survive. Thus, although visual stimuli are usually taken to be images determined by the physical environment, they are better understood as determined by the biological properties of the eye and the rest of the visual system.

Many of these properties are already well known. For a visual stimulus to exist, photons must first be transformed into a topographical array ordered by the evolved properties of the eye. The evolved preneural properties that accomplish this are the dimensions of the eye, the shape and refractive index of the cornea, the dynamic characteristics of the lens, and the properties of ocular media, all of which serve to filter and focus photons impinging on a small region of the corneal surface. This process is continued by an arrangement of photoreceptors that restricts transduction to a limited range of photon energies, and the chain of early-level neural receptive field properties that continue to transform the biologically crafted input at the level of the retina. Although the nature of neural processing is less clear as one ascends in the primary visual system, enough is known about the organization of early-level receptive fields to provide a general idea of how they contribute to this overall strategy of relying on the frequency of occurrence of visual stimuli to generate successful perceptions, as described in the following section. The major role of the physical world in this understanding of vision is simply to provide empirical feedback regarding which perceptions and behaviors promoted reproductive success, and which did not.

An Example: The Perception of Lightness

To illustrate how this concept of vision works, consider the biological transformation of radiant energy into stimuli at an early stage where the preneural and neural events are best understood. Because increasing the luminance of any region of a retinal image increases the number of photons captured by the relevant photoreceptors, common sense suggests that physical measurements of light intensity and its perceived lightness should be proportional, and that two regions returning the same amount of light should appear to be equally light or dark. Perceptions of lightness, however, do not meet these expectations: In psychophysical experiments, the apparent lightness elicited by the luminance values at any particular region of a retinal image is clearly nonlinear and depends heavily on the surrounding luminance values (20, 21, 24).

To understand the significance of these discrepancies, take a typical luminance pattern on the retina arising from photons that are ordered by the evolved properties of the eye. For all intents and purposes, an image such as the example in Fig. 2A will have occurred only once; it is highly unlikely that the retina of an observer would ever again be activated by exactly the same pattern of luminance values falling on the same topographical array of millions of receptors. Because patterns like this are effectively unique, even a large catalog of such images would be of little or no help in promoting useful visual behavior on an empirical (trial and error) basis. However, smaller regions of the image, such as those sampled by the templates in Fig. 2A, would have occurred more than once, some many times, as shown by the distributions in Fig. 2B.

There is, of course, a lower limit to the size of samples that would be useful. If, for example, the size of the sample were reduced to a single point, the frequency of occurrence of the "pattern" would be maximal, but the resulting perceptions and behaviors would be based on a minimum of information. The greatest biological success would presumably arise from frequently occurring samples that comprised relatively small patterns in which the responses of the relevant neurons used information supplied by both the luminance value at any point and a tractable number of surrounding luminance values. This arrangement corresponds to the way retinal images are in fact processed by the receptive fields of early-level visual neurons, which, in the central vision of rhesus macaques (and
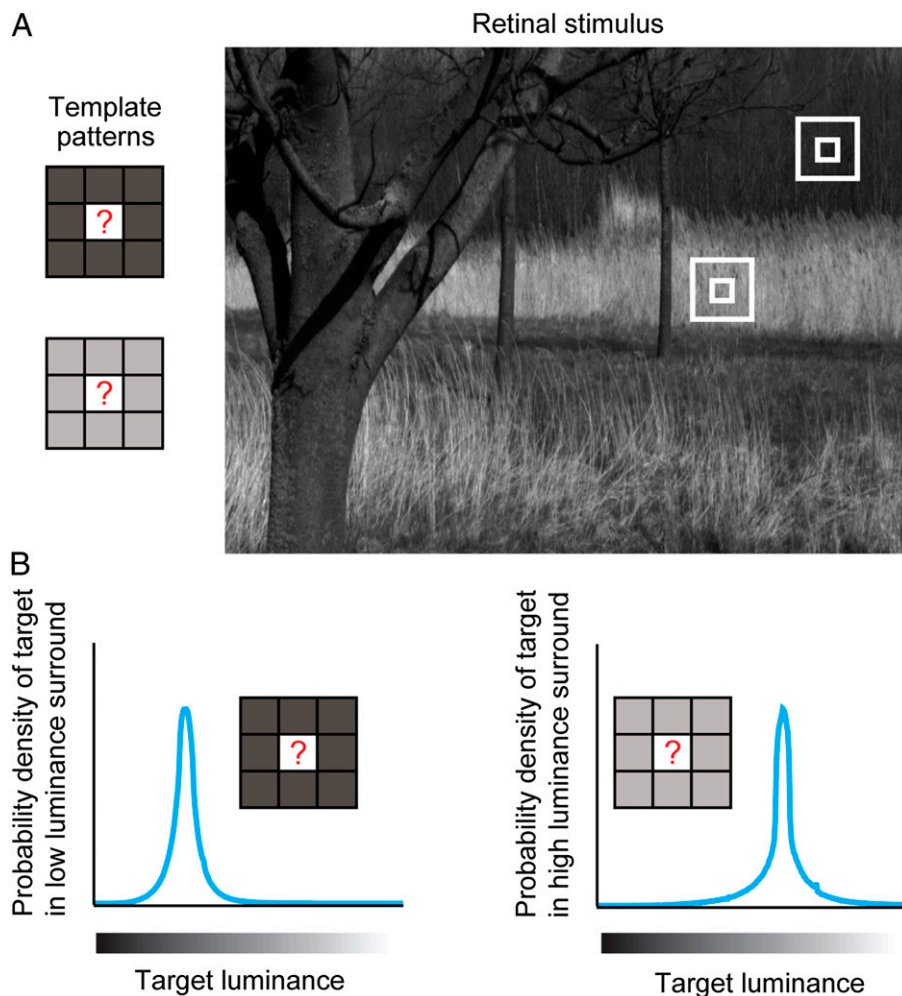
A

Retinal stimulus

Template patterns

B

Fig. 2. Accumulated human experience with luminance patterns. (*A*) To evaluate the concept that perception arises as a function of accumulated experience over evolutionary time, calibrated digital photographs can be sampled with templates about the size of visual receptive fields to measure how often different patterns of luminance occur in visual stimuli. (*B*) By repeated sampling, the frequency of occurrence of the luminance of any target region in a pattern of luminance values (indicated by a question mark) can be represented as a frequency distribution. The frequency of occurrence of the central region's luminance is different in the two surrounds, as would be true for any other pattern of luminance values assessed in this way. (The background image in *A* is from ref. 50; the data in *B* are after ref. 27).

presumably humans), are on the order of a degree or less of visual arc (25, 26)—roughly the size of the templates used in Fig. 2*A*.

To explore the merits of this concept of vision, templates like those in Fig. 2*A* can be used to sample the patterns that are routinely processed at the early stages of the visual pathway (the information extracted at other stages would, in principle, work as well). If perceptions of lightness indeed depend on the frequency of occurrence of small patterns of luminance values, then these data should predict what we see. One way of representing the frequency of occurrence of such stimuli is by transforming the distributions in Fig. 2*B* into cumulative distribution functions, thereby allowing the target luminance values in different surrounds to be ranked relative to one

another (Fig. 3). In this way, the lightness values that would be elicited by the luminance value of any region of a pattern in the context of surrounding luminance values can be specified. In the present concept of vision, the differences in these ranks account for the perceived differences in lightness of the identical target luminance values in Fig. 3.

Similar analyses have been used to explain not only the perception of simple luminance patterns like those in Figs. 2 and 3 but also perceptions elicited by a variety of complex luminance patterns (20, 27), geometrical patterns (18), spectral patterns (28), and moving stimuli (29, 30). In addition, artificial neural networks that evolve on the basis of ranking the frequency of luminance patterns can rationalize major aspects of early-level

receptive field properties in experimental animals (31, 32).

Why Stimulus Frequency Predicts Perception and Behavior

Missing from this account, however, is why the frequencies of occurrence of visual stimuli sampled in this way predict perception. The reason, we maintain, is that the relative number of times biologically generated patterns are transduced and processed in accumulated experience tracks reproductive success. In Fig. 3, for example, the frequencies of occurrence of the patterns at the stage of photoreception have caused the central luminance value to occur more often when in the lower luminance surround than in the higher one, resulting in a steeper slope at that point on the cumulative distribution function. If the relative ranking along this function corresponds to the perception of lightness, then the higher the rank of a target luminance (T) in a given surround relative to another target luminance with the same surround, the lighter the target should appear. Therefore, because the target luminance in a darker surround (Fig. 3, *Left*) has a higher rank than the same target luminance in a lighter surround (Fig. 3, *Right*), the former should be seen as lighter than the latter, as it is. Because the frequency of occurrence of patterns is an evolved property—and because these relative rankings along the function correspond to perception—the visually guided behaviors that result will in varying degrees have contributed to reproductive success. Thus, by aligning the frequencies of occurrence of light patterns over evolutionary time with perceptions of light and dark and the behaviors they elicit, this strategy can explain vision without solving the inverse optics problem.

Visual Perception on This Basis

Despite the inclination to do so, it would be misleading to imagine that the perceptions predicted by the relative ranking of luminance or other patterns depend on information about the "statistics of the environment." It is, of course, true that because physical objects tend to be uniform in their local composition, nearby luminance values in evolved retinal image patterns tend to be similar; indeed, the work of Brünswik (4) and, later, Gibson (33), which focused on how constraints of the environment might be conveyed in the structure of images, relied on this and other statistical information to explain vision. However, as illustrated in Fig. 1, the relationship between properties of the physical world and retinal images conflates such information, undermining
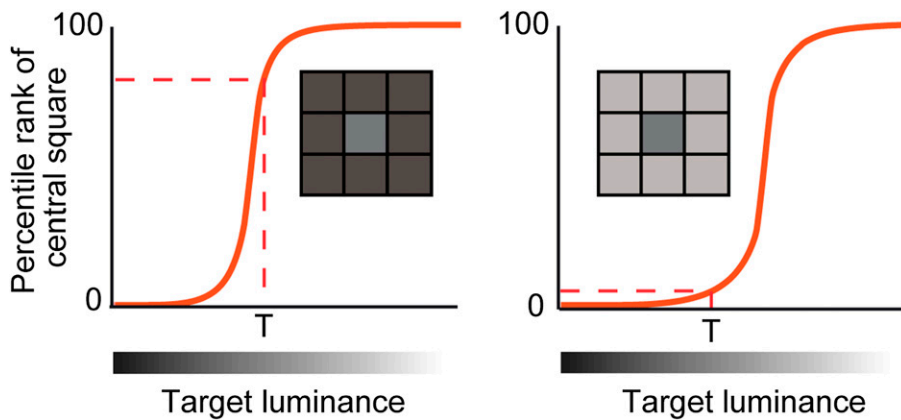
Fig. 3. Predicting lightness percepts based on the frequency of occurrence of stimulus patterns. The frequency distributions from Fig. 2B are here transformed to distribution functions that indicate the cumulative frequency of occurrence of the central target luminance given the luminance of the (*Inset*) surround. The dashed lines show the percentile rank of a specific central luminance value (*T*) in each distribution. As *Insets* show, central squares with identical photometric values elicit different lightness percepts (called "simultaneous lightness contrast") predicted by their relative rankings (relative frequencies of occurrence).

strategies that rely on statistical features of the environment to explain perception.

Although circumventing the inverse problem empirically gives the subjective impression that we perceive the actual properties of objects and conditions in the world, this is not the case. Nor does responding to luminance values (or other image attributes) according to the frequency of occurrence of local patterns reveal reality or bring subjective values "closer" to objective ones. It therefore follows that these discrepancies between lightness and luminance—or any other visual qualities and their physical correlates—are not "illusions" (22, 23) but simply signatures of the strategy we and, presumably, other visual animals have evolved to promote useful behaviors despite the inability of biological visual systems to measure physical parameters.

In sum, successful perceptions and behavior arise not because the actual properties of the world are recovered from images, but because the perceptual values assigned by the frequency of occurrence of visual stimuli accord with the reproductive success of the species and individual. As a result, the visual qualities that we see are better understood as signifying perceptions and behaviors that led to reproductive success in the past rather than encoding information, statistical or otherwise, about the world in the present.

Other Interpretations of Vision
What, then, can be said about other concepts of vision, and how they compare with the strategy of vision presented here? Three current frameworks are considered: vision as detecting and representing image features, vision as probabilistic inference, and vision as efficient coding.

Vision as Feature Detection and Representation. An early and still widely accepted idea is that visual (and other) sensory systems operate analytically, detecting behaviorally important features in retinal images that are then used to construct neural representations of the world at the level of the visual cortex. This interpretation of visual processing accords with electrophysiological evidence that demonstrates the selectivity of neuronal receptive fields, as well as with the compelling impression that what we see is external reality. Although attractive on these grounds, this interpretation of vision is ruled out by the inability of the visual system to measure the physical parameters of the world (Fig. 1), as well as its inability to explain a host of phenomena in luminance, color, form, distance, depth, and motion psychophysics on this basis (20).

Vision as Probabilistic Inference. More difficult to assess is the idea that vision is based on a strategy of probabilistic inference. Helmholtz introduced the idea of unconscious inference in the 19th century to explain how vision might improve responses to retinal images that he took to be inherently inadequate stimuli (3). In the first half of the 20th century, visual inferences were conceived in terms of gestalt laws or other heuristics. More recently, many mathematical psychologists and computer scientists have endorsed the idea of vision as statistical inference by proposing that images map back onto the properties of objects and conditions in the world as Bayesian probabilities (13, 15, 16, 34–37).

Bayes' theorem (38) states that the probability of a conditional inference about *A* given *B* being true (the posterior probability) is determined by the probability of *B* given *A* (the likelihood function) multiplied by the ratio of the independent probabilities of *A* (the prior probability) and *B*. This way of making rational predictions in the face of uncertainty is widely and successfully used in applications ranging from weather forecasting and medical diagnosis to poker and sports betting.

The value of Bayes' theorem as a tool to understand vision, however, is another matter. To be biologically useful, the posterior probability would have to indicate the probability of a property of the world (e.g., surface reflectance or illumination values) underlying a given visual stimulus. This, in turn, would depend on the probability of the visual stimulus given the physical property (the likelihood) and the prior probability of that state of the world. Although this approach is logical, information about the likelihood and prior probabilities is simply not available to the visual system given the inverse problem, thereby negating the biological feasibility of this explanation. In contrast, the empirical concept of vision described here avoids these problems by pursuing a different goal: fomenting reproductive success despite an inability to recover properties of the physical world in which behavior must take place. Although the frequency of occurrence of stimuli is often used to infer the probability of an underlying property of the physical world given an image, no such inferences are being made in this empirical strategy. Nor does the approach rely on a probabilistic solution: The biologically determined frequency of occurrence of visual stimuli simply generates useful perceptions and behaviors according to reproductive success.

These reservations add to other criticisms of Bayesian decision theory applied to cognitive issues, and to neuroscience generally (39, 40).

Vision as Efficient Coding. Another popular framework for understanding vision and its underlying circuitry is efficient coding (5, 41–45). A code is a rule for converting information from one form to another. In vision, coding is understood as the conversion of retinal stimulus patterns into the electrochemical signals (receptor, synaptic and action potentials) used for communication with the rest of the brain; this information is then taken to be decoded by further computational processes to achieve perceptual and behavioral effects. Given the nature of sensory transduction and the distribution of peripheral sensory effects to distant sites by action potentials, coding for the purpose of

neural computation seems an especially apt metaphor, and has been widely accepted (44, 46, 47).

Such approaches variously interpret visual circuits as carrying out optimal coding procedures based on minimizing energy use (5, 42, 43, 48–50), making accurate predictions (51–53), eliminating redundancy (54), or normalizing information (55, 56). The common theme of these overlapping ideas is that optimizing information transfer by minimizing redundancy, lowering wiring costs, and/or maximizing the entropy of sensory outputs will all have been advantageous to visual animals (57).

The importance of efficiency (whether in coding or otherwise) is clearly a factor in any evolutionary process, and the importance of these several ways of achieving it is not in doubt. However, generating perceptions by means of circuitry that contends with a world whose physical parameters cannot be measured by biological vision is a different goal, in much the same way that the goals of any organ system differ from the concurrent need to achieve them as efficiently as possible. Thus, these efforts are not explanations of visual perception, which no more depends on efficiency than the meaning of a verbal message depends on how efficiently it is transmitted.

Implications for Future Research

Given the central role it has played in modern neuroscience, the way scientists conceive vision is broadly relevant to the future direction of brain research, its potential benefits, and its economic value. An issue much debated at present is the intention to invest heavily over the coming decade in a complete analysis of human brain connectivity at both macroscopic and microscopic levels (58–60) (also http://blogs.nature.com/news/2013/04/obama-launches-ambitious-brain-map-project-with-100-million.html, accessed February 24, 2014). The impetus for this initiative is largely based on the success of the human genome project in scientific, health, technical, and financial terms. To underscore this parallel, the goal of the project is referred to as obtaining the "brain connectome."

Although neuroscientists rightly applaud this investment in better understanding brain connectivity, the related technology and possible health benefits, a weakness in the comparison with the human genome project (and with genetics in general) is that the basic functional and structural principles of genes were already well established at the outset. In contrast, the principles underlying the structure and function of the human brain

and its component circuits remain unknown. Indeed, the stated aim of the brain connectome project is the hope that additional anatomical information will help establish these principles.

Given this goal, the operation of the visual system—the brain region about which most is now known—is especially relevant. If the function of visual circuitry, a presumptive bellwether for operations in the rest of the brain, has been determined by evolutionary and individual history rather than by logical "design" principles, then understanding function by examining brain connectivity may be far more challenging than imagined. Perhaps the most daunting obstacle is that reproductive success—the driver of any evolved strategy of vision—is influenced by a very large number of factors, many of which will be difficult to discern, let alone quantify. Thus, the relation between accumulated experience and reproductive success may never be specified in more than qualitative or semiquantitative terms.

In light of these obstacles, it may be that the best way to understand the principles underlying neural connectivity is to evolve increasingly complex networks in progressively more realistic environments. Until relatively recently, pursuing this goal would have been fanciful. However, the advent of genetic and other computer algorithms has made evolving artificial neural networks in simple environments relatively easy (31, 32). This approach should eventually be able to link evolved visual functions and their operating principles with the wealth of detail already known from physiological and anatomical studies over the last 50 y.

Conclusion

A central challenge in understanding vision is that biological visual systems cannot measure or otherwise access the properties of the physical world. We have argued that vision like ours addresses this challenge by evolving the ability to form and transduce small, biologically determined image patterns whose frequencies of occurrence directly link perceptions and behaviors with reproductive success. In this way, perceptions and behaviors come to work in the physical world without sensory measurements of the environment, and without inferences or the complex computations that are often imagined. As a result, however, vision does not accord with reality but with perceptions and behaviors that succeed in a world whose actual properties are not revealed. This framework for vision, supported by evidence from human psychophysics and predictions of perceptions based on accumulated experience (i.e., the frequency of occurrence of biogenic stimuli), implies that Gustav Fechner's goal of understanding the relationship between objective (physical) and subjective (psychological) domains (61) can be met if pursued in these biological terms rather than in the statistical, logical, and computational terms that are more appropriate to physics, mathematics, and algorithm-based computer science. Although it may not be easy to relate this understanding of vision to higher-order tasks such as object recognition, if the argument here is correct, then all further uses of visual information must be built up from the way we see these foundational qualities.

1 Hubel DH, Wiesel T (2005) Brain and Visual Perception. A Story of a 25-Year Collaboration (Oxford University Press, New York).
2 Berkeley G (1975) Philosophical Works Including Works on Vision, ed Ayers MR (Everyman/JM Dent, London).
3 Helmholtz HLFv (1909) [Helmholtz's Treatise on Physiological Optics], trans Southall JPC (1924–1925) (Optical Society of America, New York), 3rd Ed, Vols I–III. German.
4 Brünswik E (1956/1997) Perception and the Psychological Design of Representative Experiments (University of California Press, Berkeley), 2nd Ed.
5 Barlow HB (1961) Possible principles underlying the transformation of sensory messages. Sensory Communication, ed Rosenblith WA (MIT Press, Cambridge, MA), pp 217–236.
6 Lindberg DC (1977) Theories of Vision from al-Kindi to Kepler (University of Chicago Press, Chicago).
7 Campbell DT (1982) The "blind-variation-and-selective-retention" theme. The Cognitive-Developmental Psychology of James Mark Baldwin: Current Theory and Research in Genetic Epistemology, eds Broughton JM, Freeman-Moir DJ (Ablex, Norwood, NJ), pp 87–97.
8 Campbell DT (1985) Pattern matching as an essential in distal knowing. Naturalizing Epistemology, ed Kornblith H (MIT Press, Cambridge, MA), pp 49–70.
9 Barlow HB (1990). What does the brain see? How does it understand? Images and Understanding, eds Barlow HB, Blakemore CB, Weston-Smith EM (Cambridge University Press, Cambridge), pp 5–25.
10 Shepard RN (1994) Perceptual-cognitive universals as reflections of the world. Psychon Bull Rev 1(1):2–28.
11 Pizlo Z (2001) Perception viewed as an inverse problem. Vision Res 41(24):3145–3161.
12 Shepard RN (2001) Perceptual-cognitive universals as reflections of the world. Behav Brain Sci 24(4):581–601, and discussion (2001) 24:652–671.
13 Knill DC, Richards W (1996) Perception as Bayesian Inference (Cambridge University Press, Cambridge).
14 Rao RPN, Olshausen BA, Lewicki MS (2002) Probabilistic Models of the Brain: Perception and Neural Function (MIT Press, Cambridge, MA).
15 Kersten D, Mamassian P, Yuille A (2004) Object perception as Bayesian inference. Annu Rev Psychol 55:271–304.
16 Vilares I, Howard JD, Fernandes HL, Gottfried JA, Kording KP (2012) Differential representations of prior and likelihood uncertainty in the human brain. Curr Biol 22(18):1641–1648.
17 Purves D, Lotto B (2003) Why We See What We Do: An Empirical Theory of Vision (Sinauer Associates, Sunderland, MA).
18 Howe CQ, Purves D (2005) Perceiving Geometry: Geometrical Illusions Explained by Natural Scenestatistics (Springer, New York).
19 Purves D, Wojtach WT, Lotto RB (2011) Understanding vision in wholly empirical terms. Proc Natl Acad Sci USA 108(Suppl 3): 15588–15595.

20 Purves D, Lotto B (2011) *Why We See What We Do Redux: A Wholly Empirical Theory of Vision* (Sinauer Associates, Sunderland, MA).

21 Stevens SS (1975) *Psychophysics: Introduction to Its Perceptual, Neural and Social Prospects* (Wiley, New York).

22 Adelson EH (2000) Lightness perception and lightness illusions. *The New Cognitive Neuroscience*, ed Gazzaniga M (MIT Press, Cambridge, MA), pp 339–351.

23 Weiss Y, Simoncelli EP, Adelson EH (2002) Motion illusions as optimal percepts. *Nat Neurosci* 5(6):598–604.

24 Gilchrist A (2006) *Seeing Black and White* (Oxford University Press, Oxford).

25 Wiesel TN, Hubel DH (1966) Spatial and chromatic interactions in the lateral geniculate body of the rhesus monkey. *J Neurophysiol* 29(6):1115–1156.

26 Hubel DH, Wiesel TN (1968) Receptive fields and functional architecture of monkey striate cortex. *J Physiol* 195(1):215–243.

27 Yang Z, Purves D (2004) The statistical structure of natural light patterns determines perceived light intensity. *Proc Natl Acad Sci USA* 101(23):8745–8750.

28 Long F, Yang Z, Purves D (2006) Spectral statistics in natural scenes predict hue, saturation, and brightness. *Proc Natl Acad Sci USA* 103(15):6013–6018.

29 Wojtach WT, Sung K, Truong S, Purves D (2008) An empirical explanation of the flash-lag effect. *Proc Natl Acad Sci USA* 105(42):16338–16343.

30 Wojtach WT, Sung K, Purves D (2009) An empirical explanation of the speed-distance effect. *PLoS ONE* 4(8):e6771.

31 Ng C, Sundararajan J, Hogan M, Purves D (2013) Network connections that evolve to circumvent the inverse optics problem. *PLoS ONE* 8(3):e60490.

32 Morgenstern Y, Venkata DR, Purves D (2013) Early level receptive field properties emerge from artificial neurons evolved on the basis of accumulated visual experience with natural images. *Journal of Vision* 13(9):1160.

33 Gibson JJ (1979) *The Ecological Approach to Visual Perception* (Lawrence Erlbaum, Hillsdale, NJ).

34 Mamassian P, et al. (2002) Bayesian modelling of visual perception. *Probabilistic Models of the Brain: Perception and Neural Function*, eds Rao RPN, et al. (MIT Press, Cambridge, MA), pp 13–36.

35 Kersten D, Yuille A (2003) Bayesian models of object perception. *Curr Opin Neurobiol* 13(2):150–158.

36 Geisler WS, Diehl RL (2003) A Bayesian approach to the evolution of perceptual and cognitive systems. *Cognit Sci* 27(3):379–402.

37 Knill DC, Pouget A (2004) The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends Neurosci* 27(12):712–719.

38 Bayes T (1763) An essay toward solving a problem in the doctrine of chances. *Philos Trans R Soc London* 53:370–418.

39 Jones M, Love BC (2011) Bayesian Fundamentalism or Enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behav Brain Sci*, 34(4):169–188, 188–231.

40 Bowers JS, Davis CJ (2012) Bayesian just-so stories in psychology and neuroscience. *Psychol Bull* 138(3):389–414.

41 Shannon CE (1948) A mathematical theory of communication. *Bell Syst Tech J* 27:379–423, 623–656.

42 Attneave F (1954) Some informational aspects of visual perception. *Psychol Rev* 61(3):183–193.

43 Laughlin S (1981) A simple coding procedure enhances a neuron's information capacity. *Z Naturforsch C* 36(9-10):910–912.

44 Marr D (1982) *Vision: A Computational Investigation into Human Representation and Processing of Visual Information* (W.H. Freeman, San Francisco).

45 Olshausen BA (2013) 20 years of learning about vision: Questions answered, questions unanswered, and questions not yet asked. *Twenty Years of Computational Neuroscience*, ed Bower JM (Springer, New York), pp 243–270.

46 Dayan P, Abbott LF (2001) *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems* (MIT Press, Cambridge, MA).

47 Brown EN, Kass RE, Mitra PP (2004) Multiple neural spike train data analysis: State-of-the-art and future challenges. *Nat Neurosci* 7(5):456–461.

48 Atick JJ, Redlich AN (1992) What does the retina know about natural scenes? *Neural Comput* 4(2):196–210.

49 Field DJ (1994) What is the goal of sensory coding? *Neural Comput* 6(4):559–601.

50 van Hateren JH, van der Schaaf A (1998) Independent component filters of natural images compared with simple cells in primary visual cortex. *Proc Biol Sci* 265(1394):359–366.

51 Srinivasan MV, Laughlin SB, Dubs A (1982) Predictive coding: A fresh view of inhibition in the retina. *Proc R Soc London Ser B* 216(1205):427–459.

52 Rao RP, Ballard DH (1999) Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nat Neurosci* 2(1):79–87.

53 Hosoya T, Baccus SA, Meister M (2005) Dynamic predictive coding by the retina. *Nature* 436(7047):71–77.

54 Olshausen BA, Field DJ (1996) Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381(6583):607–609.

55 Schwartz O, Simoncelli EP (2001) Natural signal statistics and sensory gain control. *Nat Neurosci* 4(8):819–825.

56 Carandini M, Heeger DJ (2012) Normalization as a canonical neural computation. *Nat Rev Neurosci* 13(1):51–62.

57 Sterling P, Laughlin S (2013) *Principles of Neural Design* (MIT Press, Cambridge, MA), in press.

58 Abbott A (January 23, 2013) Brain-simulation and graphene projects win billion-euro competition. Nature, 10.1038/nature.2013.12291.

59 Anonymous (February 23, 2013) Only connect. *The Economist*.

60 Anonymous (March 9, 2013) Hard cell. *The Economist*.

61 Fechner GT (1860) *Elements der psychophysik* (Brietkopf und Hartel, Leipzig, Germany); trans Adler HE (1966) [*Elements of Psychophysics*] (Holt, Rinehart & Winston, New York). German.