2. Snyder, L. H., Batista, A. P. & Andersen, R. A. *Vision Res.* **40**, 1433–1441 (2000).
3. Colby, C. L. & Goldberg, M. E. *Annu. Rev. Neurosci.* **22**, 319–349 (1999).
4. Culham, J. C. & Kanwisher, N. G. *Curr. Opin. Neurobiol.* **11**, 157–163 (2001).
5. Gold, J. I. & Shadlen, M. N. *Trends Cogn. Sci.* **5**, 10–16 (2001).
6. Calton, J. L., Dickinson, A. R. & Snyder, L. H. *Nat. Neurosci.* **5**, 580–588 (2002).
7. Sparks, D. L. *Curr. Opin. Neurobiol.* **9**, 698–707 (1999).
8. Basso, M. A. & Wurtz, R. H. *Nature* **389**, 66–69 (1997).
9. Dorris, M. C. & Munoz, D. P. *J. Neurosci.* **18**, 7015–7026 (1998).
10. Ungerleider, L. G. & Mishkin, M. in *Analysis of Visual Behavior* (eds. Ingle, D. J., Goodale, M. A. & Mansfield, R. J. W.) 549–586 (MIT Press, Cambridge, Massachusetts, 1982).
11. Goodale, M. A. *Curr. Opin. Neurobiol.* **3**, 578–585 (1993).

Illusions, perception and Bayes

Wilson S. Geisler and Daniel Kersten

A new model shows that a range of visual illusions in humans can be explained as rational inferences about the odds that a motion stimulus on the retina results from a particular real-world source.

Artists can create powerful illusions of distance, size, shape and orientation by mimicking on canvas the images that would be formed on the retina by perspective projection from a three-dimensional environment. These perceptual errors seem to reveal a rational (but automatic) perceptual system designed to correctly interpret the retinal images evoked by the world. This appealing explanation of visual illusions is often viewed as insufficient, however, because many idiosyncratic illusions seem unlike the rational solution to any problem.

One such class of illusions concerns the effect of luminance contrast and shape on the perception of motion velocity. Surprisingly, the apparent speed and direction of a moving pattern often changes substantially as the contrast and shape of the pattern is varied. Such illusions are typically interpreted as the errors or epiphenomena of some imprecise neural mechanism that is attempting to compute one of the quantities relevant for motion perception. However, using the tools and logic of Bayesian statistical decision theory, in this issue Weiss *et al.*[1] show that many of these seemingly idiosyncratic motion illusions, in fact, may be exactly what one would expect from a rational perceptual system. Their results suggest that the human visual system may be closer to optimal than once believed.

Wilson Geisler is in the Center for Perceptual Systems, University of Texas Austin, Seay Psychology Building, Austin, Texas 78712-7789, USA. Daniel Kersten is in the Psychology Dept., University of Minnesota, 75 E. River Rd., Minneapolis, Minnesota 55455, USA. e-mail: geisler@.psy.utexas.edu

This study is an excellent example of how Bayesian concepts are transforming perception research by providing a rigorous mathematical framework for representing the physical and statistical properties of the environment, describing the tasks that perceptual systems are trying to perform, and deriving appropriate computational theories of how to perform those tasks, given the properties of the environment and the costs and benefits associated with different perceptual decisions.

The Bayesian framework had its beginnings in Helmholtz's notion of 'unconscious inference'—the idea that the visual system incorporates implicit knowledge of the environment and image formation, and uses this knowledge to infer, automatically and unconsciously, object properties from the ambiguous images they form on the retina (see also refs. 2–4). As Helmholtz knew, retinal images are ambiguous because of natural variations in viewpoint and lighting: very different objects can give rise to similar retinal images, and the same object can give rise to very different retinal images. For example, a circle and an ellipse can produce exactly the same retinal image if the circle is slanted appropriately in depth, and the same circle slanted in depth by different amounts can produce many different images. Thus, the ambiguous stimulus in **Fig. 1** can be seen as an ellipse in the frontal plane or as a circle on a slanted plane. Bayesian statistical decision theory prescribes a framework for optimally interpreting such ambiguous retinal images.

Theoretical devices that use Bayesian statistical decision theory to make optimal interpretations are called 'ideal observers'. As one might guess intuitively, the primary objective of an ideal observer is to compute the probability of each possible true state of the environment given the stimulus on the retina (the 'posterior probability distribution'). According to Bayes' theorem, the posterior probability is proportional to the product of the probability of each possible state of the environment before receiving the stimulus (the prior probability) and the probability of the stimulus given each possible state of the environment (the likelihood). Thus, when an ideal observer (**Fig. 1**) receives a stimulus, it computes the likelihood and then multiplies by the prior probability distribution to obtain the posterior probability distribution. In many applications, prior probability distributions over the space of possible objects, events and/or lightings represent the ideal observer's knowledge of the environment, and likelihood distributions represent the ideal observer's knowledge of projective geometry and the space of possible viewpoints. Once the posterior probability distribution is determined, an ideal observer convolves the posterior distribution with a utility function (or loss function), which specifies the costs and benefits associated with the different possible errors in the perceptual decision. The result of this operation is the expected utility (or Bayes' risk) associated with each possible interpretation of the stimulus. Finally, the ideal observer picks the interpretation that has the maximum expected utility.

Weiss *et al.*[1] derived such an ideal observer and showed that it displays many of the same 'illusions' perceived by human observers. Their analysis was based on the plausible assumptions that low velocities are more likely than high velocities, and that there is more variability at low contrasts (that is, retinal images are less reliable). Motion illusions can be understood as optimal adaptations rather than mistakes in other cases as well; for example, a Bayesian analysis predicts correctly that changing the motion of a shadow alone can change the apparent three-dimensional motion of

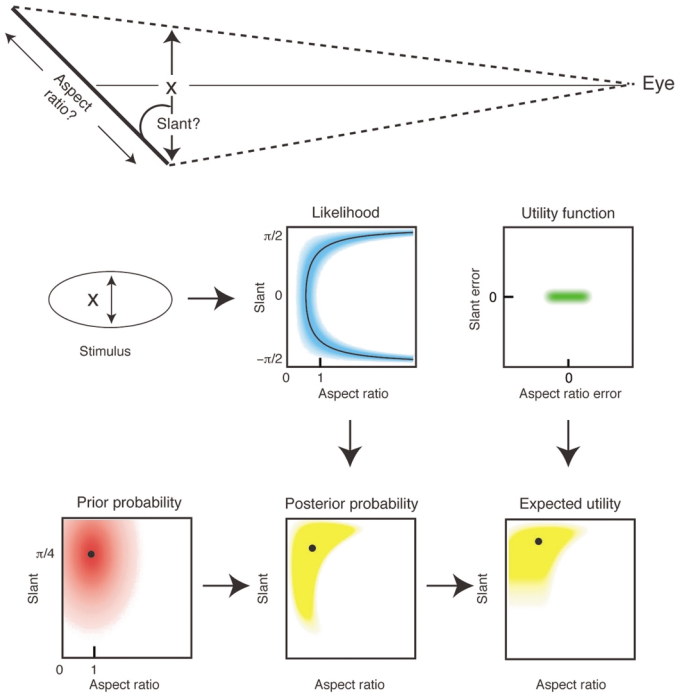

Fig. 1. Bayesian ideal observers for tasks involving the perception of objects or events that differ along two physical dimensions, such as aspect ratio and slant, size and distance, or speed and direction of motion. When a stimulus is received, the ideal observer computes the likelihood of receiving that stimulus for each possible pair of dimension values (that is, for each possible interpretation). It then multiplies this likelihood distribution by the prior probability distribution for each pair of values to obtain the posterior probability distribution—the probability of each possible pair of values given the stimulus. Finally, the posterior probability distribution is convolved with a utility function, representing the costs and benefits of different levels of perceptual accuracy, to obtain the expected utility associated with each possible interpretation. The ideal observer picks the interpretation that maximizes the expected utility. (Black dots and curves indicate the maxima in each of the plots.) As a tutorial example, the figure was constructed with a specific task in mind; namely, determining the aspect ratio and slant of a tilted ellipse from a measurement of the aspect ratio (x) of the image on the retina. The black curve in the likelihood plot shows the ridge of maximum likelihood corresponding to the combinations of slant and aspect ratio that are exactly consistent with x; the other non-zero likelihoods occur because of noise in the image and in the measurement of x. The prior probability distribution corresponds to the assumption that surface patches tend to be slanted away at the top and have aspect ratios closer to 1.0. The asymmetric utility function corresponds to the assumption that it is more important to have an accurate estimate of slant than aspect ratio.

an object[5]. However, the examples of Weiss *et al.*[1] are particularly interesting because they seemed unlikely to be optimal adaptations before the Bayesian analysis was done. For example, this analysis explains the odd combination of facts that a thin horizontally moving rhombus appears to move diagonally at low contrasts and horizontally at high contrasts, whereas a fat rhombus appears to move horizontally at all contrasts.

The Bayesian approach has advantages over other approaches in perception research. To begin with, it prescribes a principled method for determining optimal performance in a given perceptual task. This can be a very useful exercise because it forces one to consider carefully the various constraints that apply in the perceptual task, and the Bayesian ideal observer provides an appropriate benchmark against which to compare human performance. Furthermore, ideal observers are often easily modified by incorporating anatomical, physiological and other constraints, making them an excellent starting point for developing testable models[6,7].

Another advantage of the Bayesian approach is that it divides perceptual tasks into convenient and intuitive pieces that can be considered singly and then combined to understand the whole. For example, the Bayesian approach naturally partitions the 'generative model' for retinal images into a prior probability distribution on object shapes and materials, a prior probability distribution on lighting, and a stimulus likelihood distribution that incorporates the interactions between objects and lighting, and the effects of perspective projection and viewpoint. This Bayesian generative model has recently motivated a number of novel and testable hypotheses. To give just one example, Bloj *et al.*[8] demonstrated that perceived surface color can depend upon the perceived three-dimensional surface configuration in a perceptually bistable image, even when the retinal images never change. They predicted this effect from a Bayesian ideal observer that understands the effects of mutual surface illumination.

Finally, the Bayesian approach allows one to understand precisely how the reliability of different sources of information, including prior knowledge, should be combined by a perceptual system. Different sources of information do not always keep the same relative reliability, and hence a rational perceptual system should adjust the weights that it assigns to different information sources contingent upon their current relative reliabilities. This sort of weight adjustment is at the heart of the account of motion illusions from Weiss *et al.*[1]. When contrast is low, retinal image information becomes less reliable, and so the Bayesian ideal observer shifts more weight to the prior probability distribution on motion velocity; this shift in relative weight alters the optimal estimate of speed and direction. Several other recent applications of Bayesian analysis show that the human visual system can dynamically adjust weights on different information sources, often in a near-optimal way[9–11].

Although the Bayesian model proposed by Weiss *et al.*[1] makes the correct qualitative predictions for many motion illusions, the quantitative predictions are not perfect. This is not surprising for a nearly parameter-free model, but the potential reasons for less-than-perfect quantitative predictions are worth considering. First, the predictions depend on the exact shape of the prior probability and likelihood distributions. The assumption that low velocities are more probable than high velocities is likely to be correct qualitatively, but the specific prior probability distribution that Weiss *et al.* assumed is still just an educated guess. Undoubtedly, the prior probability and likelihood distributions incorporated implicitly into the visual system arise through a combination of evolution and perceptual learning, and thus it would be appropriate to estimate these distributions by measuring and analyzing natural scene statistics. Prior probability and likelihood distributions measured in the natural environment may lead to more accurate predictions of perceptual performance. For example, a Bayesian model derived from co-occurrence statistics for the geometrical rela-

tionships between edges in natural scenes quantitatively predicts human ability to detect contours on complex backgrounds, under a range of conditions[12].

Second, Weiss *et al.*[1] used the standard utility function that corresponds to picking the interpretation with the maximum posterior probability (**Fig. 1**, black dot in lower middle panel). The choice of utility function should have a minor effect in their situation, but the choice can be important in some situations[13,14]. For example, if smaller errors have greater utility than larger errors, and if the posterior probability distribution contains a tall narrow peak and a short wide peak, then the maximum expected utility may be at the short peak rather than the tall peak. Also, the most appropriate utility function when considering biological vision is arguably one based on fitness[15].

Third, as Weiss *et al.*[1] point out, their Bayesian model does not consider certain fundamental biological constraints, such as the limited dynamic range and limited speed of neural responses. Such constraints undoubtedly arise (ultimately) from biophysical limitations, and from the design compromises required because biological systems must perform many tasks. In supplementary material, Weiss *et al.* demonstrate that plausible neural constraints bring the predictions of their model into better quantitative agreement with the data.

Finally, real perceptual systems are designed through natural selection and perceptual learning, which sometimes find local optima in design space, rather than the global optimum of the Bayesian ideal observer. However, this fact does not diminish the value of the Bayesian approach; indeed, the concepts of Bayesian statistical decision theory lend themselves elegantly to rigorous formulations of natural selection[15]. Thus, even in the broader biological context of plasticity, learning and natural selection, the Bayesian approach may prove to be optimal.

1. Weiss, Y., Simoncelli, E. & Adelson, E. H. *Nat. Neurosci.* **5**, 598–604 (2002).
2. Freeman, W. T. *Nature* **368**, 542–545 (1994).
3. Knill, D. C., Kersten, D. & Yuille, A. in *Perception as Bayesian Inference* (eds. Knill, D. C. & Richards, R. W.) 1–21 (Cambridge Univ. Press, 1996).
4. Yuille, A. L. & Bülthoff, H. H. in *Perception as Bayesian Inference* (eds. Knill, D. C. & Richards, R. W.) 123–161 (Cambridge Univ. Press, 1996).
5. Kersten, D. in *The New Cognitive Neurosciences. 2nd Edn.* (ed. Gazzaniga, M. S.) 353–363 (MIT Press, Cambridge, Massachusetts, 1999).
6. Geisler, W. *Psychol. Rev.* **96**, 267–314 (1989).
7. Liu, Z. & Kersten, D. *Vision Res.* **38**, 2507–2519 (1998).
8. Bloj, M. G., Kersten, D. & Hurlbert, A. C. *Nature* **402**, 877–879 (1999).
9. Saunders, J. A. & Knill, D. C. *Vision Res.* **41**, 3163–3183 (2001).
10. Mamassian, P., Landy, M. S. & Maloney, L. T. in *Statistical Theories of the Brain* (eds. Rao, R., Olshausen, B. & Lewicki, M.) 13–36 (MIT Press, Cambridge, Massachusetts, 2002).
11. Ernst, M. O. & Banks, M. S. *Nature* **415**, 429–433 (2002).
12. Geisler, W. S., Perry, J. S., Super, B. J. & Gallogly, D. P. *Vision Res.* **41**, 711–724 (2001).
13. Brainard, D. H. & Freeman, W. T. *J. Opt. Soc. Am. A* **14**, 1393–1411 (1997).
14. Schrater, P. R. & Kersten, D. How optimal depth cue integration depends on the task. *Int. J. Comput. Vision* **40**, 73–91 (2000).
15. Geisler, W. S. & Diehl, R. *Phil. Trans. R. Soc. Lond. B Biol. Sci.* **357**, 419–448 (2002).

Wnt signals lead cells down the caudal path

During development, progenitor cells along the rostrocaudal axis of the neural tube are instructed to become forebrain, midbrain, hindbrain or spinal cord. Cells of the caudal brain are believed to arise through reprogramming ('caudalization') of cells that initially show characteristics of rostral brain. On page 525 of this issue, Thomas Edlund and colleagues now report that Wnt signals from the posterior mesoderm are required for caudalization. FGF and retinoic acid also contribute to the induction of midbrain, hindbrain and spinal cord, but previous studies showed that these factors were not sufficient on their own. Despite earlier indications that Wnt signals were involved in specifying caudal brain character, because they have a variety of other functions in development, it was unclear whether their role in caudalization was direct or indirect. Therefore, the present work provides a crucial additional piece in this developmental puzzle.

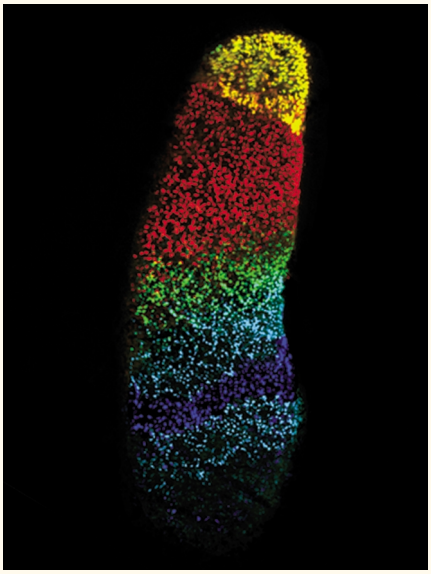

To examine the role of Wnt signaling, the authors used explant cultures of chick neural plate along with immunohistochemical labeling for expression of a combination of transcription factors which selectivity delineate the various brain regions (yellow, rostral forebrain; red/green, rostral and caudal midbrain; light/dark blue, rostral hindbrain). When explants of caudal neural plate were taken from a stage at which the cells still exhibited primarily rostral characteristics and were co-cultured with caudal mesoderm, cells expressed markers for caudal brain regions. If Wnt signaling was inhibited, however, the cells retained their rostral character. If the authors then cultured neural plate explants from a later stage when cells were already specified to eventually make rostral, middle and caudal brain regions, they found that Wnt signaling was still directly necessary for the induction of caudal character. Finally, when explants from the eventual rostral forebrain region of the neural plate were cultured in the presence of FGF and varying concentrations of Wnt conditioned medium, the authors found that increasing concentrations of Wnt resulted in expression of progressively more caudal brain markers. Therefore, the new results firmly establish a direct role for a graded Wnt signal in directing the caudalization of neural plate cells during early neural tube development.

Brian Fiske