

A reprint from  
**American Scientist**  
the magazine of Sigma Xi, The Scientific Research Society

This reprint is provided for personal and noncommercial use. For any other use, please send a request to Permissions, American Scientist, P.O. Box 13975, Research Triangle Park, NC, 27709, U.S.A., or by electronic mail to [perms@amsci.org](mailto:perms@amsci.org).  
©Sigma Xi, The Scientific Research Society and other rightsholders

# Why We See What We Do

A probabilistic strategy based on past experience explains the remarkable difference between what we see and physical reality

Dale Purves, R. Beau Lotto and Surajit Nundy

Visual illusions fascinate people. What we see—whether considered in terms of the brightness of objects, their colors or their arrangement in space—is often at odds with the underlying reality measured with photometers, spectrophotometers or rulers. In the 18th century, the Irish philosopher George Berkeley provided some insight into these discrepancies. In his “Essay Towards a New Theory of Vision,” Berkeley pointed out that the judgment of distance, for example, cannot be derived directly from the geometrical information in the retinal image. Thus, a given line in the retinal image could have been generated by the edge of a physically small object nearby, or equally well by an edge associated with a larger object farther away.

Indeed, all retinal information suffers from this inherent ambiguity. The illumination of objects and the physical properties that determine the amount and quality of light that objects return to the eye are also conflated in the retinal stimulus; thus Berkeley’s general argument applies to sensations of brightness and color, as well as to the perception of space. In each of these basic aspects of vision, the information in the retinal image

cannot directly reveal the true sources of the stimulus in the physical world. As a result, the relation between the world and our perception of it is, by its nature, an uncertain one.

In addition to providing some indication of why what we see might not always tally with reality, this fundamental fact about vision presents a biological dilemma. Survival in a complex and potentially hostile environment clearly depends on responding appropriately to the physical reality that underlies the images on the retina. For example, mistaking a smaller object nearby for a larger but more distant one could obviously be disastrous for a deluded observer. If, however, the image on the retina cannot uniquely define the underlying reality that the observer must respond to, how then does the visual system generate behavior that usually deals successfully with a world that it cannot directly apprehend?

As we show here, a growing body of evidence indicates that the visual system of humans—and presumably many other visual animals—solves Berkeley’s dilemma by generating perceptions on a wholly empirical basis. Rather than analyzing the components of the retinal image as such, percepts are determined probabilistically, using feedback from the outcome of visually guided behavior in the past to progressively improve performance in the face of the inevitable uncertainty of retinal information. The result of this process, and indeed the evidence for it, is that what we perceive accords not with the features of the retinal stimulus or the properties of the underlying objects, but with what the same or similar stimuli have typically signified in both the experience of the species over the eons and the experience of individuals over their lifetimes.

## The Basis of Brightness

The physical intensity of a light stimulus elicits sensations of relative lightness and darkness, which are arguably the most fundamental aspect of vision. Although a sensible expectation is that the sense of brightness should scale directly with the intensity of light, such that a more intense light coming to the eye always corresponds to a stronger sensation of brightness, this is not the case. In fact, two surfaces reflecting the same physical amount of light to the eyes typically look differently bright—or light—if the surfaces are observed in surrounds that are themselves returning different amounts of light. This phenomenon is called simultaneous brightness contrast.

In the past, neurobiologists based the explanation of this well-known effect on the fact that retinal neurons that send information from the eye to the visual part of the brain happen to respond more vigorously to a gray patch in a dark surround than the same gray patch on a light surround—for reasons that have to do with optimizing edge detection. If the firing rate of retinal neurons determined the apparent brightness of the patches, then the patch on a dark background would be expected to look brighter than the same patch on a lighter background.

The problem with this interpretation is that, among other things, patches embedded in scenes that have exactly the same surrounds can also be made to look differently bright. Indeed, as first shown by the 19th-century physiologist Wilhelm von Bezold, a target surrounded by territory of predominantly higher luminance can—under the right circumstances—look brighter than the same target surrounded by territory of lower average luminance. This is just

---

Dale Purves is the George Barth Geller Professor for Research in Neurobiology and professor of psychology and brain sciences at Duke University. R. Beau Lotto is an assistant professor of ophthalmology at University College London. Surajit Nundy is a graduate student in neurobiology at Duke University. The authors are all neurobiologists interested in vision and how the phenomenology of perception is linked to the structure and function of the visual system. Address for Purves: Duke University Medical Center, Department of Neurobiology, Box 3209, Durham, NC 27710. Internet: purves@neuro.duke.edu



Figure 1. Perceptions are often at odds with the measured properties of objects in a scene. In this image, for example, the white tile in the shadow of the table looks brighter than the gray one to its right, even though they are physically identical. This article indicates the reason why these seemingly maladaptive discrepancies between perception and reality occur. (Images courtesy of the authors.)




Figure 2. Several features of light determine what we see. The light that reaches the retina depends on the illumination of an object, how much and what component of the illumination gets reflected, and how the reflected light is transmitted through intervening media, such as air or glass (lower left). In addition, the same projection on the retina can be generated by a smaller object nearby, a larger one farther away or an object at an intermediate distance that is tilted (upper right). These basic facts make all visual stimuli ambiguous.

the opposite of the standard simultaneous-brightness-contrast effect, and the opposite of what the retinal-firing-rate explanation of brightness predicts.

How, then, can these puzzling facts about the relationship between the physical intensity of light and ensuing sensation of brightness be explained? Recall that the identical intensities of light arising from the two surface patches in question are inherently ambiguous. That is, similarly reflective surfaces under the same illuminant or differently reflective target surfaces under different amounts of illumination can generate identical stimuli at the eye. Suppose that this uncertainty is resolved entirely on the basis of past experience with what the source of such stimuli usually turned out to be, determined by the success or failure of the related behavior. Then to the extent that a stimulus of this sort is consistent with past experience of similarly reflective target surfaces under the same illuminant, the targets will tend to appear similarly bright, because things that are the same need to look the same to be behaviorally useful. However, insofar as the stimulus is consistent with the experience of differently reflective objects in different levels of illumination, the targets will tend to appear differently bright, because things that are different need to look different to be useful to the observer. Because the in-






Figure 3. Simultaneous brightness contrast is defined as the same surface looking differently bright in different surrounds. Here the same gray target—the circular patch—looks brighter in a darker surround than in a lighter one (left). The cartoons above show that this stimulus can come from different physical situations: physically identical patches on differently painted surfaces (upper left) and physically different patches under different illuminants (upper right). Because the standard stimulus for simultaneous brightness—the gray targets (left)—contains information that is consistent with both similar surfaces under similar illuminants and different surfaces under different illuminants, the observer sees brightness values that take both possibilities into account.




Figure 4. Cornsweet effect makes identical regions look differently bright when separated by opposing luminance gradients that meet at an edge. Various stimuli, including a spinning black and white disk (a), can create this effect. The region adjacent to the lighter gradient appears brighter than the region next to the darker gradient (b), which is the opposite effect of the standard simultaneous brightness contrast. A graph of relative luminance (c) shows that the two areas on either side of the edge are physically identical, but a graph of the perception of brightness (d) indicates that the right side seems brighter than the left, which is the Cornsweet effect. Blocks showing the Cornsweet effect in a scene (e) enhance the illusion because the information presented increases the probability that the two blocks are differently reflective surfaces under different illuminants. The upper block now looks much darker than the lower one, even though they are identical.

formation in the standard stimulus for simultaneous brightness contrast is consistent with both different surfaces under different illuminants and similar surfaces under similar illuminants, then what the observer sees will reflect both possibilities. In statistical terms, the stimulus is in some degree consistent with surfaces having different reflectances, so the identical patches in the standard demonstration of simultaneous brightness contrast will look differently bright.

This may seem a strange way to generate visual percepts. Given the inevitable uncertainty of the information in the retinal image, however, this strategy may be the best—or even the only—way to resolve Berkeley’s dilemma.

### Complex Stimuli

If this general explanation is correct, then the same perceptual effect should be elicited by any stimulus in which target territories of the same lumi-

nance have typically turned out to be differently reflective objects in different amounts of light. A particularly interesting challenge is the percept generated by a more complex stimulus called the Cornsweet edge, which is named after Tom Cornsweet, the psychologist who described this effect in the late 1960s.

In the Cornsweet effect, opposing luminance gradients that meet at an edge make physically identical adjoining regions look differently bright. Specifically, the region adjacent to the lighter gradient appears brighter than the region next to the darker gradient. Because this perceptual effect is the opposite of the effect of standard simultaneous brightness contrast, the Cornsweet stimulus provides yet another example of why explanations based on local contrast relations do not work.

Despite its seemingly complicated structure, the Cornsweet-edge effect can also be explained in empirical

terms. The common denominator of the Cornsweet stimulus and conventional simultaneous-brightness-contrast stimuli is that the percepts can both be understood in terms of the possible sources of the physically identical target territories. Thus, the equiluminant regions bordering the gradients that comprise a Cornsweet edge could have been generated by similarly reflective surfaces under the same illuminant—including painted gradients on the surface of a piece of paper on which light falls uniformly—or differently reflective surfaces under different intensities of illumination—including a cube with rounded edges placed so that one side is in light and another in shadow. Since both scenarios, and a host of others, are real possibilities, the percept elicited by the stimulus will, according to a wholly probabilistic theory of vision, take all the possible sources into account in proportion to their occurrence in the past. Given that




Figure 5. Various scenarios can create a Cornsweet-edge stimulus. Aligning gradients derived from equally illuminated pieces of paper (top) or gradients from differently reflective cubes with rounded edges placed so that one side is in light and another in shadow (bottom) both generate the same Cornsweet edge. Observers again see the effect because both possible sources are incorporated into the perception of the ambiguous stimulus.

the stimulus will often have been generated by differently reflective surfaces in different illumination, such as the cube scenario, the target territories will look differently bright.

If this statistical explanation based on past experience has merit, then the perceptual effect of the Cornsweet edge should be increased, decreased or even abolished simply by altering the relative probabilities of the possible sources of the stimulus without chang-

ing the stimulus as such. As various experiments show, this is the case.

#### Seeing Colors

Given that these otherwise puzzling aspects of the sensations elicited by the intensity of light can be understood as a consequence of a wholly probabilistic strategy of vision, we wondered whether the color sensations elicited by different light spectra might also arise in this way. After all, the distribution of spectral

power in a light stimulus, which is what gives rise to sensations of color, is ambiguous for exactly the same reasons as is the overall spectral intensity. Illumination, reflectance and other factors that determine the characteristics of the light that reaches the eye are inevitably intertwined in the retinal image and cannot be disentangled.

A good starting point in thinking about color sensations in these terms is simultaneous color contrast, a phenomenon similar to the brightness contrast effects already described. Two targets with the same spectral composition placed on differently colored backgrounds serve as standard stimulus for eliciting color contrast. As in brightness contrast, the two targets look different, but now in terms of their respective color qualities: hue, saturation and color brightness. In the past, most explanations of this phenomenon relied on some sort of color averaging across the entire stimulus. As in brightness contrast, however, these schemes fail to account for the fact that color contrast stimuli can be crafted in which the same average chromatic surrounds elicit different color percepts.

An explanation of color contrast can, however, be given in empirical terms. The sources of the target and surround in standard color-contrast stimuli are profoundly uncertain, because an infinite number of combinations of reflectances and illuminants—and other less crucial factors—can generate the same distributions of spectral power. As in the case of achromatic stimuli, the visual system could resolve this dilemma by using feedback from the success or failure of the past behavioral responses to spectral stimuli. The percept elicited by a given stimulus would thus be determined by the relative frequencies of occurrence of the real-world combinations of reflectances and illuminants that gave rise to that distribution of spectral power in the past. The same argument can be applied to color constancy, which describes the related phenomenon in which the same object continues to appear similar in color despite being under different illuminants.

If perceptions of color contrast and constancy are generated in this way, then the same spectral target on two differently chromatic backgrounds would be expected to give rise to different chromatic sensations. The reason is that, in addition to requiring behaviors appropriate to the same reflectances in





Figure 6. Simultaneous color contrast occurs when the same surfaces placed on differently colored backgrounds appear to be different colors. Here the same chromatic target looks reddish on a blue background and more purplish on a redder background (left). The cartoon above shows the ambiguity of the stimulus, which could arise from physically identical targets (upper left)—shown alone on the pedestals—placed on differently colored surrounds under equal light or from physically different targets illuminated with differently colored light (upper right). Again, simultaneous color contrast arises because the visual system incorporates into perception these different possible sources according to the frequency of their occurrence in the past.








Figure 7. Color contrast and constancy arise from the same empirical generation of visual perceptions. In this computer-generated example, the authors carefully controlled the spectral information in the scene. The upper images show the cubes as if in yellowish (top left) or bluish (top right) illumination. The lower images show specific tiles of interest in the absence of these contexts. The yellow-looking tiles depicted as if under blue light and blue-looking tiles depicted as if under yellow light are actually a gray on their own—as shown in the lower boxes marked “blue” and “yellow.” This is a striking example of color contrast. On the other hand, the red-looking tiles depicted as if under blue and yellow light both come from tiles that are actually purplish and orangish, respectively, as shown in the lower boxes marked “red.” This demonstrates color constancy. These remarkable effects show that the same targets can be made to look like very different colors, and that different colors can be made to look the same by manipulating the context.




Figure 8. Perceptions of angles do not usually correspond to the angles of the underlying objects. An angle projected on a surface, such as the retina, can arise from objects having a variety of angles and arm lengths, arranged in infinitely many three-dimensional orientations. The three angular objects shown here have subtenses of 120 (left), 90 (middle) and 60 (right) degrees and various arm lengths, but they can be arranged to make identical projections, as shown by their shadows.

the same illuminant, such stimuli would in other instances have required behaviors appropriate to targets that arise from different reflectances in different illuminants. Consequently, a spectral stimulus should elicit a sensation that incorporates all possible underlying sources in proportion to their past occurrence in human experience.

To get across the merits of this way of understanding color percepts, we devised a stimulus that looks something like a Rubik's cube. By understanding the effects of spectral differences in this probabilistic way, we could generate color contrast and con-

stancy effects that are much more dramatic than the usual textbook illustrations of these phenomena. For example, when all the information in a scene with the cube was made consistent with either yellowish illumination or bluish illumination, tiles on the surface of the cube that appear the same shade of gray in a neutral context could be made to appear either blue or yellow, respectively. This manipulation provides an impressive example of color contrast made especially dramatic by empirical manipulation of the information in the scene. Conversely, tiles that appear differently colored in a neutral

setting could, by changing the probability of their possible sources, be made to look the same color, thus providing an equally dramatic demonstration of color constancy. These demonstrations show not only that color contrast and constancy are determined probabilistically, but also that these seemingly opposite effects are both manifestations of the same empirical generation of visual percepts.

### Perceiving Scene Geometry

Vision scientists noted long ago that the perception of lines does not always accord with the real-world geometry of the underlying objects. For instance, the angles formed by lines making, or implying, an acute angle are seen as being a few degrees larger than they really are, whereas obtuse angles are seen as being a few degrees smaller. Despite a great deal of speculation about this anomaly dating from the latter part of the 19th century, there has been no consensus about its origin. Thus we asked whether these and other geometrical misperceptions might be explained in much the same empirical terms as brightness and color.

Much like luminance or spectral power, the stimulus that gives rise to a perceived angle is profoundly ambiguous. An angle projected onto a surface—the retina, for example—can arise from objects having a variety of measured angles and arm lengths, arranged in infinitely many three-dimensional orientations. In interacting with the objects that give rise to particular angle projections on the retina, observers through the ages would have experienced great variation between a given angle in retinal projection and the angles of its real-world sources. Moreover, the variations are systematic. In consequence, the perceptions elicited by different angles projected onto the retina would be expected to correspond to these frequency distributions.

To test this interpretation, we first needed to determine the probability distribution of all the possible three-dimensional sources of a projected angle. When these distributions for all possible angles were computed using the principles of projective geometry, we found that acute-angle projections usually come from sources with larger angles than the projections. Conversely, the sources of obtuse-angle projections are typically generated by sources that are somewhat smaller than the project-

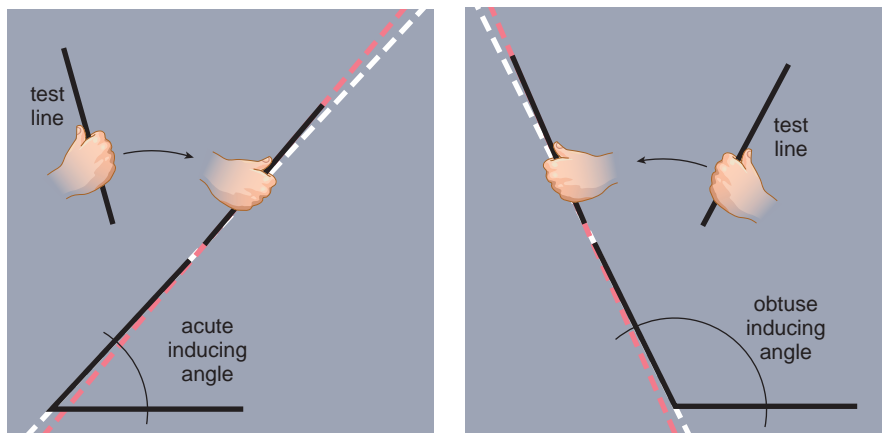


Figure 9. Systematic errors arise in judging angles. The distributions of the possible sources for all possible angles reveal that acute-angle projections usually come from sources with larger angles than the projections, and that obtuse-angle projections are on average generated by sources that are somewhat smaller than the projected angle. The visual system should incorporate these statistical facts. To test this, the authors asked subjects to line up a test line with one arm of an “inducing angle.” With an acute inducing angle (left), subjects oriented the test line as if the inducing angle was larger than it really was. With an obtuse angle (right), subjects oriented the test line as if the inducing angle was smaller than reality. Overall, the results reflected the probability distribution of the possible sources of the stimuli. Like the perception of brightness and color, the angles we see are evidently based on the empirical significance of the stimulus, not its physical dimensions.



ed angle. Right-angle projections and straight lines are generated by sources that, on average, have the angular subtense of the object itself. If percepts are determined empirically, then the visual system should generate perceptions of angles that incorporate and reflect these statistical facts of projective geometry.

We assessed this prediction by asking subjects to report their perceptions of different angular stimuli in a series of tests in which the adjustment of a test line indicated the angular subtense they were actually seeing. For example, if the subject perceived the angle to be bigger than it actually is, then the test line would be set in a position that revealed this discrepancy—being not quite parallel with the angle's arm. The results derived from such tests tallied quite well with the probability distribution of the possible sources of the corresponding stimuli, indicating that the spatial arrangement observers see is neither the retinal projection nor its real-world source, but its empirical, or past, significance.

Taken together, this evidence drawn from the perception of brightness, color

and geometry supports the idea that the problem first emphasized by Berkeley is resolved by generating visual percepts according to the probability distribution of the possible sources of the visual stimulus, whatever it may be. As a result, observers see what a visual scene typically signified in the past, rather than what it actually is in the present. We see what we do, therefore, because the statistics of past experience is the basis on which the visual system contends with the dilemma posed by the inherent ambiguity of visual stimuli.

#### Acknowledgments

We are grateful to Mark Williams for assistance with many of the figures and to the National Institutes of Health for support.

#### Bibliography

- Lotto, R. B., and D. Purves. 1999. The effects of color on brightness. *Nature Neuroscience* 2:1010–1014.
- Lotto, R. B., and D. Purves. 2000. An empirical explanation of color contrast. *Proceedings of the National Academy of Sciences USA* 97:12834–12839.
- Lotto, R. B., and D. Purves. 2002. A rationale for the structure of color space. *Trends in Neuroscience* 25:84–88.

Nundy, S., R. B. Lotto, D. Coppola, A. Shimpi and D. Purves. 2000. Why are angles misperceived? *Proceedings of the National Academy of Sciences USA* 97(10):5592–5597.

Purves, D., A. Shimpi and R. B. Lotto. 1999. An empirical explanation of the Cornsweet effect. *Journal of Neuroscience* 19:8542–8551.

Purves, D., R. B. Lotto, S. M. Williams, S. Nundy and Z. Yang. 2001. Why we see things the way we do: Evidence for a wholly empirical strategy of vision. *Philosophical Transaction of the Royal Society of London B* 356:285–297.

Williams, S. M., A. N. McCoy and D. Purves. 1998. The influence of depicted illumination on perceived brightness. *Proceedings of the National Academy of Sciences USA* 95:13296–13300.

Williams, S. M., A. N. McCoy and D. Purves. 1998. An empirical explanation of brightness. *Proceedings of the National Academy of Sciences USA* 95:13301–13306.

Links to Internet resources for further exploration of “Why We See What We Do” are available on the American Scientist Web site:

<http://www.americanscientist.org/articles/02articles/purves.html>