WILLIAM R. POWELL

REALITY, PERCEPTION, AND SIMULATION: A PLAUSIBLE THEORY

Vision is such an automatic process that few people think about it. In this article, I examine some new and well-known visual phenomena and suggest a framework for understanding them. I conclude with some speculation about the nature of experienced reality and the classical philosophical problem of free will. The proposition advanced is that all perception is contained within a simulation of physical reality created by the brain and, further, that this simulated reality has at least as strong a claim to being "real" as the physical reality that is inferred on the basis of the brain's simulation and analysis.

INTRODUCTION

When most people think about vision, they are usually concerned only with how an image is focused on their retinas. But that part of the visual process is relatively simple. Even the electrochemical process that converts the photons of light into a two-dimensional (2D) pattern of neural activity in retinal neurons is simple in comparison to what is required to convert those signals (retinal cell output pulses) into a visual experience of recognized objects in a perceived three-dimensional (3D) world. To avoid immediately becoming philosophical, I assume that there is a 3D world "out there," and that I somehow represent it inside my brain with reasonable fidelity under most conditions. The central problem of visual perception is: How does one take the 2D array of information present in the retina and process it to form a 3D perception of the physical world? This problem is too complex to address directly, so I will consider several smaller, but related, problems.

PHYSIOLOGICAL PROBLEMS IN PERCEPTION

I ignore here all the processes that focus the eye and select one direction of viewing with a common point of fixation for both eyes. We'll begin with the retina. When we examine the retina, we notice that it seems to be built backwards. That is, the photosensitive cells are on the back surface. Thus, light must pass through two distinct layers of processing cells, pass by the nerves that fan out over the retina from the optic nerve to collect signals from these processing cells, and pass around all the blood vessels of the retina before it can reach the photosensitive cells on the back surface of the retina.[1]

Why don't we see the shadows these blood vessels cast as the light passes them on its way to the photosensitive cells? When the eyes move from one fixation point to another, why doesn't the world appear to be a moving blur? And, after we fix on a new point in space, why doesn't the world appear to lurch as object images jump to new locations on the retina several times each second?

Yet another puzzle: How does the brain recognize that an object that is activating a set of neurons in one small part of the periphery of the retina on one fixation is the same object that was creating different neural activity in another part of the retina during the previous fixation? The problem becomes even more complex when we recognize that the shadows of the retinal blood vessels fragment the object images into many different parts with different shapes on each successive fixation.

Consider also that the density of nerve cells sampling the retinal images decreases rapidly as the image is displaced from the fovea, that very small region of the retina where the image is most highly resolved.[2] Yet, the entire perceived world appears to be in the same sharp focus with full detail even though the resolution or density of neurons in the periphery of the eye is much lower than in the fovea.

A consistent explanation for all these phenomena seems to force the assumption that we do not really "see" the image of an object on our retinas but, instead, construct some "inner vision" using the database of temporal and spatial information reflected to our eyes as optical energy by objects in our environment. That is, we do not really see the image in our eyes when we experience a visual percept; visual perception is a creation of the brain, not the eyes. This view is similar to Neisser's contention that a conversation is understood not with the ears, but by active construction in the brain.[3]

This premise immediately suggests an interesting question: Is the retinal data array the only database from which we construct this inner vision or visual perception? When we dream, we are "forming a visual perception" or "having a visual experience" or "seeing" with our eyes closed. People hallucinate visual experiences even in conflict with retinal images. Is it possible that the blind can see—that is, have a visual experience or form a visual perception?

Clearly, previously sighted blind persons can still generate inner visions or remember scenes, just as I can when I visualize the front of my house. However, without a retinal image of my house present in my eyes, I cannot construct a high-resolution, full 3D representation of it. No matter how hard I try, I cannot make myself think I am looking at my house unless I am in front of it with my eyes open. Only the eyes, with their millions of densely packed receptors accessible to external stimuli, can provide the large and constantly changing database required for the construction of normal inner vision.

Thus, the visual percept or inner vision possible with functioning eyes is much richer and more complete than that created otherwise, but the process of visual perception in the cases just discussed is fundamentally the same. Indeed, there is evidence that the blind can see, albeit with very low resolution.

Twenty years ago, researchers conducted experiments in which they used computer-controlled probes (400 or less) to stimulate the tactile nerves on the backs and bellies of blind subjects.[4] The degree of stimulation was directly related to the intensity of an image presented to an equal number of electronic photosensors in the focal plane of a camera. (A miniature TV camera was mounted on one side of the subject's head.) Initially, the stimulation was experienced only as sensations on the skin, but after a few hours of freely moving in the experimental area, the blind subjects' perception changed. They perceived not the stimulation, but fuzzy (low-resolution) images or scenes corresponding to the TV camera's view. Even more startling, the subjects reacted to a sudden zoom of the TV lens by ducking! That is, when the experimenter zoomed the TV lens in without a blind subject's knowledge, so that distant objects in the TV scene suddenly grew larger as if thrown toward the subject's head, the subject ducked to avoid being struck, just as a normally sighted person would. Yet, the only actual change was in the stimulation of the subject's back.

The change in the blind subjects' perception from an experience of tactile stimulation to a low-resolution visual experience is analogous to the perception of a blind person using a cane. He or she does not perceive vibration and pressures in the hand but the world of objects located at the tip of the cane. It is quite possible that technology will some day provide low-resolution electronic "tactile eyes" for the blind.

PSYCHOLOGICAL PROBLEMS IN PERCEPTION

Often the perceived characteristics of viewed objects are in direct conflict with physical characteristics of the images formed on the retina. When a man stands three times farther away than his half-grown son, the image of the man on the viewer's retina is smaller than that of his son, yet the viewer clearly perceives that the man is larger than the boy (size constancy problem). When we look at a phonograph record, we recognize it as a circular disk, yet its retinal image is an ellipse except for one unusual viewing angle (shape constancy problem). A green leaf viewed through rose-colored glasses still appears to be green (color constancy problem).[5] When we mistakenly

recognize an approaching stranger as a known friend, we maintain this perception until it suddenly collapses.

Likewise, when we view an intentionally ambiguous illustration, such as the well-known Necker cube shown in Figure 1A, we can only see one version at a time (uniqueness problem). Most urban-dwelling people viewing the Necker cube do not perceive the two triangles, four trapezoids, and square that compose it. The 3D perception experienced is a construct of the brain, not an accurate representation of the actual 2D object presented to the viewer. On the other hand, Figure 1B is ambiguous in its dimensionality—is it a 2D or 3D object? This ambiguity is a perception error at least as problematic as the depth reversal for which the Necker cube is better known.

Our ability to perceive a movie as a 3D world, not a flat screen, is even more amazing when we realize that we are viewing the movie from a different perspective than the camera that filmed it. For example, if the scene filmed contains a long, rectangular box with a square cross section placed symmetrically to extend across the entire film stage, and the camera is in line with the left end of the box, as shown in Figure 2A, then the right end of the box recorded on the film will be shorter than the left edge (Figure 2B, top). Yet, a viewer in any location in the theater perceives the box as an object with equally large ends.

According to the accepted theory of size constancy, the movie viewer on the left side of the theater automatically increases the perceptual size of objects at greater distance. The viewer symmetrically seated on the right relative to the camera position of Figure 2A will be just remote enough from the projected left end of the luminous trapezoid on the screen to exactly compensate for the size difference recorded on the film, and will thereby receive a rectangular retinal image. But the movie goer viewing a luminous (screen) trapezoid whose near end is actually smaller than the far end doesn't automatically increase the size of the remote left end to achieve size constancy, as in the "computed" perception of Figure 2B.
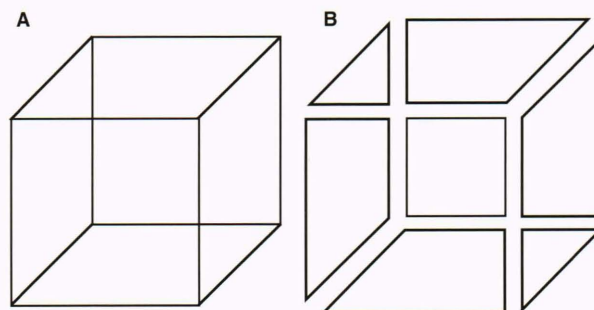
Figure 1. A. The Necker cube. **B.** Geometric component structure of the Necker cube. These deliberately ambiguous two-dimensional structures are perceived as bistable three-dimensional objects. The representation in part A appears to reverse depth, whereas the one in part B can be perceived as either two-dimensional or three-dimensional. Only one version can be seen at a time, demonstrating the "uniqueness" feature of perception.

A

B

Figure 2. A. Schematic of a movie theater showing positions of a box and camera during filming of the box, and the screen, projector, and viewer during viewing of the resulting film. **B.** Although the right end of the box as filmed and projected is smaller than the left end (top two views), a viewer seated anywhere in the movie theater automatically adjusts his or her perception to see a symmetrical box (bottom view), demonstrating the size constancy feature of perception.
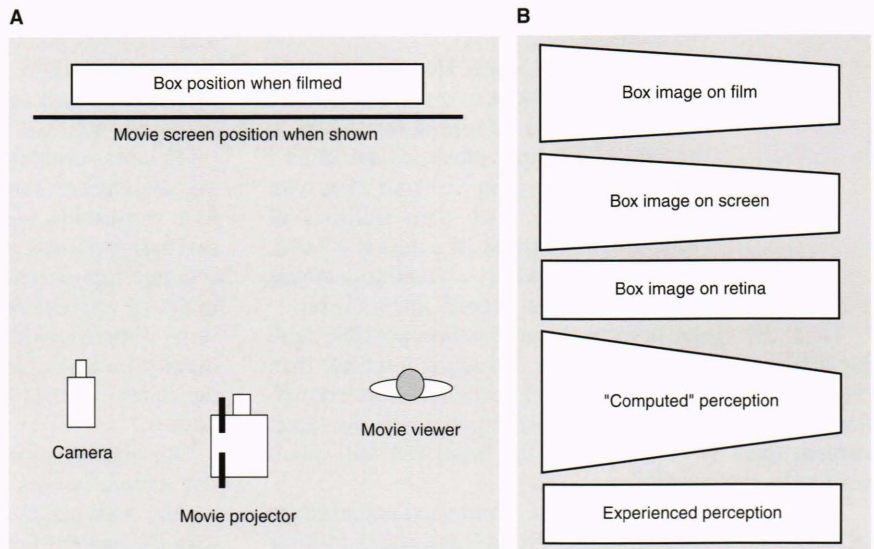
Instead, he or she *reduces* the perceptual size of the remote left end (trapezoid to rectangular) to perceive a rectangular box.

No one can seriously maintain that purchasing a movie ticket and sitting in a dark room deactivates the automatic computational processes that are conventionally evoked to explain size and shape constancy. Clearly, Descartes had it right and the modern theory of vision as the emerging end result of automatic computational transforms of retinal data is wrong. We construct our visual percepts from our "opinions and knowledge,"[6] although this statement does not explain how we perform this amazing feat.

Other illusions present other problems. For example, in the illusions of Figure 3, we perceive a vertical boundary when only horizontal lines are present (Fig. 3A) and a rectangle when no rectangle is present (Fig. 3B).[7]

COGNITIVE EXPLANATIONS

Illusions and other visual "defects" can help us to understand how the visual system is working.[8] But first, we must agree upon what constitutes an explanation of perception.

The cognitive scientist would insist that we may not explain a process by giving it a new name or restating the observed process in other terms. (To tell someone that morphine makes him sleepy because it contains a narcoleptic agent is hardly an explanation.) Instead, we must break complex processes down into sets of simpler processes and further refine each of these until we have identified processes that can conceivably be accomplished by sets of nerve cells. To illustrate this operation, let us consider an explanation of how we perceive a horizontal line and distinguish it from a vertical line and other objects.

Neurological Processes

Many retinal cells can stimulate a single cell in the visual cortex. The geometric pattern of retinal cells that stimulates a particular cortical cell is called its receptive field.[9] For example, Figure 4A illustrates a receptive field
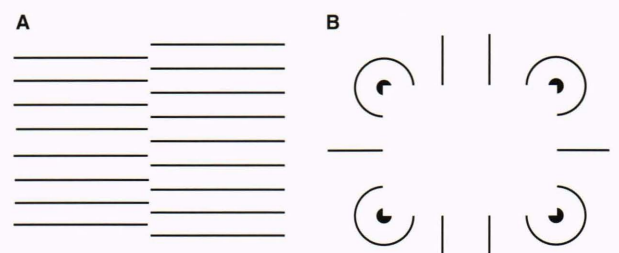
A

B

Figure 3. Illusory contours. **A.** Illusion of a vertical line. **B.** Illusion of a rectangle.

of retinal cells arranged as a vertical ellipse, all of which stimulate a hypothetical cortical cell. The nerve cells are represented as uniformly dense inside the ellipse only because of drawing limitations; real neurons would be more irregular. Figure 4B represents a set of horizontally arranged cells that stimulate a different cortical cell. Clearly, a horizontal line centered on the receptive field of the cells in Figure 4B will stimulate the corresponding cortical cell to a higher rate of output activity than a horizontal line on the field in Figure 4A (Fig. 4C). However, many other illumination patterns will stimulate both cells, e.g., Figures 4D and 4E. Consequently, a high rate of discharge by the cortical cell stimulated by the field in Figure 4B is not a reliable indicator that a horizontal line of illumination is present.

The output of one cell can also decrease the activity of another, i.e., can inhibit it. Therefore, if an inhibitory receptive field of the retinal cells in Figure 4B surrounds the excitatory field ("center on, surround off"), and these areas are approximately equal in influence (Fig. 5A), then the circular and square illumination patterns illustrated in Figures 4D and 4E would not significantly activate the corresponding cortical cell. But a horizontal line of illumination would still activate it because only a small portion of the inhibitory receptive field is illuminated (Fig. 5B). (Light is represented by black areas in these illustrations.)
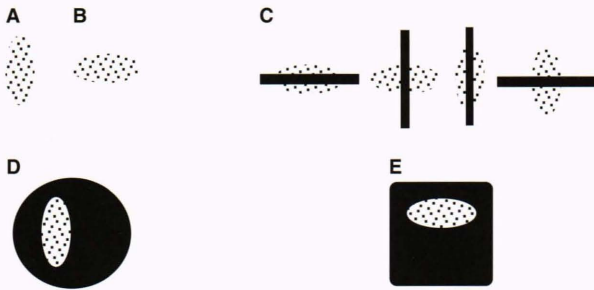
Figure 4. A.,B. Retinal cell patterns comprising vertical and horizontal receptive fields stimulate different cortical cells. **C.** Horizontal and vertical illuminated line segments produce different patterns of stimulation for horizontal and vertical receptive fields. **D.,E.** Both vertical and horizontal receptive fields will respond to some illumination patterns.
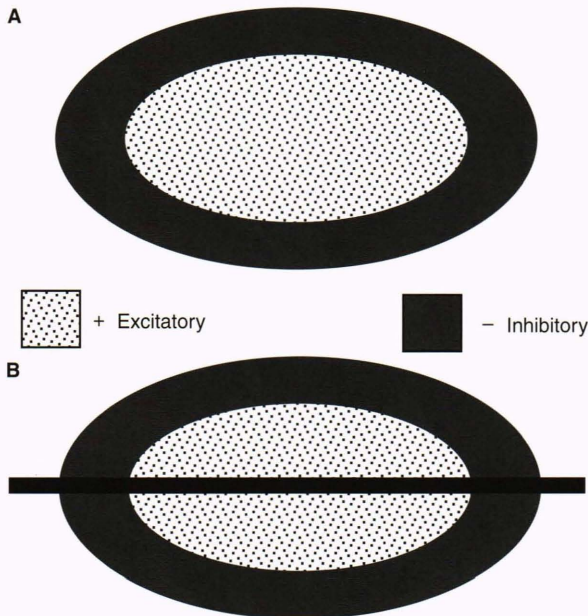

Figure 5. A. An inhibitory receptive field surrounding an excitatory receptive field ("center on, surround off") can decrease the activity of the corresponding cortical cell. **B.** A horizontal line of illumination activates more excitatory than inhibitory cells in a horizontal "center on, surround off" receptive field pattern.

Figure 6A illustrates another type of retinal receptive field, one that activates a cortical cell that is sensitive to edges in the retinal image. Receptive fields of specialized detector cells exist in different shapes and sizes in the retina but are relatively simple, consisting primarily of the oriented edge detectors illustrated in Figure 6 and the "center on, surround off" cells illustrated in Figure 5.

Hubel and Wiesel received the Nobel Prize in 1981 for their work on these feature-detector cells.[10] Even more specialized cells have been discovered. Popper and Eccles[11] reported research in which very small electrodes were placed inside brain cells of alert monkeys or cats and the most excitatory shapes were determined. For example, cells with maximal response to the shape of a hand have been discovered, even though there is no

evidence for receptive fields in the retina shaped like hands. These very complex feature detectors operate by processing the outputs of other, less complex, cortical cells.

Edge-detector cells, such as that illustrated in Figure 6A, are sensitive to the direction of contrast. In other words, the cortical cell activated by such a receptive field responds strongly when the left side of the receptive field is illuminated and the right side is not. If the receptive field of this cell were illuminated on the right and not on the left side, then the cell's activity would essentially stop. Opposing contrast edge-detector cells can be combined to make boundary-detection cells that are not sensitive to the direction of contrasts. Although Figure 6A represents the receptive field of a vertical-edge detector, the visual cortex contains edge detectors for all different orientations to process information from each location in the retina.

Individual retinal cells can stimulate many cells in the visual cortex, usually via the lateral geniculate nucleus, a primary relay point in the brain where retinal ganglion cells synapse with neurons that have projections to the visual cortex. The same retinal cell could, for example, be part of both a horizontal- and a vertical-edge detector's receptive field. It is the more central cortical cell, not the retinal cell, that indicates the orientation of the edge. Note that the illumination edge does not need to be precisely aligned with the orientation of the receptive field to produce some activity. Also, a strong contrast boundary slightly misaligned with the preferred orientation of an edge-detector cell could produce more activity in that cell than a perfectly aligned weaker boundary. Thus, the meaning of any given level of activity in a particular edge-detecting cell is somewhat uncertain. However, some edge-detector cell (probably many) corresponding to the same retinal location will be perfectly aligned with the illumination edge and will be even more active. Since the edge must have some unique orientation, these well-aligned edge-detector cells will attempt to suppress activity in other, worse-aligned edge detectors corresponding to the same retinal location.

This competition tends to sharpen the response characteristics and make the edge detectors more accurate indicators of the orientation of the illumination edge present. Likewise, all edge-detector cells with the same preferred edge direction tend to stimulate each other and nearby detectors of the same or nearly the same orientation. Before considering this mutual stimulation of similar adjacent edge detectors, let us address the related question of how we "know" which active edge-detector cells to associate or bind together to form the continuous edge of the objects we see.

Binding Problem

When we view a flat black-and-white drawing, all our brains have to work with is a 2D array of separate, active edge-detector cells. No one is directing how these cells should be grouped together to form the boundaries of separate objects. For example, many different parts of the retina are illuminated by edge-contrast boundaries as we look at Figure 7.

Count the number of objects in Figure 7 and then the number in Figure 8. Although one object was added to Figure 7 to create Figure 8, we now perceive only 7 objects in Figure 8 versus 15 in Figure 7. We recognize the "L" objects in Figure 8, probably because like-oriented edge detectors in close proximity tend to reinforce each other,[12] especially when there is another adjacent edge that can be associated with a closer opaque object by the same mutual reinforcement process. Thus, we associate the lower boundary of the top horizontal rectangle and the upper boundary of the lower horizontal rectangle of Figure 7 with the exterior of the box seen in Figure 8 rather than with these two rectangles. Looking at Figure 7, we have no other way to form closed contours for all of the objects. In Figure 8, however, we can see that these same two horizontal boundaries are part of the exterior boundary of the box. In the absence of depth data (such as stereoscopic information), we associate the largest scale closed boundaries together to define the outline of an object. We then perceive this object as being in front of other objects that can share a boundary with it.

Multiple-microelectrode studies have recently produced new evidence for mutual reinforcement by demonstrating separated neurons with synchronized discharges or firing patterns in the 40- to 70-Hz range.[13] Synchronized firing patterns may be part of the brain's mechanism for assembling concurrent environmental patterns and events into coherently perceived objects and actions. Thus, in Figure 8, we may perceive that the various parts of a fragmented "L" extend behind the box to make a unified "L" because the detectors for the missing portions of the "L" boundary (that is, those behind the box) are excited by the nearby and generally aligned portions of the "L" boundary that are visible. The firing patterns of these hidden-edge detectors are probably synchronized with the cells responding to visible boundaries. The separate parts of the "L's" cannot form closed-boundary contours in Figure 8 because the box has claimed part of the boundary that allowed us to perceive closed contours around the smaller fragments of the "L's" in Figure 7.

This concept is illustrated in Figure 9, where the activity of the hidden-boundary detectors is illustrated explicitly by white lines across the box.

Evidence is also available that receptive fields come in several different sizes. We know that the larger-scale feature detectors of the visual cortex respond more rapidly than those for finer detail,[14] which may explain why, when presented with a flat drawing such as Figure 8, we see the boundary of the larger box object first; then, higher (more central) processes accept the weak activity (white lines in Figure 9) as boundaries unifying the "L" objects.

Our minds have no way to know what is present out there in the world. We can make judgments (calculations) based only on the pattern of neural activity present in our brains. If a set of boundary detectors is active and discharging synchronously, then it is not surprising that the brain judges, for example, that the "L" fragments in Figure 8 continue behind the box. The box also produces synchronized boundary oscillations, but not with the same phase or frequency as the "L's."

What I am suggesting is far from proven. Basically, I propose that when a set of boundary detectors can form a closed contour and synchronize their firing patterns by mutual reinforcement, then the area inside that contour will be taken as part of the surface of some unique object. When cells indicating a contrast boundary can oscillate synchronously with either of two different sets of boundary cells, as is illustrated by the white lines in Figure 10 for the boundary parts shared by both the box and the "L's," then larger aggregates of cells appear to lock together in synchronous oscillation first.

Figure 11 presents two different versions of the well-known Kanizsa square.[15] Note that the illusion of a square is present even in the reversed contrast version (right), indicating that boundary cells extend the boundaries (discussed earlier) rather than edge-detector cells that are sensitive to the direction of contrast. The effectiveness of the line-only version in Figure 3B in producing the illusion of a rectangle, even though there are no boundaries (on a large scale) to extend, indicates that other mechanisms are also at work. The nature of these processes and a possible neurological explanation are provided by Grossberg.[16]

Recognition of Objects

When we see a horizontal line, we are aware only of the horizontal line; we are not aware that some set of neurons that are selective for horizontal lines is actively discharging. Nevertheless, such neurons are indeed activated when a horizontal line is presented to the retina in the position associated with this set of neurons. Higher-level sets of neurons process the activity of this selective set of neurons to generate the perception of the horizontal line. In other words, we perceive not a retinal image, with its shadows of blood vessels, nor the specialized feature detectors in the visual cortex. Rather, these parts of the visual system collect data, analyze it into features, and transform it into other forms used by more central neurons[17] to construct objects in our own personal 3D representational worlds. It is these constructed 3D represen-
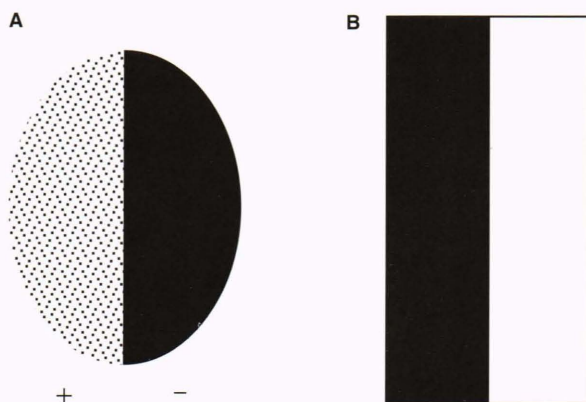

Figure 6. A. A vertical edge-detector receptive field with excitatory (+) and inhibitory (–) regions activates a cortical cell that is sensitive to edges in the retinal image. **B.** Example of an illumination pattern (dark area represents illumination) that would maximally activate the edge detector in part A.
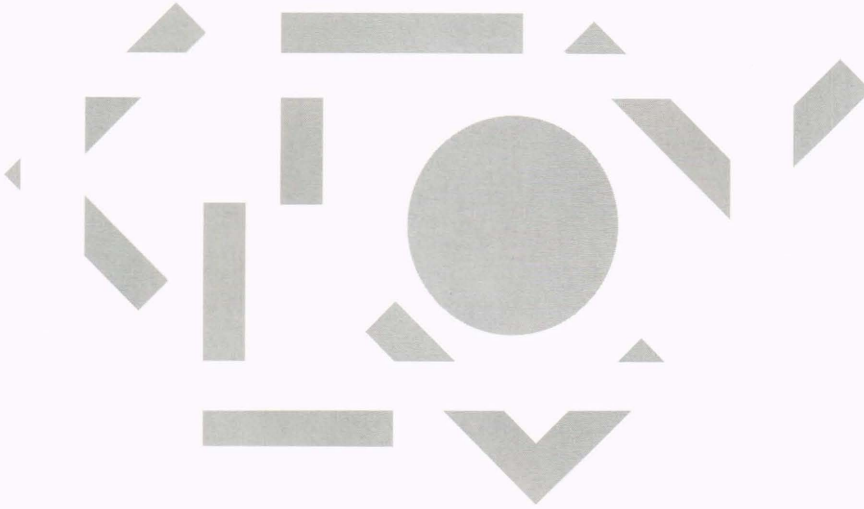
Figure 7. A two-dimensional array of seemingly disparate geometric shapes activates various sets of edge-detector receptive fields. **Count the number of objects seen before turning the page.**

tational worlds that we perceive. Thus, each of us lives in a world of our own construction, although evolutionary selection has caused us to construct our personal world so that it closely resembles the physical world. Figures 1, 3, and 11 are examples of conflict between the perceptual world and the physical world.

How do we recognize an object once lower-level neurons have separated it from other objects and the background? Let us consider something a little more complex and interesting than a horizontal line, such as a fully visible black cat. Recognition of a multicolored cat, partially hidden behind some grass, is much too complex to explain, but obviously our visual systems can usually cope. (If they could not do so most of the time, we would not exist because too many of our ancestors would have been eaten by tigers.)

We could logically assume that since we have seen black cats before in many different sizes and in many different positions, we simply store in memory all these different images of black cats. Then, when we need to identify the black object that is activating a special combination of lower-level neurons, we compare the outputs of this combination of active neurons to image sets stored in memory. This concept is almost surely wrong.

Since we do not initially know that the black object we need to identify is a black cat, we would need to check our memory for all black things that we have seen, including, for example, black telephones and black tires. In addition, we would have to consider each of the black objects in all of the possible views stored in our memories. Also, we normally cannot restrict our search of memory to just images of black objects, since many things, including cats, come in several colors or in many different solid colors. Clearly, the storage capacity in the finite volume of a human brain is not large enough to store all views of everything that we have ever seen.

Instead, we must do some form of analysis (neural computation) that leads to the selection of "cat" and all the things we know about cats for inclusion in our 3D representation of the world. In this process, we respond to the physical world by identifying the stimulus as a cat

and add a cat representation to our personal 3D representation of the world, resulting in our perception and belief that there is a cat present in the physical world. No one is sure how we create this view. I propose one plausible object identification procedure, and illustrate it with the following experiment.

For his work on motion perception, Johansson[18] affixed a set of small lights to a dancer's ears, shoulders, elbows, wrists, hands, hips, knees, ankles, and feet and had him dance on a dark stage. Viewers had no trouble in following the dance movements because the distances between some pairs of lights (e.g., the knee and ankle lights) remained constant regardless of how the dancer moved. Thus, viewers had enough data about the position and orientation of the "natural" segments of the dancer (or object) to identify it. If the lights had not been at the major joints of the body, viewers would probably not have been able to recognize the figure as human.

Natural divisions in objects often occur where the surface characteristics, such as color or texture, change abruptly. Camouflage paints exploit this feature to hinder recognition. (I will explain the natural division of objects with uniform surfaces later.)

In much the same way as for our dancer, we may recognize a cat by a procedure that first parses the cat into major components, such as the head, torso, legs, feet, and tail (Fig. 12). Prior to identification of the object, we don't know that the cat is a cat. The identification procedure must be able to subdivide any object into subcomponents that are usually present when the object is in full view and properly associate these subdivisions. That is, the cat's head must be directly associated with the neck, not the leg. The relative sizes of the various subcomponents, as well as their interconnections, must also be retained during the subdivision process. All three—relative locations, relative sizes, and interconnections—are necessary for identification.

I suggest that we do not store any image of any cat in memory. Rather, we store the prescription for constructing a cat's image, i.e., a set of cat construction rules that is flexible enough to allow construction of images of cats

Figure 8. When a dark rectangle is superimposed on the array of geometric shapes shown in Figure 7, the viewer now perceives a rectangle surrounding a circular object and occluding five L-shaped objects.
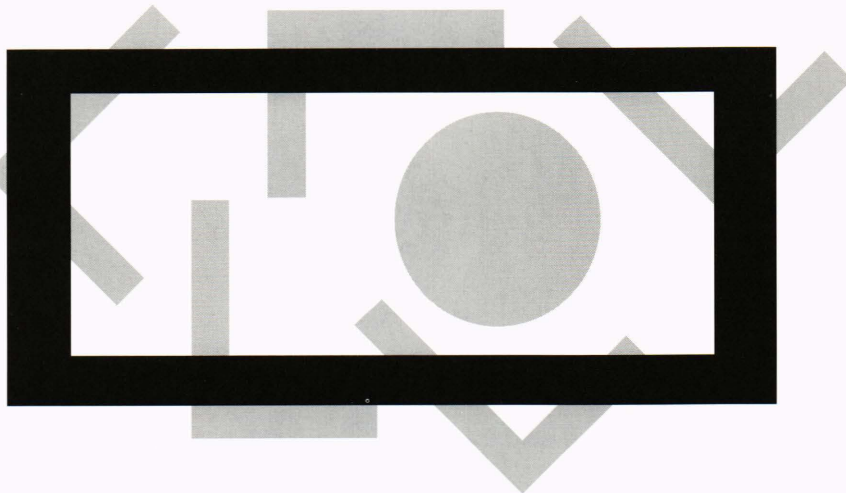
Figure 9. The white lines across the box illustrate the activity of hidden-edge detectors, which probably are activated in synchrony with cells responding to visible boundaries. The synchronous activity induces the brain to judge that the "L" fragments continue behind the box.
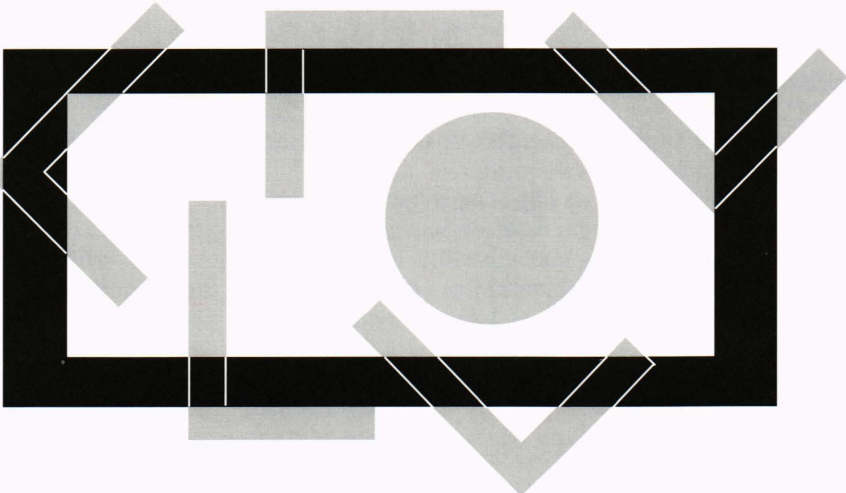
Figure 10. When objects share contrast boundaries, such as the rectangle and L-shaped objects of Figure 8 (white lines), shared-boundary edge detectors may be activated in synchrony with the larger aggregates of coaligned edge detectors (box) rather than with the smaller aggregates (L-shaped objects). The viewer then perceives that the larger object occludes the more remote smaller ones. Note that the drawing is all in one plane with no actual occultation of any object.
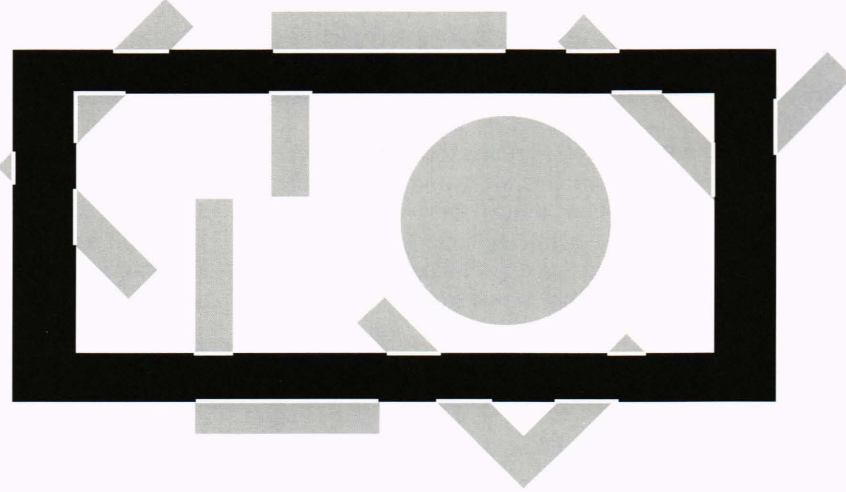
from different viewpoints and with various parts hidden. By parsing the unidentified object, we develop a list of subcomponents and rules about interconnections. Our memory search, then, is for a similar list with associated interconnection rules, as well as other characteristics, such as surface texture. The surface of the cat, for exam-
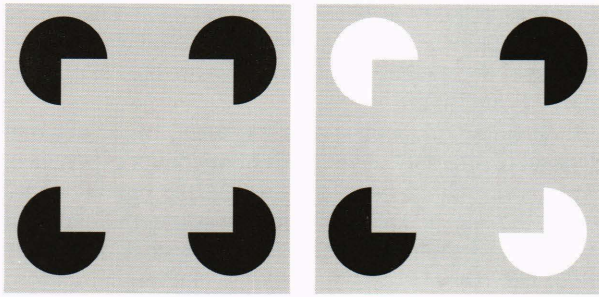
Figure 11. The illusion of a square covering the circles demonstrates the activation of coaligned detectors in regions without illumination edges, as also noted in Figure 9. Note that the illusion persists even in the reversed contrast version (right), indicating that boundary detectors are activated rather than edge-detector cells that are sensitive to direction of contrast. The perception of a rectangle in the line-only drawing in Figure 3B indicates that other mechanisms may play a role in creating the illusion.

ple, would be analyzed into active neural feature sets quite different from those of a chicken: "black fur, not white feathers," the selective surface-feature detectors would indicate to the higher processing centers. Many other characteristics would also be represented by activity in neural sets selective for those features—for example, movement. This identification procedure is probably hierarchical. That is, the same approach is used to identify some of the cat's more complex parts, such as the face.

The analyses used to activate these special sets of neurons related to the visual image of a cat are too complex to describe, but they are no different in principle from the type of neural analysis used for the horizontal line. Just as the horizontal-line detectors are active if and only if a horizontal line is present, the black-surface detectors are active only when a black surface is present in the physical world.

One visual cue that we may use when parsing a black cat or other uniformly colored unknown object into separate subcomponents is the short radius of curvature that occurs at the interface between natural divisions of most objects. When we look at a cat, the interfaces between its various natural subcomponents are often marked by abrupt changes in the orientation of the active sets of boundary detector neurons. Recall from our discussion of receptive fields that a particular orientation of a contrast edge will stimulate more central sets of nerve cells representing this edge orientation. By the same token, an abrupt change in the orientation of the object boundary will stimulate even more central sets of cortical neurons that represent potential natural subdivision points.

At least one current theory of visual object recognition is consistent with our postulated one. In "recognition-by-components" theory, object images are segmented into geometric components called "geons" at regions of deep concavity. Geons are simple shapes, such as blocks, cylinders, wedges, and cones, that can be modeled as generalized cones and can be derived from a small set of easily detectable properties of edges in 2D images, namely, curvature, collinearity, symmetry, parallelism, and cotermination. Different arrangements of a relatively small "vocabulary" of such components can produce a large number of multi-geon objects, where an object's representation in memory is a structural description of the relationships among its component geons.[19]

Our observations about the dancer with the lighted joints dancing in a dark room tend to support this view of object parsing: fixed-length subcomponents must exist

Figure 12. Parsing the cat. One way in which the brain may identify an object, such as a cat, is to parse (subdivide) the object into its major components and subcomponents, whose interfaces are marked by a short radius of curvature. Using the relative location, size, and interconnections of the parts as clues, the brain searches its memory for a list of similar components and associated interconnection rules for known objects ("cat construction" rules, for example), as well as other characteristics, such as surface texture and movement.

if their kinematic motion is being used to infer that the dancer exists.

PHILOSOPHICAL IMPLICATIONS AND SPECULATIONS

This concept of visual processing and perception leads to a rather strange conclusion, but one that I believe provides a satisfactory resolution to some ancient philosophical problems.

If we have normal sight, vision is our dominant modality for knowing the world. Many experiments have shown that we will ignore correct, but conflicting, information from the other senses.[20] Because of the dominance of vision and my concept of the visual process, I believe that the world "I" experience (or the world "I" dwell in) is a virtual world my brain has constructed, not the physical world. As Figures 1, 2, 3, and 11 illustrate, my experienced world and the physical world are not always the same. (By "I" or "we" in quotes, I mean our psychological selves, not our bodies that exist in the physical world.)

My body, including the brain, is governed by the same rules that regulate inanimate objects. That is, my body and brain are complex, deterministic biological mechanisms, explainable by advances in chemistry and physics. "I," however, dwell in a non-physical world constructed primarily with the information provided by my eyes and my memory. I believe "we" all are part of the reality constructed by our own brains, that is, part of a simulation of the physical world.

This simulation runs in (or is conducted by) the world's most advanced computer, the human brain. But what is represented in this computer is not determined by the physics of the brain, any more than the meaning of the symbols manipulated in a desktop computer is determined by the construction or processing details of that computer. Thus, the non-physical, simulated world in which "I" exist is not necessarily a deterministic world. When the brain is not running (conducting) this simulation, "I" and all of my world cease to exist. There is no "I," no personality, while my body is in deep dreamless sleep.

Accurately predicting the behavior of physical world objects such as falling rocks and other people has great survival value, both for our bodies and the "I" of the simulation. Thus, evolutionary selection has forced our brains to simulate the physical world with a high degree of fidelity under most conditions (no drugs). When the brain represents other people or the higher animals in the simulation, they have mental states and feelings as well as physical characteristics.

The brain's simulation of physical reality can most readily predict the behavior of other people by modeling their actions as if they were motivated by these mental states and feelings instead of by applying the laws of chemistry and physics. Consequently, my brain represents other people in the only reality "I" directly know (my simulation) as agents with free will, not as complex biological machines. Whether this representation of other humans is accurate is impossible for me to know with complete certainty. It is at least conceivable* that all the humans I observe are only complex machines behaving as if they have mental states and feel pain, but "I" directly know "I" feel pain.

I can infer (or deduce) the existence of an external physical reality only on the basis of my direct experience. There is no reality more certain than the one "I" experience, the one "I" know. I believe in an external physical reality, and I think that part of that physical world, my brain, has constructed the world "I" know and created "me." I may be mistaken in all or part of this belief, but not in the things "I" (and only "I") directly know. Likewise, there are things "you" directly know without possibility of error that "I" can only surmise, such as whether "you" are angry or in pain.

"I" have a unique and special relationship to my body. Not only does it exert great influence upon the simulation, such as making "me" thirsty at appropriate times, but "I" can also control it. "I" can, for example, cause my finger to move as "I" wish. Since this is a fact of the physical world (as is the fact that rocks fall), and evolution has forced my brain to be reasonably faithful in constructing the representation of the physical world when it creates "me" and my total reality in the simulation, "I" am represented in the simulation with special unique powers over my body and with uniquely privileged communication from its sensors. For example, "I" alone hear with my body's ears, a fact that is represented in the simulation.

Except for this special relationship to my body and the unique memory of my past, "I" am represented in my brain's simulation much like any other human. "I" am modeled with physical characteristics such as weight, height, and position in space and have feelings, dynamic mental states, purposeful behavior, moods, and free will. For "me," all of this is directly known and true, if anything is true, despite my belief (which is inferred and possibly false) that "I" and all of my perceived world exist only as a simulation conducted by the world's best computer.

This resolution of the ancient philosophical conflict between free will (or moral responsibility for one's actions) and the deterministic mechanical nature of one's body appears to be strange because it defines "I" as a nonmaterial informational process. A simple thought experiment will show, however, that the conventional concept of "I" as physical body is even more strange. If technology were to advance to the point where it became possible to replace every cell in my body with the exact functional but artificial equivalent of the original cell, then technology could copy "me" if "I" am physically defined. It would probably be necessary to destroy the original "me" as my body is disassembled for examination and replication, one cell at a time. If this advanced technology could produce an exact functional copy of a physical "me," it could produce a dozen copies, each of

---

*It is more than conceivable. Behaviorism, which rejected mental states as a matter of principle, was the orthodox view among American scientists for more than three decades in the early part of this century.

which could equally well claim to be the true "me." Therefore, the conventional view, which identifies "me" with my physical body, becomes increasingly untenable, although its flaws are not widely recognized because of the current limitations of technology.

As more and more artificial organs are developed and implanted, some people have begun to question what constitutes "a person." Most people conclude that all parts of an individual except for the brain could conceptually be replaced without destroying the person, but what is materially unique about the brain? What, from a material point of view, prevents "a person" from continuing to exist when all of his or her brain cells have been replaced by functionally equivalent artificial brain cells?
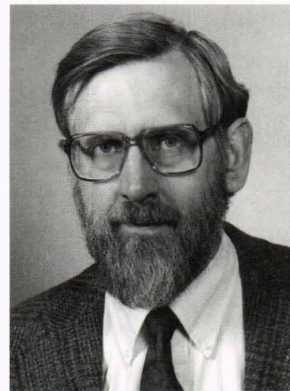
I contend that advanced technology could not duplicate the informational process "me" of my theory because the process of disassembling my brain to characterize the function of each of my brain cells and their mutual interconnections would destroy the informational process (i.e., the "me" in my theory). I suspect that, in analogy with the quantum mechanical uncertainty principle, there is a "biological uncertainty principle," which essentially states that the more detailed and accurate the measurements designed to characterize the informational processes of a complex biological system, the more the measurements disrupt those informational processes. An EEG, for example, interferes little with brain processes, but provides little detail about the activity of brain cells. A set of invasive electrodes is more disruptive but provides more detailed information about the disrupted activity of the few brain cells near the electrode tips. If this biological uncertainty principle is true—and it is difficult to imagine that it is not when applied to a system as complex as the human brain—then it is not possible in principle to duplicate "me" as an informational process existing in a human brain. My body may be deterministic, but "I" cannot be copied; "I" am unique; "I" have free will.

REFERENCES

[1] Gregory, R. L., *Eye and Brain*, 4th ed., Princeton University Press, Princeton, N.J., Figure 5.9, p. 70 (1979).

[2] DeValois, R. L., and DeValois, K. K., *Spatial Vision*, Oxford University Press, New York, Figure 2.10, p. 50 (1988).

[3] Neisser, U., *Cognitive Psychology*, Appleton-Century-Crofts, New York, pp. 212-218 (1967).

[4] Bach-y-Rita, P., *Brain Mechanisms in Sensory Substitution*, Academic Press, New York (1972).

[5] McBurney, D. H., and Collings, V. B., *An Introduction to Sensation/Perception*, Prentice-Hall, Englewood Cliffs, N.J., pp. 204-206 (1977).

[6] In 1637, Descartes described this shape constancy phenomenon in his *Dioptrics*: "Judgments of shape clearly come from our knowledge, or opinion, as to the position of the various parts of the objects and not in accordance with the pictures in the eye; for these pictures normally contain ovals and diamonds when they cause us to see circles and squares." [Quoted by R. L. Gregory in *Eye and Brain*, 4th ed., Princeton University Press, Princeton, N.J., p. 161 (1979).]

[7] von der Heydt, R., and Peterhans, E., "Mechanisms of Contour Perception in Monkey Visual Cortex. I. Lines of Pattern Discontinuity," *J. Neurosci.* **9**(5), 1731-1748, Figure 1 (1989).

[8] Frisby, J. P., *Seeing: Illusion, Brain and Mind*, Oxford University Press, New York (1979).

[9] DeValois, R. L., and DeValois K. K., *Spatial Vision*, Oxford University Press, New York, Figure 3.16, pp. 84-85 (1988).

[10] Gregory, R. L., *Eye and Brain*, 4th ed., Princeton University Press, Princeton, N.J., p. 55 (1979).

[11] Popper, K. R., and Eccles, J. C., *The Self and Its Brain*, Springer International, New York, p. 268 (1981).

[12] von der Malsburg, C., "Self-Organization of Orientation Sensitive Cells in the Striate Cortex," *Kybernetik*a **14**, 85-100 (1973).

[13] Gray, C. M., König, P., Engel, A. K., and Singer, W., "Oscillatory Responses in Cat Visual Cortex Exhibit Inter-Columnar Synchronization Which Reflects Global Stimulus Properties," *Nature* **338**, 334-337 (1989).

[14] Kaufman, L., *Perception: The World Transformed*, Oxford University Press, New York, Figure 5-26, pp. 143-146 (1979).

[15] Kanizsa, G., "Subjective Contours," *Sci. Am.* **234**(4), 48-52 (1976).

[16] Grossberg, S. "Cortical Dynamics of Three-Dimensional Form, Color, and Brightness Perception; I: Monocular Theory," *Percept. Psychophys.* **41**(2), 87-116 (1987).

[17] Hubel, D. H., and Wiesel, T. N., "Receptive Fields, Binocular Interactions and Functional Architecture in the Cat's Visual Cortex," *J. Physiol.-London* **160**, 106-154 (1962).

[18] Johansson, G., "Visual Motion Perception," *Sci. Am.* **232**(6), 76-88 (1975).

[19] Biederman, I., "Recognition-by-Components: A Theory of Human Image Understanding," *Psychol. Rev.* **94**, 115-147 (1987).

[20] Rock, I., and Harris, C. S., "Vision and Touch," *Sci. Am.* **216**(5), 96-104 (1967).

THE AUTHOR

WILLIAM R. POWELL received a B.A. degree in engineering physics from Cornell University in 1959 and a Ph.D. degree in physics from The Johns Hopkins University in 1966. He joined APL's Research Center in 1966, transferring to the Space Department in 1974, where he designed several biomedical monitors and sensors. He was a member of the Naval Warfare Analysis Department from 1983 until his retirement in December 1993, working on battlegroup coordination and over-the-horizon targeting. He also spent one year in The Johns Hopkins University Cognitive Science Center as the Hafstad Fellow and actively participated in visual attention and cognitive science seminars in the Psychology Department until his retirement. The holder of several patents, he currently resides in Brazil and is researching the perceptual capacities of the blind in affiliation with the University of São Paulo and two smaller universities.