

The Problem of Consciousness

It can now be approached by scientific investigation of the visual system. The solution will require a close collaboration among psychologists, neuroscientists and theorists

by Francis Crick and Christof Koch



The overwhelming question in neurobiology today is the relation between the mind and the brain. Everyone agrees that what we know as mind is closely related to certain aspects of the behavior of the brain, not to the heart, as Aristotle thought. Its most mysterious aspect is consciousness or awareness, which can take many forms, from the experience of pain to self-consciousness. In the past the mind (or soul) was often regarded, as it was by Descartes, as something immaterial, separate from the brain but interacting with it in some way. A few neuroscientists, such as Sir John Eccles, still assert that the soul is distinct from the body. But most neuroscientists now believe that all aspects of mind, including its most puzzling attribute—consciousness or awareness—are likely to be explainable in a more materialistic way as the behavior of large sets of interacting neurons. As William James, the father of American psychology, said a century ago, consciousness is not a thing but a process.

Exactly what the process is, however, has yet to be discovered. For many years after James penned *The Principles of Psychology*, consciousness was a taboo concept in American psychology because of the dominance of the behaviorist movement. With the advent of cognitive science in the mid-1950s, it became possible once more for psychologists to consider mental processes as opposed to merely observing behavior. In spite of these changes, until recently most cognitive scientists ignored consciousness, as did almost all neuroscientists. The problem was felt to be either purely “philosophical” or too elusive to study experimentally. It would not have been easy for a neuroscientist to get a grant just to study consciousness.

In our opinion, such timidity is ridiculous, so a few years ago we began to think about how best to attack the problem scientifically. How to explain mental events as being caused by the firing of large sets of neurons? Although there are those who believe such an approach is hopeless, we feel it is not productive to worry too much over aspects of the problem that cannot be solved scientifically or, more precisely, cannot be solved solely by using existing scientific ideas. Radically new concepts may indeed be needed—recall the modifications of scientific thinking forced on us by quantum mechanics. The only sensible approach is to press the experimental attack until we are confronted with dilemmas that call for new ways of thinking.

There are many possible approaches to the problem of consciousness. Some psychologists feel that any satisfactory theory should try to explain as many aspects of consciousness as possible, including emotion, imagination, dreams, mystical experiences and so on. Although such an all-embracing theory will be necessary in the long run, we thought it wiser to begin with the particular aspect of consciousness that is likely to yield most easily. What this aspect may be is a matter of personal judgment. We selected the mammalian visual system because humans are very visual animals and because so much experimental and theoretical work has already been done on it.

It is not easy to grasp exactly what we need to explain, and it will take many careful experiments before visual consciousness can be described scientifically. We did

VISUAL AWARENESS primarily involves seeing what is directly in front of you, but it can be influenced by a three-dimensional representation of the object in view retained by the brain. If you see the back of a person’s head, the brain infers that there is a face on the front of it. We know this is true because we would be very startled if a mirror revealed that the front was exactly like the back, as in this painting, *Reproduction Prohibited* (1937), by René Magritte.

Reprinted from the September 1992 issue

The Authors

FRANCIS CRICK and CHRISTOF KOCH share an interest in the experimental study of consciousness. Crick is the co-discoverer, with James Watson, of the double helical structure of DNA. While at the Medical Research Council Laboratory of Molecular Biology in Cambridge, he worked on the genetic code and on developmental biology. Since 1976, he has been at the Salk Institute for Biological Studies in San Diego. His main interest lies in understanding the visual system of mammals. Koch was awarded his Ph.D. in biophysics by the University of Tübingen. After spending four years at the Massachusetts Institute of Technology, he joined the California Institute of Technology, where he is now professor of computation and neural systems. He is studying how single brain cells process information and the neural basis of motion perception, visual attention and awareness. He also designs analog VLSI vision chips for intelligent systems.

Lewis Carroll’s vanishing cat can be used to study awareness.



not attempt to define consciousness itself because of the dangers of premature definition. (If this seems like a copout, try defining the word “gene”—you will not find it easy.) Yet the experimental evidence that already exists provides enough of a glimpse of the nature of visual consciousness to guide research. In this article, we will attempt to show how this evidence opens the way to attack this profound and intriguing problem.

Visual theorists agree that the problem of visual consciousness is ill posed. The mathematical term “ill posed” means that additional constraints are needed to solve the problem. Although the main function of the visual system is to perceive objects and events in the world around us, the information available to our eyes is not sufficient by itself to provide the brain with its unique interpretation of the visual world. The brain must use past experience (either its own or that of our distant ancestors, which is embedded in our genes) to help interpret the information coming into our eyes. An example would be the derivation of the three-dimensional representation of the world from the two-dimensional signals falling onto the retinas of our two eyes or even onto one of them.

Visual theorists also would

agree that seeing is a constructive process, one in which the brain has to carry out complex activities (sometimes called computations) in order to decide which interpretation to adopt of the ambiguous visual input. “Computation” implies that the brain acts to form a symbolic representation of the visual world, with a mapping (in the mathematical sense) of certain aspects of that world onto elements in the brain.

Ray Jackendoff of Brandeis University postulates, as do most cognitive scientists, that the computations carried out by the brain are largely unconscious and that what we become aware of is the result of these computations. But while the customary view is that this awareness occurs at the highest levels of the computational system, Jackendoff has proposed an intermediate-level theory of consciousness.


What we see, Jackendoff suggests, relates to a representation of surfaces that are directly visible to us, together with their outline, orientation, color, texture

and movement. (This idea has similarities to what the late David C. Marr of the Massachusetts Institute of Technology called a “ $2\frac{1}{2}$ -dimensional sketch.”) It is more than a two-dimensional sketch because it conveys the orientation of the visible surfaces. It is less than three-dimensional because depth information is not explicitly represented.) In the next stage this sketch is processed by the brain to produce a three-dimensional representation. Jackendoff argues that we are not visually aware of this three-dimensional representation.

An example may make this process clearer. If you look at a person whose back is turned to you, you can see the back of the head but not the face. Nevertheless, your brain infers that the person has a face. We can deduce as much because if that person turned around and had no face, you would be very surprised.

The viewer-centered representation that corresponds to the visible back of the head is what you are vividly aware of. What your brain infers about the front would come from some kind of three-dimensional representation. This does not mean that information flows only from the surface representation to the three-dimensional one; it almost certainly flows in both directions.

AMBIGUOUS IMAGES were frequently used by Salvador Dali in his paintings. In *Slave Market with the Disappearing Bust of Voltaire* (1940), the bust of the French philosopher Voltaire is apparent from a distance but transforms into figures of three people when viewed at close range. Studies with ambiguous figures in the behaving monkey have found that many neurons in higher cortical areas respond only to the currently “perceived” figure; the neuronal response to the “unseen” image is suppressed.



©1997 DEMART PRO ARTE (R), GENEVA/ARTISTS RIGHTS SOCIETY (ARS), NEW YORK, © SALVADOR DALI MUSEUM, INC., ST. PETERSBURG, FLORIDA

When you imagine the front of the face, what you are aware of is a surface representation generated by information from the three-dimensional model.

It is important to distinguish between an explicit and an implicit representation. An explicit representation is something that is symbolized without further processing. An implicit representation contains the same information but requires further processing to make it explicit. The pattern of colored dots on a television screen, for example, contains an implicit representation of objects (say, a person's face), but only the dots and their locations are explicit. When you see a face on the screen, there must be neurons in your brain whose firing, in some sense, symbolizes that face.

We call this pattern of firing neurons an active representation. A latent representation of a face must also be stored in the brain, probably as a special pattern of synaptic connections between neurons. For example, you probably have a representation of the Statue of Liberty in your brain, a representation that usually is inactive. If you do think about the Statue, the representation becomes active, with the relevant neurons firing away.


An object, incidentally, may be represented in more than one way—as a visual image, as a set of words and their related sounds, or even as a touch or a smell. These different representations are likely to interact with one another. The representation is likely to be distributed over many neurons, both locally and more globally. Such a representation may not be as simple and straightforward as uncritical introspection might indicate. There is suggestive evidence, partly from studying how neurons fire in various parts of a monkey's brain and partly from examining the effects of certain types of brain damage in humans, that different aspects of a face—and of the implications of a face—may be represented in different parts of the brain.

First, there is the representation of a face as a face: two eyes, a nose, a mouth and so on. The neurons involved are usually not too fussy about the exact size or position of this face in the visual field, nor are they very sensitive to small changes in its orientation. In monkeys, there are neurons that respond best when the face is turning in a particular direction, while others seem to be more concerned with the direction in which the eyes are gazing.

Then there are representations of the

parts of a face, as separate from those for the face as a whole. Further, the implications of seeing a face, such as that person's sex, the facial expression, the familiarity or unfamiliarity of the face, and in particular whose face it is, may each be correlated with neurons firing in other places.

What we are aware of at any moment, in one sense or another, is not a simple matter. We have suggested that there may be a very transient form of fleeting awareness that represents only rather simple features and does not require an attentional mechanism. From this brief awareness the brain constructs a viewer-centered representation—what we see



WILLIAM JAMES, the father of American psychology, observed that consciousness is not a thing but a process.

ividly and clearly—that does require attention. This in turn probably leads to three-dimensional object representations and thence to more cognitive ones.

Representations corresponding to vivid consciousness are likely to have special properties. William James thought that consciousness involved both attention and short-term memory. Most psychologists today would agree with this view. Jackendoff writes that consciousness is “enriched” by attention, implying that whereas attention may not be essential for certain limited types of consciousness, it is necessary for full consciousness. Yet it is not clear exactly which forms of memory are involved. Is long-term memory needed? Some forms of acquired knowledge are so

embedded in the machinery of neural processing that they are almost certainly used in becoming aware of something. On the other hand, there is evidence from studies of brain-damaged patients that the ability to lay down new long-term episodic memories is not essential for consciousness to be experienced.

It is difficult to imagine that anyone could be conscious if he or she had no memory whatsoever of what had just happened, even an extremely short one. Visual psychologists talk of iconic memory, which lasts for a fraction of a second, and working memory (such as that used to remember a new telephone number) that lasts for only a few seconds unless it is rehearsed. It is not clear whether both of these are essential for consciousness. In any case, the division of short-term memory into these two categories may be too crude.

If these complex processes of visual awareness are localized in parts of the brain, which processes are likely to be where? Many regions of the brain may be involved, but it is almost certain that the cerebral neocortex plays a dominant role. Visual information from the retina reaches the neocortex mainly by way of a part of the thalamus (the lateral geniculate nucleus); another significant visual pathway from the retina is to the superior colliculus, at the top of the brain stem.

The cortex in humans consists of two intricately folded sheets of nerve tissue, one on each side of the head. These sheets are connected by a large tract of about half a billion axons called the corpus callosum. It is well known that if the corpus callosum is cut, as is done for certain cases of intractable epilepsy, one side of the brain is not aware of what the other side is seeing. In particular, the left side of the brain (in a right-handed person) appears not to be aware of visual information received exclusively by the right side. This shows that none of the information required for visual awareness can reach the other side of the brain by traveling down to the brain stem and, from there, back up. In a normal person, such information can get to the other side only by using the axons in the corpus callosum.

A different part of the brain—the hippocampal system—is involved in one-shot, or episodic, memories that, over weeks and months, it passes on to the neocortex. This system is so placed that it receives inputs from, and projects to, many parts of the brain. Thus, one might suspect that the hippocampal system is

the essential seat of consciousness. This is not the case: evidence from studies of patients with damaged brains shows that this system is not essential for visual awareness, although naturally a patient lacking one is severely handicapped in everyday life because he cannot remember anything that took place more than a minute or so in the past.

In broad terms, the neocortex of alert animals probably acts in two ways. By building on crude and somewhat redundant wiring, produced by our genes and by embryonic processes, the neocortex draws on visual and other experience to slowly “rewire” itself to create categories (or “features”) it can respond to. A new category is not fully created in the neocortex after exposure to only one example of it, although some small modifications of the neural connections may be made.

The second function of the neocortex (at least of the visual part of it) is to respond extremely rapidly to incoming signals. To do so, it uses the categories it has learned and tries to find the combinations of active neurons that, on the basis of its past experience, are most likely to represent the relevant objects and events in the visual world at that moment. The formation of such coalitions of active neurons may also be influenced by biases coming from other parts of the brain: for example, signals telling it what best to attend to or high-level expectations about the nature of the stimulus.

Consciousness, as James noted, is always changing. These rapidly formed coalitions occur at different levels and interact to form even broader coalitions. They are transient, lasting usually for only a fraction of a second. Because coalitions in the visual system are the basis of what we see, evolution has seen to it that they form as fast as possible; otherwise, no animal could survive. The brain is handicapped in forming neuronal coalitions rapidly because, by computer standards, neurons act very slowly. The brain compensates for this relative slowness partly by using very many neurons, simultaneously and in parallel, and partly by arranging the system in a roughly hierarchical manner.

If visual awareness at any moment corresponds to sets of neurons firing, then the obvious question is: Where are these neurons located in the brain, and in what way are they firing? Visual awareness is highly unlikely to occupy all the neurons in the neocortex that are

firing above their background rate at a particular moment. We would expect that, theoretically, at least some of these neurons would be involved in doing computations—trying to arrive at the best coalitions—whereas others would express the results of these computations, in other words, what we see.

Fortunately, some experimental evidence can be found to back up this theoretical conclusion. A phenomenon called binocular rivalry may help identify the neurons whose firing symbolizes awareness. This phenomenon can be seen in dramatic form in an exhibit prepared by Sally Duensing and Bob Miller at the Exploratorium in San Francisco.

Conflicting Inputs

Binocular rivalry occurs when each eye has a different visual input relating to the same part of the visual field. The early visual system on the left side of the brain receives an input from both eyes but sees only the part of the visual field to the right of the fixation point. The converse is true for the right side. If these two conflicting inputs are rivalrous, one sees not the two inputs superimposed but first one input, then the other, and so on in alternation.

In the exhibit, called “The Cheshire Cat,” viewers put their heads in a fixed place and are told to keep the gaze fixed. By means of a suitably placed mirror, one of the eyes can look at another person’s face, directly in front, while the other eye sees a blank white screen to the side. If the viewer waves a hand in front of this plain screen at the same location in his or her visual field occupied by the face, the face is wiped out. The movement of the hand, being visually very salient, has captured the brain’s attention. Without attention the face cannot be seen. If the viewer moves the eyes, the face reappears.

In some cases, only part of the face disappears. Sometimes, for example, one eye, or both eyes, will remain. If the viewer looks at the smile on the person’s face, the face may disappear, leaving only the smile. For this reason, the effect has been called the Cheshire Cat effect, after the cat in Lewis Carroll’s *Alice’s Adventures in Wonderland*.

Although it is very difficult to record activity in individual neurons in a human brain, such studies can be done in monkeys. A simple example of binocular rivalry has been studied in a monkey by Nikos K. Logothetis and Jeffrey D. Schall, both then at M.I.T. They

The Cheshire Cat Experiment

This simple experiment with a mirror illustrates one aspect of visual awareness. It relies on a phenomenon called binocular rivalry, which occurs when each eye has a different input from the same part of the visual field. Motion in the field of one eye can cause either the entire image or parts of the image to be erased. The movement captures the brain’s attention.



trained a macaque to keep its eyes still and to signal whether it is seeing upward or downward movement of a horizontal grating. To produce rivalry, upward movement is projected into one of the monkey's eyes and downward movement into the other, so that the two images overlap in the visual field. The monkey signals that it sees up and down movements alternatively, just as humans would. Even though the motion stimulus coming into the monkey's eyes is always the same, the monkey's percept changes every second or so.

Cortical area MT (which some researchers prefer to label V5) is an area mainly concerned with movement. What do the neurons in MT do when the monkey's percept is sometimes up and sometimes down? (The researchers studied only the monkey's first response.) The simplified answer—the actual data are rather more messy—is that whereas the firing of some of the neurons correlates with the changes in the percept, for others the average firing rate is relatively unchanged and independent of which direction of movement the monkey is seeing at that moment. Thus, it is unlikely that the firing of all the neurons in the visual neocortex at one particular moment corresponds to the monkey's

visual awareness. Exactly which neurons do correspond to awareness remains to be discovered.

We have postulated that when we clearly see something, there must be neurons actively firing that stand for what we see. This might be called the activity principle. Here, too, there is some experimental evidence. One example is the firing of neurons in a specific cortical visual area in response to illusory contours. Another and perhaps more striking case is the filling in of the blind spot. The blind spot in each eye is caused by the lack of photoreceptors in the area of the retina where the optic nerve leaves the retina and projects to the brain. Its location is about 15 degrees from the fovea (the visual center of the eye). Yet if you close one eye, you do not see a hole in your visual field.

Philosopher Daniel C. Dennett of Tufts University is unusual among philosophers in that he is interested both in psychology and in the brain. This interest is much to be welcomed. In a recent book, *Consciousness Explained*, he has argued that it is wrong to talk about filling in. He concludes, correctly, that “an absence of information is not the same as information about an absence.” From this general principle he argues that the

brain does not fill in the blind spot but rather ignores it.

Dennett's argument by itself, however, does not establish that filling in does not occur; it only suggests that it might not. Dennett also states that “your brain has no machinery for [filling in] at this location.” This statement is incorrect. The primary visual cortex lacks a direct input from one eye, but normal “machinery” is there to deal with the input from the other eye. Ricardo Gattass and his colleagues at the Federal University of Rio de Janeiro have shown that in the macaque some of the neurons in the blind-spot area of the primary visual cortex do respond to input from both eyes, probably assisted by inputs from other parts of the cortex. Moreover, in the case of simple filling in, some of the neurons in that region respond as if they were actively filling in.

Thus, Dennett's claim about blind spots is incorrect. In addition, psychological experiments by Vilayanur S. Ramachandran [see “Blind Spots,” *SCIENTIFIC AMERICAN*, May 1992] have shown that what is filled in can be quite complex depending on the overall context of the visual scene. How, he argues, can your brain be ignoring something that is in fact commanding attention?

Filling in, therefore, is not to be dismissed as nonexistent or unusual. It probably represents a basic interpolation process that can occur at many levels in the neocortex. It is, incidentally, a good example of what is meant by a constructive process.


How can we discover the neurons whose firing symbolizes a particular percept? William T. Newsome and his colleagues at Stanford University have done a series of brilliant experiments on neurons in cortical area MT of the macaque's brain. By studying a neuron in area MT, we may discover that it responds best to very specific visual features having to do with motion. A neuron, for instance, might fire strongly in response to the movement of a bar in a particular place in the visual field, but only when the bar is oriented at a certain angle, moving in one of the two directions perpendicular to its length within a certain range of speed.

It is technically difficult to excite just a single neuron, but it is known that neurons that respond to roughly the same position, orientation and direction of movement of a bar tend to be located near one another in the cortical sheet. The experimenters taught the monkey a

To observe the effect, a viewer divides the field of vision with a mirror placed between the eyes (a). One eye sees the cat; the other eye a reflection in the mirror of a white wall or background. The viewer then waves the hand that corresponds to the eye looking at the mirror so that the hand passes through the area in which the image of the cat appears in the other eye (b). The result is that the cat may disappear. Or if the viewer was attentive to a specific feature before the hand was waved, those parts—the eyes or even a mocking smile—may remain (c). —F.C. and C.K.



PHOTOGRAPHS BY JASON GOLTZ



JOHNNY JOHNSON

OPTICAL ILLUSION devised by Vilayanur S. Ramachandran illustrates the brain's ability to fill in, or construct, visual information that is missing because it falls on the blind spot of the eye. When you look at the patterns of broken green bars, the visual system produces two illusory contours defining a vertical strip. Now shut your right eye and focus on the white square in the green series of bars. Move the page toward your eye until the blue dot disappears (roughly six inches in front of your nose). Most observers report seeing the vertical strip completed across the blind spot, not the broken line. Try the same experiment with the series of just three red bars. The illusory vertical contours are less well defined, and the visual system tends to fill in the horizontal bar across the blind spot. Thus, the brain fills in differently depending on the overall context of the image.

simple task in movement discrimination using a mixture of dots, some moving randomly, the rest all in one direction. They showed that electrical stimulation of a small region in the right place in cortical area MT would bias the monkey's motion discrimination, almost always in the expected direction.

Thus, the stimulation of these neurons can influence the monkey's behavior and probably its visual percept. Such experiments do not, however, show decisively that the firing of such neurons is the exact neural correlate of the percept. The correlate could be only a subset of the neurons being activated. Or perhaps the real correlate is the firing of neurons in another part of the visual hierarchy that are strongly influenced by the neurons activated in area MT.

These same reservations apply also to cases of binocular rivalry. Clearly, the problem of finding the neurons whose firing symbolizes a particular percept is not going to be easy. It will take many careful experiments to track them down even for one kind of percept.

It seems obvious that the purpose of vivid visual awareness is to feed into the cortical areas concerned with the implications of what we see; from there the information shuttles on the one hand to the hippocampal system, to be encoded (temporarily) into long-term episodic memory, and on the other to the planning levels of the motor system. But is it possible to go from a visual input to a behavioral output without any relevant visual awareness?

That such a process can happen is demonstrated by the remarkable class of patients with "blindsight." These patients, all of whom have suffered damage to their visual cortex, can point with fair accuracy at visual targets or track them with their eyes while vigorously denying seeing anything. In fact, these patients are as surprised as their doctors by their abilities. The amount of information that "gets through," however, is limited: blindsight patients have some ability to respond to wavelength, orientation and motion, yet they cannot distinguish a triangle from a square.

It is naturally of great interest to know which neural pathways are being used in these patients. Investigators originally suspected that the pathway ran through the superior colliculus. Recent experiments suggest that a direct albeit weak connection may be involved between the lateral geniculate nucleus and other visual areas in the cortex. It is unclear

whether an intact primary visual cortex region is essential for immediate visual awareness. Conceivably the visual signal in blindsight is so weak that the neural activity cannot produce awareness, although it remains strong enough to get through to the motor system.

Normal-seeing people regularly respond to visual signals without being fully aware of them. In automatic actions, such as swimming or driving a car, complex but stereotypical actions occur with little, if any, associated visual awareness. In other cases, the information conveyed is either very limited or very attenuated. Thus, while we can function without visual awareness, our behavior without it is rather restricted.

Clearly, it takes a certain amount of time to experience a conscious percept. It is difficult to determine just how much time is needed for an episode of visual awareness, but one aspect of the problem that can be demonstrated experimentally is that signals received close together in time are treated by the brain as simultaneous.

A disk of red light is flashed for, say, 20 milliseconds, followed immediately by a 20-millisecond flash of green light in the same place. The subject reports that he did not see a red light followed by a green light. Instead he saw a yellow light, just as he would have if the red and the green light had been flashed simultaneously. Yet the subject could not have experienced yellow until after the information from the green flash had been processed and integrated with the preceding red one.

Experiments of this type led psychologist Robert Efron, now at the University of California at Davis, to conclude that the processing period for perception is about 60 to 70 milliseconds. Similar periods are found in experiments with tones in the auditory system. It is always possible, however, that the processing times may be different in higher parts of the visual hierarchy and in other parts of the brain. Processing is also more rapid in trained, compared with naive, observers.

Because it appears to be involved in some forms of visual awareness, it would help if we could discover the neural basis of attention. Eye movement is a form of attention, since the area of the visual field in which we see with high resolution is remarkably small, roughly the area of the thumbnail at arm's length. Thus, we move our eyes to gaze directly at an object in order to see it more clear-

ly. Our eyes usually move three or four times a second. Psychologists have shown, however, that there appears to be a faster form of attention that moves around, in some sense, when our eyes are stationary.

The exact psychological nature of this faster attentional mechanism is at present controversial. Several neuroscientists, however, including Robert Desimone and his colleagues at the National Institute of Mental Health, have shown that the rate of firing of certain neurons in the macaque's visual system depends on what the monkey is attending to in the visual field. Thus, attention is not solely a psychological concept; it also has neural correlates that can be observed. A number of researchers have found that the pulvinar, a region of the thalamus, appears to be involved in visual attention. We would like to believe that the thalamus deserves to be called "the organ of attention," but this status has yet to be established.

Attention and Awareness

The major problem is to find what activity in the brain corresponds directly to visual awareness. It has been speculated that each cortical area produces awareness of only those visual features that are "columnar," or arranged in the stack or column of neurons perpendicular to the cortical surface. Thus, the primary visual cortex could code for orientation and area MT for motion. So far experimentalists have not found one particular region in the brain where all the information needed for visual awareness appears to come together. Dennett has dubbed such a hypothetical place "The Cartesian Theater." He argues on theoretical grounds that it does not exist.

Awareness seems to be distributed not just on a local scale, but more widely over the neocortex. Vivid visual awareness is unlikely to be distributed over every cortical area because some areas show no response to visual signals. Awareness might, for example, be associated with only those areas that connect back directly to the primary visual cortex or alternatively with those areas that project into one another's layer 4.

Rhythmic and synchronized firing may be the neural correlate of awareness and might bind together activity [concerning the same object] in different cortical areas.

(The latter areas are always at the same level in the visual hierarchy.)

The key issue, then, is how the brain forms its global representations from visual signals. If attention is indeed crucial for visual awareness, the brain could form representations by attending to just one object at a time, rapidly moving from one object to the next. For example, the neurons representing all the different aspects of


the attended object could all fire together very rapidly for a short period, possibly in rapid bursts.

This fast, simultaneous firing might not only excite those neurons that symbolized the implications of that object but also temporarily strengthen the relevant synapses so that this particular pattern of firing could be quickly recalled—a form of short-term memory. If only one representation needs to be held in short-term memory, as in remembering a single task, the neurons involved may continue to fire for a period.

A problem arises if it is necessary to be aware of more than one object at exactly the same time. If all the attributes of two or more objects were represented by neurons firing rapidly, their attributes might be confused. The color of one might become attached to the shape of another. This happens sometimes in very brief presentations.

Some time ago Christoph von der Malsburg, now at the Ruhr-Universität Bochum, suggested that this difficulty would be circumvented if the neurons associated with any one object all fired in synchrony (that is, if their times of firing were correlated) but out of synchrony with those representing other objects. Recently two groups in Germany reported that there does appear to be correlated firing between neurons in the visual cortex of the cat, often in a rhythmic manner, with a frequency in the 35- to 75-hertz range, sometimes called 40-hertz, or *g*, oscillation.

Von der Malsburg's proposal prompted us to suggest that this rhythmic and synchronized firing might be the neural correlate of awareness and that it might serve to bind together activity concerning the same object in different cortical areas. The matter is still undecided, but at present the fragmentary experimen-



JASON GOLTZ

BRIEF FLASHES of colored light enable researchers to infer the minimum time required for visual awareness. A disk of red light is projected for 20 milliseconds (a), followed immediately by a 20-millisecond flash of green light (b). But the observer reports seeing a single flash of yellow (c), the color that would be apparent if red and green were projected simultaneously. The subject does not become aware of red followed by green until the length of the flashes is extended to 60 to 70 milliseconds.

tal evidence does rather little to support such an idea. Another possibility is that the 40-hertz oscillations may help distinguish figure from ground or assist the mechanism of attention.

Correlates of Consciousness

Are there some particular types of neurons, distributed over the visual neocortex, whose firing directly symbolizes the content of visual awareness? One very simplistic hypothesis is that the activities in the upper layers of the cortex are largely unconscious ones, whereas the activities in the lower layers (layers 5 and 6) mostly correlate with consciousness. We have wondered whether the pyramidal neurons in layer 5 of the neocortex, especially the larger ones, might play this latter role.

These are the only cortical neurons that project right out of the cortical system (that is, not to the neocortex, the thalamus or the claustrum). If visual awareness represents the results of neural computations in the cortex, one might expect that what the cortex sends elsewhere would symbolize those re-

sults. Moreover, the neurons in layer 5 show a rather unusual propensity to fire in bursts. The idea that layer 5 neurons may directly symbolize visual awareness is attractive, but it still is too early to tell whether there is anything in it.

Visual awareness is clearly a difficult problem. More work is needed on the psychological and neural basis of both attention and very short term memory. Studying the neurons when a percept changes, even though the visual input is

constant, should be a powerful experimental paradigm. We need to construct neurobiological theories of visual awareness and test them using a combination of molecular, neurobiological and clinical imaging studies.

We believe that once we have mastered the secret of this simple form of awareness, we may be close to understanding a central mystery of human life: how the physical events occurring in our brains while we think and act in the world relate to our subjective sensations—that is, how the brain relates to the mind.

Postscript: There have been several relevant developments since this article was first published. It now seems likely that there are rapid “on-line” systems for stereotyped motor responses such as hand or eye movement. These systems are unconscious and lack memory. Conscious seeing, on the other hand, seems to be slower and more subject to visual illusions. The brain needs to form a conscious representation of the visual scene that it then can use for many different actions or thoughts. Exactly how all these pathways work and how they interact is far from clear.

There have been more experiments on the behavior of neurons that respond to bistable visual percepts, such as binocular rivalry, but it is probably too early to draw firm conclusions from them about the exact neural correlates of visual consciousness. We have suggested on theoretical grounds based on the neuroanatomy of the macaque monkey that primates are not directly aware of what is happening in the primary visual cortex, even though most of the visual information flows through it. This hypothesis is supported by some experimental evidence, but it is still controversial. SA

Further Reading

- PERCEPTION. Irvin Rock. Scientific American Library, 1984.
 - CONSCIOUSNESS AND THE COMPUTATIONAL MIND. Ray Jackendoff. MIT Press/Bradford Books, 1987.
 - COLD SPRING HARBOR SYMPOSIA ON QUANTITATIVE BIOLOGY, VOL. LV: THE BRAIN. Cold Spring Harbor Laboratory Press, 1990.
 - TOWARDS A NEUROBIOLOGICAL THEORY OF CONSCIOUSNESS. Francis Crick and Christof Koch in *Seminars in the Neurosciences*, Vol. 2, pages 263–275; 1990.
 - THE COMPUTATIONAL BRAIN. Patricia S. Churchland and Terrence J. Sejnowski. MIT Press/Bradford Books, 1992.
 - THE VISUAL BRAIN IN ACTION. A. David Milner and Melvyn A. Goodale. Oxford University Press, 1995.
 - ARE WE AWARE OF NEURAL ACTIVITY IN PRIMARY VISUAL CORTEX? Francis Crick and Christof Koch in *Nature*, Vol. 375, pages 121–123, May 11, 1995.
-