

Reprint Series
4 November 1988, Volume 242, pp. 741-745

SCIENCE

Perspectives on Cognitive Neuroscience

PATRICIA S. CHURCHLAND AND TERRENCE J. SEJNOWSKI

Perspectives on Cognitive Neuroscience

PATRICIA S. CHURCHLAND AND TERRENCE J. SEJNOWSKI

How is it that we can perceive, learn and be aware of the world? The development of new techniques for studying large-scale brain activity, together with insights from computational modeling and a better understanding of cognitive processes, have opened the door for collaborative research that could lead to major advances in our understanding of ourselves.

NEUROSCIENCE AND COGNITIVE SCIENCE SHARE THE GOAL of trying to understand how the mind-brain works. In the past, discoveries at the neuronal level and explanations at the cognitive level were so distant that each often seemed of merely academic significance to the other. Symbol processing models based on the digital computer have been unpromising as a means to bridge the gap between neuroscience and cognitive science, because they did not relate to what was known about nervous systems at the level of signal processing. However, there is now a gathering conviction among scientists that the time is right for a fruitful convergence of research from hitherto isolated fields. The research strategy developing in cognitive neuroscience is neither exclusively from the top down, nor exclusively from the bottom up. Rather, it is a coevol-

P. S. Churchland is in the Department of Philosophy at the University of California at San Diego, La Jolla, CA 92093. T. J. Sejnowski is at The Salk Institute, La Jolla, CA 92093, and in the Department of Biology, University of California at San Diego, La Jolla, CA 92093.

tionary strategy, typified by interaction among research domains, where research at one level provides constraints, corrections, and inspiration for research at other levels (1).

Levels

There are in circulation at least three different notions of the term "levels," as it is used to describe scientific research, each notion carving the landscape in a different way—levels of analysis, levels of organization, and levels of processing.

Levels of analysis concern the conceptual division of a phenomenon in terms of different classes of questions that can be asked about it. A framework articulated by Marr and Poggio (2) drew upon the conception of levels in computer science and identified three levels: (i) the computational level of abstract problem analysis, decomposing the task into its main constituents (for example, determination of the three-dimensional structure of a moving object from successive views); (ii) the level of the algorithm, specifying a formal procedure to perform the task by providing the correct output for a given input; and (iii) the level of physical implementation. Marr (3) maintained that computational problems of the highest level could be analyzed independently of understanding the algorithm that performs the computation. Similarly, he thought the algorithmic problem of the second level was solvable independently of understanding its physical implementation.

Some investigators have used the doctrine of independence to conclude that neuroscience is irrelevant to understanding cognition.




Fig. 1. Structural levels of organization in the nervous system. The spatial scale at which anatomical organizations can be identified varies over many orders of magnitude. (Left) Drawing by Vesalius (33) of the human brain, the spinal column, and the peripheral nerves. (Right) Schematic diagrams illustrate (top) a processing hierarchy of visual areas in monkey visual cortex (34); (center) a small network model for the synthesis of oriented receptive fields of simple cells in visual cortex (35); and (bottom) the structure of a chemical synapse (36). Relatively little is known about the properties at the network level in comparison with the detailed knowledge we have of synapses and the general organization of pathways in sensory and motor system.

However, the independence that Marr emphasized pertained only to the formal properties of algorithms, not to how they might be discovered (4). Computational theory tells us that algorithms can be run on different machines and in that sense, and that sense alone, the algorithm is independent of the implementation. The formal point is straightforward: since an algorithm is formal, no specific physical parameters (for example, vacuum tubes or Ca^{2+}) are part of the algorithm. That said, it is important to see that the purely formal point cannot speak to the issue of how best to discover the algorithm used by a given machine, nor how best to arrive at the neurobiologically adequate task analysis. Certainly it cannot tell us that the discovery of the algorithms relevant to cognitive functions will be independent of a detailed understanding of the nervous system. Moreover, different implementations display enormous differences in speed, size, efficiency, and elegance. The formal independence of algorithm from architecture is something we can exploit to build other machines once we know how the brain works, but it is not a guide to discovery when we do not yet know how the brain works. Knowledge of brain architecture also can be the essential basis and invaluable catalyst for devising likely and powerful algorithms—algorithms that might explain how in fact the brain does its job.

Levels of organization. How do the three levels of analysis map onto the nervous system? There is organized structure at different scales: molecules, synapses, neurons, networks, layers, maps, and systems (5) (Fig. 1). The range of structural organization implies, therefore, that there are many levels of implementation and that each has its companion task description. But if there are as many types of task description as there are levels of structural organization, this diversity will be reflected in a multiplicity of algorithms that characterize how the tasks are accomplished. This in turn means that the notion of the algorithmic level is as oversimplified as the notion of the implementation level. Structure at every scale in the nervous system—molecules, synapses, neurons, networks, layers, maps, and systems (Fig. 1)—is separable conceptually but not detachable

physically. Psychological phenomena may be associated with a variety of levels. Some perceptual states such as the “raw” pain of a toothache might be a low-level effect, whereas attention may depend on a variety of mechanisms, some of which can be found at the level of local neural networks and others at the level of larger neural systems that reside in many different locations in the brain.

Levels of processing. This concept could be described as follows: The greater the distance from cells responding to sensory input, the higher the degree of information processing. Thus the level assigned is a function of synaptic distance from the periphery. On this measure, cells in the primary visual area of the neocortex that respond to oriented bars of light are at a higher level than cells in the lateral geniculate nucleus (LGN), which in turn are at a higher level than retinal ganglion cells.

Once the sensory information reaches the cerebral cortex it fans out through cortico-cortical projections into a multitude of parallel streams of processing. In the primate visual system 24 visual areas have been identified (6). Many (perhaps all) forward projections are accompanied by a back projection, and there are even massive feedback projections from primary visual cortex to the LGN. Given these reciprocal projections, it might seem that the processing levels do not really form a hierarchy, but there is a way to order the information flow by examining the layer of cortex into which fibers project. Forward projections generally terminate in the middle layers of cortex and feedback projections usually terminate in the upper and lower layers (7). However, we do not yet understand the function of these feedback pathways. If higher areas can affect the flow of information through lower areas, then the concept of sequential processing must be modified.

The hierarchical organization typical of earlier sensory areas is not adhered to everywhere. On the contrary, the anatomy of association areas and prefrontal cortex suggests a more “democratic” organization, and processing appears to take place in webs of strongly interacting networks (8). Decisions to act and the execution of plans and choices could be the outcome of a system with distributed control rather than a single control center. Coming to grips with systems having distributed control will require both new experimental techniques and new conceptual advances. Perhaps more appropriate metaphors for this type of processing will emerge from studying models of interacting networks of neurons.

Color Vision: A Case Study

As an illustration of fruitful interactions between psychology and physiology on a problem in perception, we have chosen several examples from color vision. Similar examples can also be found in the areas of learning and addiction (9) and sensory-motor integration (10, 11). Newton’s ingenious prism experiment demonstrated that white light can be decomposed into a mixture of wavelengths and recombined to recover the white light. This physical description of color, however, did not satisfy artists, who were well aware that the perception of color involved complex spatial and temporal effects. As Goethe pointed out in *Zur Farbenlehre*, dark shadows often appear blue. The physical description of color and the psychological description of color perception are at two different levels: The link between them is at the heart of the problem of relating brain to cognition. Three examples will be given to illustrate how such links are being made in color vision.

The knowledge that mixtures of only three wavelengths of light are needed to match any color led Young to propose in 1802 (12) that there are only three types of photoreceptors. Quite a different theory of color vision was later proposed by Hering, who suggested that color perception was based on a system of color opponents, one

for yellow versus blue, one for red versus green, and a separate system for black versus white (13). Convincing experiments and impressive arguments were marshaled by supporters of these two rival theories for nearly a century. The debate was finally settled by physiological studies proving that both theories were right—in different parts of the brain. In the retina three different types of color-sensitive photoreceptors were found, as predicted by Young, and the genes for the three cone photopigments have been sequenced (14). In the thalamus and visual cortex there are neurons that respond to Hering's color opponents (15). Evidently, even at this early stage of visual processing the complexity of the brain may lead to puzzles that can only be settled by knowing how the brain is constructed (16).


Recent progress in solving the problem of color constancy is a second example of converging physiological and psychological research. Red apples look red under a wide range of illumination even though the physical wavelengths impinging on the retina vary dramatically from daylight to interior lighting. Insights into color constancy have come from artists, who manipulate color contrasts in paintings, psychophysicists, who have quantified simultaneous contrast effects (17), and theorists, who have modeled them (18). Color constancy depends on being able to compute the intrinsic reflectance of a surface independently of the incident light. The reflectance of a patch of surface can be approximately computed by comparing the energy in wavelength bands coming from the patch of surface to the average energy in these bands from neighboring and distant regions of the visual field. The signature of a color-sensitive neuron that was performing this computation would be a long-range suppressive influence from regions of the visual field outside the conventional receptive field. Neurons with such color-selective surrounds have been reported in visual cortex area V4 (19); the first nonclassical

surrounds were found for motion-selective cells in area MT (20). If these neurons are necessary for color constancy then their loss should result in impairments of color vision. Bilateral lesions of certain extrastriate visual areas in man do produce achromatopsia—a total loss of color perception (21)—although this condition is usually found with other deficits and the damaged areas may not be homologous with area V4 in monkeys.

The third example of a link between brain and cognition comes from research on how form, motion, and color information are processed in the visual system. If different parts of the system are specialized for different tasks, for example, for motion or color, then there should be conditions under which these specializations are revealed. Suppose the "color system" is good at distinguishing colors, but not much else, and, in particular, is poor at determining shape, depth, and motion, whereas the "shape system" is not sensitive to color differences but to brightness differences. When boundaries are marked only by color differences—all differences in brightness are experimentally removed—shape detection should be impaired. Psychophysical research has shown that this is indeed the case. The perceived motion of equiluminant contours is degraded (22); form cues such as shape-from-shading are difficult to interpret (23), and perceived depth in random-dot stereograms collapses (24). Physiological and anatomical research has begun to uncover a possible explanation for these phenomena (25). The separate processing streams in cerebral cortex mentioned earlier carry visual information about different properties of objects (6, 26). In particular the predominant pathway for color information diverges from those carrying information on motion and depth (Fig. 2). The separation is not perfect, however, but equiluminant stimuli provide physiologists with a visual "scalpel" for tracking down the correlates of perceptual coherence in different visual areas.

The lessons learned from color perception may have significance for studying other cognitive domains. So far as we know only a small fraction of the neurons in the visual system respond in a way that corresponds to our perceptual report of color. The locations in the brain where links between physiological states and perceptual states can be found vary from the retina to deep in the visual system for different aspects of color perception (27). New experimental techniques will be needed to study these links when the information is encoded in a large population of interacting neurons (10, 28).

Fig. 2. Schematic diagram of anatomical connections and response selectivities of neurons in early visual areas of the macaque monkey. Visual information from the retina is split into two parallel streams at the level of the lateral geniculate nucleus (LGN), the parvocellular and magnocellular divisions. The parvocellular stream projects to two divisions of primary visual cortex (V1): the cytochrome oxidase-rich regions (Blob) and cytochrome oxidase-poor regions surrounding the blobs (Interblob). The magnocellular stream projects to layer 4B of V1. These three divisions of V1 project into corresponding areas of V2: the "thin stripe," "interstripe," and "thick stripes" of cytochrome oxidase-rich and -poor regions in V2. These areas in turn project to visual areas V3, V4, and MT (middle temporal area, also called V5). Heavy lines indicate robust primary connections, and thin lines indicate weaker, more variable connections. Dotted lines indicate connections that require additional verification. Not all projections from these areas to other brain areas are represented. The neurons in each visual area respond preferentially to particular properties of visual stimuli as indicated by the icons: Prism, tuned or opponent wavelength selectivity; Angle, orientation selectivity; Spectacles, binocular disparity selectivity or strong binocular interactions; Pointing hand, direction of motion selectivity. [Reprinted with permission from (6)]



Techniques and Research Strategies

Color vision is a problem that has been studied for hundreds of years; we know much less about the biological basis of other perceptual and cognitive states. Fortunately, new techniques, such as regional blood flow analysis with positron emission tomography (PET) and magnetic resonance imaging (MRI) are becoming available for noninvasively measuring brain activity in humans. With these techniques the large-scale pattern of what is happening where and when in the brain can be determined; later, as techniques with higher resolution are developed they can be focused on the relevant areas to ask how the processing is accomplished.

A useful way to get an overview of the assorted techniques is to graph them with respect to temporal and spatial resolution. This permits us to identify areas where there do not yet exist techniques to get access to levels of organization at those spatio-temporal resolutions and to compare their strengths and weaknesses (Fig. 3). For example, it is apparent that we lack detailed information about processing in neural networks within cortical layers and columns over a wide range of time scales, from milliseconds to hours. There is also a pressing need for experimental techniques designed to address the dendritic and synaptic level of investigation in cerebral cortex.

Without these data it will not be possible to develop realistic models of information processing in cortical circuits.

Although we need experimental data concerning the properties of neurons and behavioral data about psychological capacities, we also need to find models that explain how patterns of activity in neurons represent surfaces, optical flow, and objects; how networks develop and learn, store, and retrieve information; and how networks accomplish sensorimotor and other types of integration. Ideally, modeling and experimental research will have a symbiotic relationship, such that each informs, corrects, and inspires the other.

Although many diverse kinds of things are presented as models for some part of the nervous system, it is useful to distinguish between realistic models, which are genuinely and strongly predictive of some aspect of nervous system dynamics or anatomy, and simplifying models, which though not so predictive, demonstrate that the nervous system could be governed by specific principles. Connectionist network models (29), which are simplifying models, are typically motivated by cognitive phenomena and are governed primarily by computational constraints, while honoring very general neurobiological constraints such as number of processing units and time required to perform a task. Accordingly, they are more properly considered demonstrations of what could be possible and sometimes what is not possible. Realistic models of actual neural networks, by contrast, are primarily motivated by biological constraints, such as the physiological and anatomical properties of specific cell types (30). Despite their different origins and sources of dominant constraints, simplifying models and realistic neural models are both based on the mathematics of nonlinear dynamical


systems in high-dimensional spaces (31). The common conceptual and technical tools used in these models should provide links between two rich sources of experimental data, and consequently, connectionist and neural models have the potential to coevolve toward an integrated, coherent account of information processing in the mind-brain.

The ultimate goal of a unified account does not require that it be a single model that spans all the levels of organization. Instead the integration will probably consist of a chain of models linking adjacent levels. When one level is explained in terms of a lower level, this does not mean that the higher level theory is useless or that the high-level phenomena no longer exist. On the contrary, explanations will coexist at all levels, as they do in chemistry and physics, genetics, and embryology.

Conclusions

It would be convenient if we could understand the nature of cognition without understanding the nature of the brain itself. Unfortunately, it is difficult if not impossible to theorize effectively on these matters in the absence of neurobiological constraints. The primary reason is that computational space is consummately vast, and there are many conceivable solutions to the problem of how a cognitive operation could be accomplished. Neurobiological data provide essential constraints on computational theories, and they consequently provide an efficient means for narrowing the search space. Equally important, the data are also richly suggestive in hints

Fig. 3. Schematic illustration of the ranges of spatial and temporal resolution of various experimental techniques for studying the function of the brain. The vertical axis represents the spatial extent of the technique, with the boundaries indicating the largest and smallest sizes of the region from which the technique can provide useful information. Thus, single-unit recording can only provide information from a small region of space, typically 10 to 50 μm on a side. The horizontal axis represents the minimum and maximum time intervals over which information can be collected with the technique. Thus, action potentials from a single neuron can be recorded with millisecond accuracy over many hours. Patch-clamp recording allows the ionic currents through single ionic channels to be measured. Optical and fluorescent dyes that can reveal membrane potential, ionic concentrations, and intracellular structure have been used with high resolution in tissue culture, where it is possible to obtain a clear view of single cells (37, 38). However, recordings from the central nervous system are limited in resolution by the optical properties of nervous tissue and only about 0.1-mm resolution has been achieved (39). Confocal microscopy is a recent development in light microscopy that could be used for improving the resolution of the technique for three-dimensional specimens (40). ERP (evoked response potential) and MEG (magnetoencephalography) record the average electrical and magnetic activity over large brain regions and are limited to events that take place over about 1 s (41). The temporal resolution of PET (positron emission tomography) depends on the lifetime of the isotope being used, which ranges from minutes to an hour. It may be possible to achieve a temporal resolution of seconds with ^{15}O to study fast changes in blood flow by using temporal binning of the gamma ray events (equivalent to the poststimulus time histogram for action potentials) (42). The 2-deoxyglucose (2-DG) technique has a time resolution of about 45 min and a spatial resolution of 0.1 mm with large pieces of tissue and 1 μm with small pieces of tissue (43). The 2-DG technique can also be applied to humans with PET (44). Lesions



allow the interruption of function to be studied immediately after ablation of the tissue and for a long period of time after the ablation (21, 45). Microlesion techniques make selective modifications with substances such as ibotenic acid, which destroys neurons but not fibers of passage, and 4-amino-phosphonobutyric acid, which selectively and reversibly blocks a class of glutamate receptors (46). Video-enhanced light microscopy has opened a window onto dynamical activity within neurons, such as the recent visualization of axonal transport of organelles on microtubules (37, 47). All of the boundaries drawn show rough regions of the spatio-temporal plane where these techniques have been used and are not meant to indicate fundamental limitations.

concerning what might really be going on and what computational strategies evolution might have chanced upon. Moreover, it is by no means settled what exactly are the functional categories at the cognitive levels, and theories of lower level function may well be crucial to the discovery of the nature of higher level organization. Accordingly, despite the fact that the brain is experimentally demanding, basic neurobiology is indispensable in the task of discovering the theories that explain how we perform such activities as seeing, thinking, and being aware.

On the other hand, the possibility that cognition will be an open book once we understand the details of each and every neuron and its development, connectivity, and response properties is likewise misconceived. Even if we could simulate, synapse for synapse, our entire nervous system, that accomplishment, by itself, would not be the same as understanding how it works. The simulation might be just as much of a mystery as the function of the brain currently is, for it may reveal nothing about the network and systems properties that hold the key to cognitive effects. Even simulations of small network models have capabilities that are difficult to understand (32). Genuine theorizing about the nature of neurocomputation is therefore essential.

Many major questions remain to be answered. Although some problems in vision, learning, attention, and sensorimotor control are yielding, this will be harder to achieve for more complex psychological phenomena such as reasoning and language. Nonetheless, once we understand some fundamental principles of brain function, we may see how to reformulate the outstanding problems and address them in ways that are impossible now to predict. Whatever the outcome, the results are likely to surprise us.

REFERENCES AND NOTES

1. P. S. Churchland, *Neurophilosophy: Toward a Unified Science of the Mind-Brain* (MIT Press, Cambridge, MA, 1986); J. LeDoux and W. Hirst, *Mind and Brain: Dialogues in Cognitive Neuroscience* (Cambridge Univ. Press, Cambridge, 1986); S. M. Kosslyn, *Science* **240**, 1621 (1988).
2. The original conception of levels of analysis can be found in D. Marr and T. Poggio [*Neurosci. Res. Program Bull.* **15**, 470 (1977)]. Although Marr (3) emphasized the importance of the computational level, the notion of a hierarchy of levels grew out of earlier work by W. Reichardt and T. Poggio [*Q. Rev. Biophys.* **9**, 311 (1976)] on the visual control of orientation in the fly. In a sense, the current view on the interaction between levels is not so much a departure from the earlier views as a return to the practice that was previously established by Reichardt, Poggio, and even by Marr himself, who published a series of papers on neural network models of the cerebellar cortex and cerebral cortex. See, for example, D. Marr, *J. Physiol. (London)* **202**, 437 (1969); *Proc. R. Soc. London B* **176**, 161 (1970). The emphasis on the computational level has nonetheless had an important influence on the problems and issues that concern the current generation of neural and connectionist models [T. J. Sejnowski, C. Koch, P. S. Churchland, *Science* **241**, 1299 (1988)].
3. D. Marr, *Vision* (Freeman, San Francisco, 1982).
4. P. M. Churchland, *Dialogue* **21**, 223 (1982).
5. F. H. C. Crick, *Sci. Am.* **241**, 219 (September 1979); P. S. Churchland and T. J. Sejnowski, in *Neural Connections and Mental Computation*, L. Nadel, Ed. (MIT Press, Cambridge, MA, 1988); G. M. Shepherd, *Yale J. Biol. Med.* **45**, 584 (1972); P. Grobstein, *Brain Behav. Evol.* **31**, 34 (1988).
6. E. A. DeYoe and D. C. Van Essen, *Trends Neurosci.* **11**, 219 (1988).
7. J. H. R. Maunsell and D. C. Van Essen, *J. Neurosci.* **3**, 2563 (1983).
8. P. S. Goldman-Rakic, *Annu. Rev. Neurosci.* **11**, 137 (1988); V. B. Mountcastle, *Trends Neurosci.* **9**, 505 (1986).
9. T. H. Brown, P. Chapman, E. W. Kairiss, C. L. Keenan, *Science* **242**, 724 (1988); G. F. Koob and F. E. Bloom, *ibid.*, p. 715; L. R. Squire, *Memory and Brain* (Oxford Univ. Press, Oxford, 1987); E. R. Kandel et al., in *Synaptic Function*, G. M. Edelman, W. E. Gall, W. M. Cowan, Eds. (Wiley, New York, 1987), pp. 471-518; M. Mishkin and T. Appenzeller, *Sci. Am.* **256**, 80 (June 1987).
10. S. P. Wise and R. Desimone, *Science* **242**, 736 (1988).
11. S. G. Lisberger, *ibid.*, p. 728.
12. T. Young, *Philos. Trans. R. Soc. London* **92**, 12 (1802).
13. E. Hering, *Zur Lehre vom Lichtsinn* (Berlin, 1878).
14. J. Nathans, T. P. Piantanida, R. L. Eddy, T. B. Shows, D. S. Hogness, *Science* **232**, 203 (1986).
15. The possibility of a two-stage analysis for color vision was suggested as early as 1881, but no progress was made until physiological techniques became available for testing the hypothesis. There are still some issues that have not yet been fully resolved by this theory.
16. H. B. Barlow, *Q. J. Exp. Psychol.* **37**, 121 (1985).
17. D. Jameson and L. M. Hurvich, *J. Opt. Soc. Am.* **51**, 46 (1961).
18. E. H. Land, *Proc. Natl. Acad. Sci. U.S.A.* **83**, 3078 (1986); A. Hurlbert and T. Poggio, *Science* **239**, 482 (1988).
19. S. Zeki, *Neuroscience* **9**, 741 (1983); R. Desimone, S. J. Schein, L. G. Ungerleider, *Vision Res.* **25**, 441 (1985).
20. J. Allman, F. Miezin, E. McGuinness, *Annu. Rev. Neurosci.* **8**, 407 (1985).
21. A. R. Damasio, in *Principles of Behavioral Neurology*, M. M. Mesulam, Ed. (Davis, Philadelphia, 1985), pp. 259-288; O. Sacks, R. L. Wasserman, S. Zeki, R. M. Siegel, *Soc. Neurosci. Abstr.* **14**, 1251 (1988).
22. V. S. Ramachandran and R. L. Gregory, *Nature* **275**, 55 (1978).
23. P. Cavanagh and Y. Leclerc, *Invest. Ophthalmol. Suppl.* **26**, 282 (1985).
24. C. Lu and D. H. Fender, *ibid.* **11**, 482 (1972).
25. M. S. Livingstone and D. H. Hubel, *Science* **240**, 740 (1988).
26. S. Zeki, *Nature* **335**, 311 (1988).
27. D. Y. Teller and E. N. Pugh, Jr., in *Color Vision: Physiology and Psychophysics*, J. D. Mollon and L. T. Sharpe, Eds. (Academic Press, New York, 1983).
28. T. J. Sejnowski, *Nature* **332**, 308 (1988).
29. T. Kohonen, *Self-Organization and Associative Memory* (Springer-Verlag, New York, 1984); J. J. Hopfield and D. W. Tank, *Science* **233**, 625 (1986); S. Grossberg and M. Kuperstien, *Neural Dynamics of Adaptive Sensory-Motor Control* (North-Holland, Amsterdam, 1986); M. A. Arbib, *Brains, Machines and Mathematics* (McGraw-Hill, New York, ed. 2, 1987); D. E. Rumelhart and J. L. McClelland, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (MIT Press, Cambridge, MA, 1986).
30. T. J. Sejnowski, C. Koch, P. S. Churchland, *Science* **241**, 1299 (1988); C. Koch and I. Segev, *Methods in Neuronal Modeling: From Synapse to Networks* (MIT Press, Cambridge, MA, in press).
31. R. F. Abraham and C. D. Shaw, *Dynamics, the Geometry of Behavior* (Aerial Press, Santa Cruz, CA, 1982).
32. G. Edelman, *Neural Darwinism* (Basic Books, New York, 1987); R. A. Andersen and D. Zipsler, *Can. J. Physiol. Pharmacol.* **66**, 488 (1988); S. R. Lehky and T. J. Sejnowski, *Nature* **333**, 452 (1988).
33. A. Vesalius, *De Humani Corporis* (Brussels, 1543).
34. J. H. R. Maunsell and W. T. Newsome, *Annu. Rev. Neurosci.* **10**, 363 (1987).
35. D. H. Hubel and T. N. Wiesel, *J. Physiol. (London)* **160**, 106 (1962).
36. E. Kandel and J. Schwartz, *Principles of Neural Science* (Elsevier, New York, ed. 2, 1984).
37. S. J. Smith, *Science* **242**, 708 (1988).
38. J. Dodd and T. M. Jessell, *ibid.* **242**, 692 (1988); A. L. Harrelson and C. S. Goodman, *ibid.*, p. 692.
39. G. G. Blasdel and G. Salama, *Nature* **321**, 579 (1986); A. Grinvald, E. Lieke, R. D. Frostig, C. D. Gilbert, T. N. Wiesel, *ibid.* **324**, 361 (1986).
40. A. Fine, W. B. Amos, R. M. Durbin, P. A. McNaughton, *Trends Neurosci.* **11**, 346 (1988).
41. S. J. Williamson, G. L. Romani, L. Kaufman, I. Modena, *Biomagnetism: An Interdisciplinary Approach* (Plenum, New York, 1983); 2-DG: S. A. Hillyard and T. W. Picton, in *Handbook of Physiology*, section 1, *Neurophysiology*, F. Plum, Ed. (American Physiological Society, New York, 1987), pp. 519-584.
42. M. E. Raichle, *Trends Neurosci.* **9**, 525 (1986); M. I. Posner, S. E. Petersen, P. T. Fox, M. E. Raichle, *Science* **240**, 1627 (1988).
43. L. Sokoloff, *Metabolic Probes of Central Nervous System Activity in Experimental Animals and Man* (Sinauer, Sunderland, MA, 1984).
44. M. E. Phelps and J. C. Mazziotta, *Science* **228**, 799 (1985).
45. M. M. Mesulam, in *Principles of Behavioral Neurology*, M. M. Mesulam, Ed. (Davis, Philadelphia, 1985), pp. 125-168; B. Milner, in *Amnesia*, C. W. M. Whitty and O. Zangwill, Eds. (Butterworth, London, 1966), pp. 109-133; R. W. Sperry and M. Gazzaniga, in *Brain Mechanisms Underlying Speech and Language*, C. Millikan and F. Darley, Eds. (Grune and Stratton, New York, 1967), pp. 108-115; E. H. Land, D. H. Hubel, M. S. Livingstone, S. H. Perry, M. M. Burns, *Nature* **303**, 616 (1983); S. M. Kosslyn, J. D. Holtzman, M. S. Gazzaniga, M. J. Farrah, *J. Exp. Psychol. General* **114**, 311 (1985).
46. P. Schiller, *Nature* **297**, 580 (1982); J. C. Horton and H. Sherk, *J. Neurosci.* **4**, 374 (1984).
47. B. J. Schnapp and T. S. Reese, *Trends Neurosci.* **9**, 155 (1986).
48. We thank F. Crick whose insights as well as critical judgments were a major resource in writing this article. We also thank P. Churchland, R. Cone, C. Koch, and D. MacLeod for helpful suggestions. T.J.S. is supported by a grants from the National Science Foundation, the Seaver Institute, the Air Force Office of Scientific Research, the Office of Naval Research, and the General Electric Corporation; P.S.C. is supported by grants from the National Science Foundation and the James S. McDonnell Foundation. Portions of this article are based on our chapter in *Foundations of Cognitive Science* [P. S. Churchland and T. J. Sejnowski, in *Foundations of Cognitive Science*, M. Posner, Ed. (MIT Press, Cambridge, in press)].