

- phy of Mind*. Royal Institute of Philosophy Supplement 42. Cambridge: Cambridge University Press.
- Block, N., O. Flanagan, and G. Güzeldere, Eds. (1997). *The Nature of Consciousness: Philosophical Debates*. Cambridge, MA: MIT Press.
- Crick, F., and C. Koch. (1995). Are we aware of neural activity in primary visual cortex? *Nature* 375: 121–123.
- Davies, M., and G. W. Humphreys, Eds. (1993). *Consciousness: Psychological and Philosophical Essays*. Oxford: Blackwell.
- Dennett, D. C., and M. Kinsbourne. (1992). Time and the observer: The where and when of consciousness in the brain. *Behavioral and Brain Sciences* 15: 183–247.
- Metzinger, T., Ed. (1995). *Conscious Experience*. Paderborn: Schöningh.
- Nelkin, N. (1996). *Consciousness and the Origins of Thought*. Cambridge: Cambridge University Press.
- Peacocke, C. (1998). Conscious attitudes, attention and self-knowledge. In B. C. Smith, C. Wright, and C. Macdonald, Eds., *Knowing Our Own Minds*. Oxford: Oxford University Press.
- Rolls, E. T. (1997). Consciousness in neural networks? *Neural Networks* 10: 1227–1240.
- Schacter, D. L. (1989). On the relation between memory and consciousness: Dissociable interactions and conscious experience. In H. Roediger and F. Craik, Eds., *Varieties of Memory and Consciousness: Essays in Honor of Endel Tulving*. Hillsdale, NJ: Erlbaum.
- Shear, J., Ed. (1997). *Explaining Consciousness: The Hard Problem*. Cambridge, MA: MIT Press.
- Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge, MA: MIT Press.

Consciousness, Neurobiology of

After a hiatus of fifty years or more, the physical origins of CONSCIOUSNESS are being once again vigorously debated, in hundreds of books and monographs published in the last decade. What sparse facts can we ascertain about the neurobiological basis of consciousness, and what can we reasonably assume at this point in time?

By and large, neuroscientists have made a number of working assumptions that need to be justified more fully, in particular,

1. There is something to be explained, that is, the subjective content associated with a conscious sensation (what philosophers refer to as QUALIA; see also WHAT-IT'S-LIKE) does exist and has its physical basis in the brain.

2. Consciousness is one of the principal properties of the human brain, a highly evolved system; it must therefore have a useful function to perform. Crick and Koch (1995) assume that the function of visual consciousness is to produce the best current interpretation of the visual scene—in the light of past experiences—and to make it available for a sufficient time to the parts of the brain that contemplate, plan, and execute voluntary motor outputs (including language). This needs to be contrasted with the on-line systems that bypass consciousness but can generate stereotyped behaviors (see below).

3. At least some animal species (i.e., non-human primates such as the macaque monkey) are assumed to possess some aspects of consciousness. Consciousness associated with sensory events is likely to be very similar in humans and

monkeys for several reasons. First, trained monkeys behave as humans do under controlled conditions for most sensory tasks (e.g., visual motion discrimination; see MOTION, PERCEPTION OF; Wandell 1995). Second, the gross neuroanatomy of humans and nonhuman primates is the same, once the difference in size has been accounted for. Finally, MAGNETIC RESONANCE IMAGING in humans is confirming the existence of a functional organization very similar to that discovered by single-cell electrophysiology in the monkey (Tootell et al. 1996). As a corollary, it follows that language is not necessary for consciousness to occur (although it greatly enriches human consciousness). In the following, we will mainly concentrate on sensory consciousness, and, in particular, on visual consciousness, because it is experimentally the most accessible and the best understood.

Cognitive and clinical research demonstrates that much complex information processing can occur without involving consciousness, both in normals as well as in patients. Examples of this include BLINDSIGHT (Weiskrantz 1997), priming, and the implicit recognition of complex sequences (Velmans 1991; Berns, Cohen, and Mintun 1997). Milner and Goodale (1995) have made a masterful case for the existence of so-called on-line visual systems that bypass consciousness, and that serve to mediate relative stereotype visual-motor behaviors, such as eye and arm movements as well as posture adjustments, in a very rapid manner. On-line systems work in egocentric coordinate systems and lack both certain types of perceptual ILLUSIONS (e.g. size illusion) and direct access to WORKING MEMORY. Milner and Goodale (1995; see also Rossetti forthcoming) hypothesize that on-line systems are associated with the dorsal stream of visual information in the CEREBRAL CORTEX, originating in the primary VISUAL CORTEX (V1) and terminating in the posterior parietal cortex (see VISUAL PROCESSING STREAMS). This contrasts well with the function of consciousness alluded to above, namely, to synthesize information from many different sources and use it to plan behavioral patterns over time.

What is the neuronal correlate of consciousness? Most popular has been the belief that consciousness arises as an emergent property of a very large collection of interacting neurons (Popper and Eccles 1981; Libet 1995). An alternative hypothesis is that there are special sets of “consciousness” neurons distributed throughout cortex (and associated systems, such as the THALAMUS and the BASAL GANGLIA) that represent the ultimate neuronal correlate of consciousness (NCC), in the sense that activity of an appropriate subset of them is both necessary and sufficient to give rise to an appropriate conscious experience or percept (Crick and Koch 1995). NCC neurons would, most likely, be characterized by a unique combination of molecular, biophysical, pharmacological, and anatomical traits. It is also possible, of course, that all cortical neurons may be capable of participating in the representation of one percept or another, though not necessarily doing so for all percepts. The secret of consciousness would then consist of all cortical neurons representing that particular percept at that moment (see BINDING BY NEURONAL SYNCHRONY).

Where could such NCC neurons be found? Based on clinical evidence that small lesions of the intralaminar

nuclei of the thalamus (ILN) cause loss of consciousness and coma and that ILN neurons project widely and reciprocally into the cerebral cortex, ILN neurons have been proposed as the site where consciousness is generated (Bogen 1995; Purpura and Schiff 1997). It is more likely, however, that ILN neurons provide an enabling or arousal signal without which no significant cortical processing can occur. The great specificity associated with the content of our consciousness at any point in time could only be mediated by neurons in the cerebral cortex, its associated specific thalamic nuclei, and the basal ganglia. It is here, among the neurons whose very specific response properties have been extensively characterized by SINGLE-NEURON RECORDING, that we have to look for the NCC.

What, if anything, can we infer about the location of these neurons? In the case of visual consciousness, Crick and Koch (1995) surmised that these neurons must have access to visual information and project to the planning stages of the brain, that is, to premotor and frontal areas (Fuster 1997). Because in the macaque monkey, no neurons in primary visual cortex project to any area anterior to the central sulcus, Crick and Koch (1995) proposed that neurons in V1 do not directly give rise to consciousness (although V1 is necessary for most forms of vision, just as the retina is). Current electrophysiological, psychophysical, and imaging evidence (He, Cavanagh, and Intriligator 1996; Engel, Zhang, and Wandell 1997) supports the hypothesis that the NCC is not to be found among V1 neurons.

A promising experimental approach to locate the NCC has been the use of bistable percepts, that is, pairs of percepts, alternating in time, that arise from a constant visual stimulus as in a Necker cube (Crick and Koch 1992). In one such case, a small image, say of a horizontal grating, is presented to the left eye and another image, say of a vertical grating, is presented to the corresponding location in the right eye. In spite of the constant retinal stimulus, observers “see” the horizontal grating alternate every few seconds with the vertical one, a phenomenon known as “binocular rivalry” (Blake 1989). The brain does not allow for the simultaneous perception of both images.

It is possible, though difficult, to train a macaque monkey to report whether it is currently seeing the left or the right image. The distribution of the switching times and the way in which changing the contrast in one eye affects these times leaves little doubt that monkeys and humans experience the same basic phenomenon (Myerson, Miezins, and Allman 1981). In a series of elegant experiments, Logothetis and colleagues (Logothetis and Schall 1989; Leopold and Logothetis 1996; Sheinberg and Logothetis 1997) recorded from a variety of monkey cortical areas during this task. In early visual cortex, only a small fraction of cells modulated their response as a function of the percept of the monkey, while 20 to 30 percent of neurons in MT and V4 cells did. The majority of cells increased their firing rate in response to one or the other retinal stimulus with no regard to what the animal perceived at the time. In contrast, in a high-level cortical area, such as the inferior temporal cortex (IT), almost all neurons responded only to the perceptual dominant stimulus (in other words, a “face” cell only fired

when the animal indicated by its performance that it saw the face and not the sunburst pattern in the other eye). This makes it likely that the NCC is located among—or beyond—IT neurons.

Finding the NCC would only be the first, albeit critical, step in understanding consciousness. We also need to know where these cells project to, their postsynaptic action, and what happens to them in various diseases known to affect consciousness, such as schizophrenia or AUTISM, and so on. And, of course, a final theory of consciousness would have to explain the central mystery—why a physical system with a particular architecture gives rise to feelings and qualia. (Chalmers 1996).

See also ATTENTION; ATTENTION IN THE ANIMAL BRAIN; ATTENTION AND THE HUMAN BRAIN; SENSATIONS

—Christof Koch and Francis Crick

References

- Berns, G. S., J. D. Cohen, and M. A. Mintun. (1997). Brain regions responsive to novelty in the absence of awareness. *Science* 276: 1272–1275.
- Blake, R. (1989). A neural theory of binocular rivalry. *Psychol. Rev.* 96: 145–167.
- Bogen, J. E. (1995). On the neurophysiology of consciousness: 1. An overview. *Consciousness and Cognition* 4: 52–62.
- Chalmers, D. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.
- Crick, F., and C. Koch. (1992). The problem of consciousness. *Scientific American* 267(3): 153–159.
- Crick, F., and C. Koch. (1995). Are we aware of neural activity in primary visual cortex? *Nature* 375: 121–123.
- Engel, S., X. Zhang, and B. Wandell (1997). Colour tuning in human visual cortex measured with functional magnetic resonance imaging. *Nature* 388: 68–71.
- Fuster, J. M. (1997). *The Prefrontal Cortex: Anatomy, Physiology, and Neuropsychology of the Frontal Lobe*. 3rd ed. Philadelphia: Lippincott-Raven.
- He, S., P. Cavanagh, and J. Intriligator. (1996). Attentional resolution and the locus of visual awareness. *Nature* 383: 334–337.
- Leopold, D. A., and N. K. Logothetis. (1996). Activity changes in early visual cortex reflect monkeys’ percepts during binocular rivalry. *Nature* 379: 549–553.
- Libet, B. (1995). *Neurophysiology of Consciousness: Selected Papers and New Essays*. Boston: Birkhäuser.
- Logothetis, N., and J. Schall. (1989). Neuronal correlates of subjective visual perception. *Science* 245: 761–763.
- Milner, D., and M. Goodale. (1995). *The Visual Brain in Action*. Oxford: Oxford University Press.
- Myerson, J., F. Miezins, and J. Allman. (1981). Binocular rivalry in macaque monkeys and humans: A comparative study in perception. *Behav. Anal. Lett.* 1: 149–156.
- Popper, K. R., and J. C. Eccles. (1981). *The Self and Its Brain*. Berlin: Springer.
- Purpura, K. P., and N. D. Schiff (1997). The thalamic intralaminar nuclei: a role in visual awareness. *Neuroscientist* 3: 8–15.
- Rossetti, Y. (Forthcoming). Implicit perception in action: Short-lived motor representations of space evidenced by brain-damaged and healthy subjects. In P. G. Grossenbacher, Ed., *Finding Consciousness in the Brain*. Philadelphia: Benjamins.
- Sheinberg, D. L., and N. K. Logothetis. (1997). The role of temporal cortical areas in perceptual organization. *Proc. Natl. Acad. Sci. U.S.A.* 94: 3408–3413.

- Tootell, R. B. H., A. M. Dale, M. I. Sereno, and R. Malach. (1996). New images from human visual cortex. *Trends Neurosci.* 19: 481–489.
- Velmans, M. (1991). Is human information processing conscious? *Behavioral Brain Sci.* 14: 651–726.
- Wandell, B. A. (1995). *Foundations of Vision*. Sunderland, MA: Sinauer.
- Weiskrantz, L. (1997). *Consciousness Lost and Found*. Oxford: Oxford University Press.

Further Readings

- Crick, F., and C. Koch. (1998). Consciousness and neuroscience. *Cerebral Cortex* 8: 97–107.
- Jackendoff, R. (1987). *Consciousness and the Computational Mind*. Cambridge, MA: MIT Press.
- Zeki, S. (1993). *Vision of the Brain*. Oxford: Blackwell.

Constraint Satisfaction

A *constraint satisfaction problem* (CSP) defined over a *constraint network* consists of a finite set of *variables*, each associated with a *domain* of values, and a set of *constraints*. A *solution* is an assignment of a value to each variable from its domain such that all the constraints are satisfied. Typical constraint satisfaction problems are to determine whether a solution exists, to find one or all solutions, and to find an optimal solution relative to a given cost function. A well-known example of a constraint satisfaction problem is *k*-colorability, where the task is to color, if possible, a given graph with *k* colors only, such that any two adjacent nodes have different colors. A constraint satisfaction formulation of this problem associates the nodes of the graph with variables, the possible colors are their domains, and the inequality constraints between adjacent nodes are the constraints of the problem. Each constraint of a CSP may be expressed as a relation, defined on some subset of variables, the legal combinations of whose values are noted. Constraints can also be described by mathematical expressions or by computable procedures. Another typical constraint satisfaction problem is *SAT*isfiability, the task of finding the truth assignment to propositional variables such that a given set of clauses is satisfied. For example, given the two clauses $(A \vee B \vee \neg C)$, $(\neg A \vee D)$, the assignment of *false* to *A*, *true* to *B*, *false* to *C*, and *false* to *D*, is a satisfying truth value assignment.

The structure of a constraint network is depicted by a constraint graph whose nodes represent the variables and in which any two nodes are connected if the corresponding variables participate in the same constraint. In the *k*-colorability formulation, the graph to be colored is the constraint graph. In our SAT example the constraint graph has *A* connected to *D*, and *A*, *B*, and *C* are connected to each other.

Constraint networks have proven successful in modeling mundane cognitive tasks such as vision, language comprehension, default reasoning, and abduction, as well as in applications such as scheduling, design, diagnosis, and temporal and spatial reasoning. In general, constraint satisfaction tasks are computationally intractable (“NP-hard”; see COMPUTATIONAL COMPLEXITY).

ALGORITHMS for processing constraints can be classified into two interacting categories: (1) search and (2) consistency

inference. Search algorithms traverse the space of partial instantiations, while consistency inference algorithms reason through equivalent problems. Search algorithms are either systematic and complete or stochastic and incomplete. Likewise, consistency inference algorithms have either complete solutions (e.g., variable-elimination algorithms) or incomplete solutions (i.e., local consistency algorithms).

Local consistency algorithms, also called “consistency-enforcing” or “constraint propagation” algorithms (Montanari 1974; Mackworth 1977; Freuder 1982), are polynomial algorithms that transform a given constraint network into an equivalent, yet more explicit network by deducing new constraints to be added onto the network. Intuitively, a consistency-enforcing algorithm will make any partial solution of a small subnetwork extensible to some surrounding network. For example, the most basic consistency algorithm, called an “arc consistency” algorithm, ensures that any legal value in the domain of a single variable has a legal match in the domain of any other selected variable. A “path consistency” algorithm ensures that any consistent solution to a two-variable subnetwork is extensible to any third variable, and, in general, *i*-consistency algorithms guarantee that any locally consistent instantiation of *i* – 1 variables is extensible to any *i*th variable. Enforcing *i*-consistency is time and space exponential in *i*. Algorithms for *i*-consistency frequently decide *inconsistency*.

A network is *globally consistent* if it is *i*-consistent for every *i*, which means a solution can be assembled by assigning values using any variable ordering without encountering any dead end, namely, in a “backtrack-free” manner. However, it is enough to possess *directional* global consistency relative to a given ordering only. Indeed, an *adaptive consistency* (variable elimination) algorithm enforces global consistency in a given order only, such that every solution can be extracted, with no dead ends along this ordering. Another related algorithm, called a “tree-clustering” algorithm, compiles the given constraint problem into an equivalent tree of subproblems (Dechter and Pearl 1989) whose respective solutions can be efficiently combined into a complete solution. Adaptive consistency and tree-clustering algorithms are time and space exponential in a parameter of the constraint graph called an “induced-width” or “tree-width” parameter (Arnborg and Proskourowski 1989; Dechter and Pearl 1987).

When a problem is computationally hard for an adaptive consistency algorithm, it can be solved by bounding the amount of consistency enforcing (e.g., arc or path consistency), and by augmenting the algorithm with a search component. Generally speaking, search will benefit from network representations that have a high level of consistency. However, because the complexity of enforcing *i*-consistency is exponential in *i*, there is a trade-off between the effort spent on consistency inference and that spent on search. Theoretical and empirical studies of this trade-off, prior to or during search, aim at identifying a problem-dependent cost-effective balance (Haralick and Elliot 1980; Prosser 1993; Sabin and Freuder 1994; Dechter and Rish 1994).

The most common algorithm for performing systematic search is the *backtracking* algorithm, which traverses the