

A neurophilosophical slant on consciousness research

Patricia Smith Churchland*

UC President's Professor of Philosophy, Philosophy Department 0119, UCSD, La Jolla, CA 92093, USA

Abstract: Explaining the nature and mechanisms of conscious experience in neurobiological terms seems to be an attainable, if yet unattained, goal. Research at many levels is important, including research at the cellular level that explores the role of recurrent pathways between thalamic nuclei and the cortex, and research that explores consciousness from the perspective of action. Conceptually, a clearer understanding of the logic of expressions such as “causes” and “correlates”, and about what to expect from a theory of consciousness are required. The logic of some terms, such as “qualia” and “reductionism”, continues to generate misunderstandings about the scientific possibilities and limits. Experimentally, a deeper understanding of the role of the thalamus in coordinating activity across cortical levels, and a readiness to reconsider the orthodox approach to thalamocortical organization are also required.

The problem

The nature of consciousness is a problem at the interface of a range of disciplines: philosophy, psychology, neuroscience, anesthesiology, genetics, ethology, and evolutionary biology. Psychology helps us understand phenomena such as attention, conscious sensation, and declarative memory at the macrolevels. It has, for example, helped us learn that eye movements, made nonconsciously, nonetheless exhibit strategy and planning. On its own, however, psychological investigation is not enough, for we need also to understand mechanisms at the network and neuronal levels. Nonhuman animal studies, both behavioral and neurobiological, are essential to discerning which conscious capacities are shared and which not, and wherein lie the neurobiological differences.

The problem of consciousness in the 21st century is not the mind–body problem Descartes struggled

with. The classical mind–body problem was how the nonphysical stuff that makes up the immaterial soul can causally interact with the material stuff that is the body. No one, including of course Descartes, made the slightest progress in solving that problem. But we can see now that interaction is a pseudo problem, like the problem of how the crystal spheres of the heavens daily rotate, or how the heart concocts animals spirits. Compelling evidence implies the extreme improbability that thinking, feeling, and experiencing are events in a nonphysical soul. Rather, they are events of the entirely physical brain. Because one party to the alleged interaction almost certainly does not exist, interactionism is a nonproblem.

The contemporary mind/brain problem, therefore, is a nest of empirical questions about the brain: for example, what are the differences in the brain between being awake and being in deep sleep, and which of these differences explain being conscious when awake and not being conscious in deep sleep? How does that condition compare to the brain during absence seizures or during complex partial seizures? How much of decision-making is conscious, and what are

*Tel.: +1-858-822-1655; Fax: +1-858-534-8566;
E-mail: pschurchland@ucsd.edu

the differences between conscious and nonconscious stages of decision-making? What is the nature and origin of top-down attention and how does it work? What exactly happens in the brain when early skill acquisition slowly becomes a polished skill, performed automatically? This is, evidently, a partial and open-ended list, a list some of whose missing items may be unimagined, and perhaps unimaginable, given the current state of science. A truncated version of the contemporary problem is this: how can psychological phenomena be explained in neurobiological terms? Traditionally in philosophy of science, this is considered roughly equivalent to this question: how can psychology be reduced to neuroscience?

What is reductionism?

Disconnected from its root meaning, the word, “reductionism” has come to be freighted with plethora of very diverse, and often incompatible, emotional meanings. For example, it is sometimes used to mean that a research strategy must proceed “from the bottom up”. That is not part of its classical meaning. Rather than sort all that out, this section explains what is meant by reduction from the author’s point of view, and why that old bruised and misused word still has traction and utility (Churchland, 2002).

Reduction is a relation between scientific theories. To a first approximation, a science has achieved a reduction when the causal powers described at the macrolevel are explained as the outcome of events and processes at the lower level. Classically, reductions involve identifying functional macrocomponents that map onto structural microcomponents. Next, one assembles and tests hypotheses about how the components interact to produce the large scale effect. Explanation is considered achieved to the degree that the hypotheses survive tough experimental testing, and are consilient with other parts of well established science. That scenario may be repeated when the functions of the microcomponents are themselves decomposed into yet lower level structural constituents and their nanoactivities.

Thus thermodynamics is said to be reduced to statistical mechanics because temperature is explained in terms of the motion of the constituent molecules.

Theory of optics was reduced to theory of electromagnetic radiation, as we came to understand that light is, in fact, electromagnetic radiation, and inhabits the same explanatory framework as radio waves, micro-waves and X-rays. Typically, macro and micro-level theories co-evolve through time, as each provides tests, problems, and ideas for the other. Nature permitting, the two theories may come to knit together more closely, until the explanatory connections are so rich that scientists write their textbooks detailing the relevant range of macrophenomena as explanatorily brought to heel. In the 20th century, this was the profile of macro and microgenetics, and we see it robustly in progress in the co-evolution of embryology and molecular biology. The co-evolution of the cognitive sciences and the neurosciences, though in a very early stage, is filling the journals with remarkable, and often puzzling, discoveries.

One particular aspect of co-evolution of theories is worth dwelling on. As discoveries are made, it is inevitable that the descriptions of various phenomena are upgraded to reflect the discoveries. Consequently the meanings of the words in the descriptions undergo a parallel semantic evolution. As Francis Crick was fond of pointing out, the meaning of the word “gene” in 2000 is much richer and much different from its meaning in 1950. Depending on how the science goes, the semantic evolution mirroring the scientific evolution may be very dramatic. In such instances, people often speak of conceptual revolution. Thus one might well say that in the period from about the 1960 through to 1980 genetics underwent a conceptual revolution. This also marked geology as it came to appreciate that the continents drift on a surface of magma. The idea that heat was not caloric fluid, or any kind of stuff at all, but merely the *motion* of molecules, was a conceptual revolution that occurred in the 19th century. More generally, as the explanatory exoskeleton emerges — that is, as the basic principles are discovered and put into the theoretical framework — quite radical changes can occur. For this is the period when folk ideas are gradually replaced by scientific ideas, and in turn, early scientific ideas are replaced by more mature hypotheses. This is the period when the ostensibly obvious gets wrecked on the shoals of scientific discovery.

One simple example of semantic evolution concerns fire. In the middle ages, formulating a precise

definition of “fire” was not possible, since physicists did not know what it really was. Certainly no scientist could say that it was rapid oxidation, since nothing was known as about oxygen as an element or about its role in the burning of wood. Had a kind alien left a message on a monk’s pillow, “guess what — burning is rapid oxidation”, no one would have known what the message could possibly mean. Such precision as there was in defining “fire” consisted in grouping together a range of phenomena, all considered paradigmatic instances of fire: the sun, comets, burning of wood, lightning, northern lights, and fireflies. The criteria were drawn from what seemed observably obvious: they all emit light or heat or both. As it happened, burning of carbon material was the first to be understood, thanks to Lavoisier and Priestly.

The reality behind what was supposed to be observably obvious turned out to be surprising: the items in the list are fundamentally different from each other and are subsumed by very different scientific subfields. The heat and light from the sun is the result not of oxidation, but of nuclear fusion; lightning is actually thermal emission; the fireflies’ display is based on biophosphorescence; comets are balls of ice reflecting sunlight; northern lights are the result of spectral emission. These items are not part of the same family at all. From the vantage point of 21st century science, the medieval categorization might seem a bit foolish. Because we learn contemporary science as children, that current science becomes second nature to us — it seems dead obvious. The medieval category was not owed to foolishness but merely to ignorance. When you do not really understand the nature of a phenomenon, you try to do justice to what you think is observably obvious.

In terms of scientific maturity, neuroscience is still wet behind the ears. More exactly, neuroscience is still in search of its basic explanatory exoskeleton — of the fundamental principles that explain how nervous systems work. Although an enormous amount is known about the molecular aspects of individual neurons (their structure and their function), the fundamentals characterizing how macro effects emerge from populations of neurons are still largely mysterious.

Can we expect conceptual revolutions? The progress of science is impossible to predict with much

accuracy, but given the history of science, simple prudence suggests that we currently misconceptualize many problems because we do not have the scientific understanding to generate and render precise the appropriate concepts. My hunch is that consciousness is probably a case in point. For example, it sometimes seems observably obvious that consciousness is a single, unitary phenomenon, that you either have it or you do not, that it is like a light — it is either off or on. But these seemingly obvious ideas may turn out to be quite wrong. With an even higher degree of uncertainty, it is conjectured that as we come to understand how *time* is managed and represented and used, and to understand the function of the many emerging and fading intrinsic rhythms displayed by individual neurons and by populations of neurons, some of the explanatory exoskeleton will begin to be discernible (Jahnsen and Llinas, 1984; Sherman and Guillery, 2001; Massimini et al., 2004). Were a kindly alien to leave a message on one’s pillow concerning the true neurobiological nature of conscious phenomena, one doubtless could not make much sense of the message.

What are the target phenomena?

The list of phenomena commonly rated as instances of “consciousness” turns out to be diverse and somewhat puzzling. Someone in coma following an epileptic seizure or a blow to the head, for example, is not conscious; someone not paying attention to a mild hunger is not conscious of that hunger; someone in deep sleep is not conscious in the way that he is conscious when awake. Stupor differs from coma in that patients can be aroused from unresponsiveness with vigorous stimuli; patients in persistent vegetative state (PVS) show sleep–wake cycles but are totally nonresponsive to external stimuli, however vigorous; coma patients are totally nonresponsive and exhibit no sleep–wake cycles. Evidently, these are quite different ways of being “not conscious” (Plum and Posner, 1982). Correlatively, being awake, paying attention, explicitly remembering, emerging from anesthesia, smelling mint, regaining awareness following an absence seizure, regaining awareness following a complex partial seizure, being aroused from sedation, dreaming and hallucinating, are, quite possibly, different ways of being conscious.

Executing an intentional action, such as picking up a hammer and pounding in a nail seems to involve consciousness of the action, though many aspects of the action sequence are conscious in one sense, but not in another. For the skilled carpenter, less attention is paid to the details of the hammering movement, and more to planning and organizing the next set of complex actions. Automatized skills are exercised consciously, but not in the way that the novice is conscious of his unskilled action. How formed our decisions are when we become aware of them (and think of ourselves as “creating” them), is unclear. For example, it is known that nonconscious antecedents to an intention, as indicated by the “readiness potential” measured by the EEG, precede a “freely-chosen” action by up to a second and a half (Libet, 1985).

We are aware of emotions, such as being angry or sad, of drives, such as sexual lust, hunger, and curiosity; of the passage of time and of spatial depth and layout. We are aware that a scene is unfamiliar, that we feel dizzy, that we are falling, that we need to pass water. Efference copy allows one to be aware that a movement is one’s movement, not the world’s movement (Churchland, 2002). Normally, we are aware of our body as our own (Damasio, 1999). Normally we are aware of our thoughts as our own; on one hypothesis, schizophrenic patients lack this sense of ownership, and attribute their thoughts to external agents, such as God, the devil, or the FBI (Frith, 1992).

Why stress the diversity of conscious phenomena? For several reasons. First, because much of the literature assumes without discussion that at bottom there is only one basic phenomenon, with various features that come and go (Chalmers, 1996). From a neural point of view, that assumption may be quite wrong. Like the medieval physicists thinking about fire, we just do not know enough to put much confidence in the “single thing” assumption about consciousness. There may be a set of inter-related but semi-independent processes, some may share some background conditions, but not others. Just as coma is different from persistent vegetative state, so paying attention is likely to be different from waking from deep sleep or emerging from an absence seizure.

The second reason for stressing varieties of conscious phenomena is that the set and its diverse

subsets remind us that there are many entry points to the puzzle and many ways of attacking the problem. At this early stage, it may be unwise to assume that one subset is more paradigmatic, more privileged and prototypical, and more accessible to scientific investigation, than others. It may be unrewarding to assume that some items in the list are really only background conditions, not consciousness per se. Visual experiences, for example, have sometimes been regarded as prototypically conscious. Not enough is known about how the brain works to give this idea unquestioned approval, and “intuitions” are neither reliable nor uniform across scientists. And of course some humans are completely blind, and some species, such as the star-nosed mole, have no vision whatever. Where the breakthroughs will come is anyone’s guess; efference copy, for example, may be a more promising entry point than visual perception. Consequently, diverse research strategies targeting distinct items on the list, with some attempts to coordinate across approaches, is probably appropriate at this stage.

What counts as a theory of conscious phenomena?

The word “theory” can be used in many ways, from something that is a loose hunch about the causes of a poorly defined phenomena (e.g., the “theory”, circa 1950 that autism is caused by cold mothering), to fact-based speculations about a moderately well-defined phenomena (Wegener’s postulation of continental drift) to theories that specify mechanism and either mesh with other parts of established science or else testably challenge other parts of science (the Hodgkin–Huxley theory of the action potential in a neuron; the theory of how proteins get produced).

For the purposes at hand, it is useful to adopt the convention that something will be considered a theory of consciousness if, like the theory of how a neuron produces an action potential, it explains the main properties in sufficient detail that the following are satisfied: (1) we understand how macroevents emerge from the properties and organization of the microevents, (2) novel phenomena can be predicted, (3) the system can be manipulated, and (4) it is clear at what level of brain organization the phenomenon resides.

According to this convention, therefore, we assume that a genuine explanation of the properties of conscious phenomena must characterize neurobiological mechanisms. A theory satisfying the four desiderata will not be solely a psychological level account relating various cognitive functions (i.e., not just an array of boxes, labeled as cognitive, with arrows connecting the boxes). Boxology at the psychological level is a crucially important step, but it does not explain the neurobiological bases for the functions in the boxes. Likewise, an explanatory theory will not just consist of detailed anatomical maps of what projects to what, though such maps are also essential to the solution. It is also assumed that finding correlations — perhaps via fMRI or single cell recordings — for certain conscious events does not as such constitute a theory, because such correlations do not, ipso facto, explain mechanism. Notice, moreover, that X can be correlated with Y for a range of reasons: X causes Y, Y causes X, they have a common cause, or X and Y are actually the same thing under different descriptions; i.e., $X = Y$. Discovery of identities (i.e., that $X = Y$) is typically needed to make the crucial step towards theoretical authority. Correlating events using different measuring instruments, such as fMRI and behavioral reports, can be extremely useful, certainly, but a roster of correlations does not constitute a theory.

The semantic convention proposed entails a fairly strong requirement for “theory-hood”. It is meant to be strong in order to emphasize the importance of testability, predictability, and consilience with other parts of science in general and neuroscience in particular. It is meant to require a theory of consciousness to be comparably powerful to the theory that light is electromagnetic radiation, or the theory of the action potential or the theory that DNA codes for proteins. Without this sort of guideline, just about everybody and his dog lay claim to a theory of consciousness.

In this strong sense, neuroscience has not yet produced a theory of any of the various conscious phenomena aforementioned. What has emerged over the last ten years, however, are fruitful prototheories regarding some aspect or other that highlight some feature(s) at some level of brain organization as being fundamental. Other names for prototheory might be “general approach” or “line of attack”.

Prototheories are more speculative, and harbor dark regions where explanation is merely hand-waving. This is not only tolerated but applauded in the hope that prototheories will mature into explanatorily competent theories as they are prodded, pushed, pounded, and goaded into experimental test (For a range of prototheories, see articles in Baars et al., 2003).

Two philosophical objections

Two further matters should be addressed concerning what we can expect from a theory. First, a common philosophical complaint is that any neurobiological theory of consciousness will always leave something out — something crucial. It will always leave out the feeling itself — the feeling of what it is like to be aware, to see blue, smell mint, and so on (Nagel, 1974; Chalmers, 1996). These are so-called qualia — the experiences themselves — and these are what are important about consciousness. Pursuing this point further, the philosopher may go on to conclude that no science can ever really explain qualia because it cannot demonstrate what it is like to see blue if you have never seen blue; consciousness is forever beyond the reach of scientific understanding.

What is the merit in this objection? It is lacking merit, for if you look closely, you will find that it rests on a misunderstanding. The argument presumes that if a conscious phenomenon, say smelling mint, were genuinely explained by a scientific theory, then a person who understood that theory should be *caused to have that experience*; e.g., should be caused to smell mint. Surely, however, the expectation is unwarranted. Why should anyone expect that understanding the theory must result in the production of the phenomenon the theory addresses? Consider an analogy. If a student really understands the nature of pregnancy by learning all there is to know about the causal nature of pregnancy, no one would expect the student to become pregnant thereby. If a student learns and really understands Newton’s laws, we should not expect the student, like Newton’s fabled apple, to thereby fall down.¹ To smell mint, a certain range of neuronal activities have to obtain,

¹This example is owed to Ed Hubbard.

particularly, let us assume, in olfactory cortex. Understanding that the olfactory cortex must be activated in manner β will not itself activate the olfactory cortex in manner β . We are asking too much of a neuroscientific theory if we ask it not only to explain and predict, but also to *cause* its target phenomenon, namely the smell of mint, simply by virtue of understanding the theory.

A second and related complaint raised by certain philosophers is that even if neuroscience were to discover with what brain states being aware of a burning pain on one's left ear is identical, we would still not understand why just *those* brain states are identical with precisely *that* sensation, as opposed, say, to feeling a desire to void. Neuroscience, it will be averred, will never be able to explain why conscious states $Y =$ brain states X , rather than say, brain state Z . For those who are keen on qualia as metaphysical simples forever beyond the scope of science, the next step may be to infer that we cannot ever hope to understand that identity in neurobiological terms (Chalmers, 1996). Awareness, the claim goes, will always be ineffable and metaphysically basic. This means neuroscience cannot ever really explain consciousness.

This complaint too rests on a misunderstanding. What is an example where a science — any subfield of science — explains why $X = Y$? Not how we *know* or why we *believe* that $X = Y$, but why X *is* identical to Y , rather than to Z . Using the examples already at hand, the corresponding questions would be these: why is temperature mean molecular kinetic energy, rather than, say, caloric fluid or something else entirely? Why is visible light actually electromagnetic radiation rather than, say, something else entirely, say, “intrinsic photonicness”? By and large science does not offer explanations for fundamental identities. Rather, the discovery is that two descriptions refer to one and the same thing — or that two different measuring instruments are in fact measuring one and the same thing. Why is that thing, the thing it is? It just is. Science discovers fundamental identities, but the identities it discovers just are the way things are. There is no fundamental set of laws from which to derive that temperature is mean molecular kinetic energy or light is electromagnetic radiation.

Reflection shows this logical point to be acknowledged in an everyday setting. If someone discovers

that The Morning Star (Venus) is identical to The Evening Star (Venus), he will explain why he believes this by citing his evidence. But if asked, “why is The Morning Star (Venus) identical to The Evening Star (Venus)”, no answer is appropriate; that is just the way the world is. The question itself is based on the false assumption that identities ought to follow from general laws. But they don't. We may get an explanation of why people mistakenly thought what they saw in the dusk was not the same as the planet they saw in the dawn, and how they came to realize that what they saw in the dusk is identical to what they saw in the dawn. There is, however, no explanation of why Venus is Venus; of why the Morning Star is identical to the Evening Star. It just is. Or, to put it as the medieval philosophers sagely noted, everything is what it is, and not another thing. Correspondingly, assuming we discover that a certain pattern of activity Naj is identical to smelling mint, there will be no further explanation of why *that* pattern of activity is identical to smelling mint (why $Naj = C$).

Such merit as there is in the complaints probably comes merely to this: given the current state of neuroscience, it is very hard to predict what the explanation of conscious phenomena will look like — *very* hard. But so what? It is always hard to predict the course of a science, and especially hard to predict what an immature science will look like when it matures.

Theories of consciousness: Where is the action?

During the last decade, research targeting the problems of consciousness has intensified. By and large, the efforts have been directed toward cortical regions, cortical pathways, and cortical activity. Under the characterization, “seeking the *neural* correlates of consciousness”, the research has largely been looking for the *cortical* correlates of consciousness. Clinical studies of human patients with cortical lesions have inspired this line of attack, owing to suggestive correlations between deficits in specific kind of experiences and region-specific lesions. Correlations such as (a) middle temporal lesions and loss of visual experience of motion, (b) ventral stream lesions resulting in the inability visually to detect shapes, and (c) fusiform lesions resulting in

prosopagnosia, have been extremely important in providing a guidance for this research. The ease of imaging cortical activity with fMRI has probably also played some role in the focus on the cortex. Finally, because humans have proportionally more cortex than our closest relative, and because humans regard themselves as conscious beings *par excellence*, it may seem to follow that the cortex holds the answers.

True enough, human cortex is really big, relative to the thalamus, striatum, etc. But the logic in the “it must be cortex” argument is seriously flawed. Almost seventy years ago, neurosurgeons Wilder Penfield and Edwin Boldrey made a cautionary point regarding the exclusive focus on the cortex: “All parts of the brain may well be involved in normal conscious processes, but the indispensable substratum of consciousness lies outside of the cerebral cortex, probably in the diencephalon [thalamus]” (1937, p. 241).

One drawback to cortical chauvinism is that it tends concentrate on conscious *perception*, to the neglect of the perspective of behavior. This focus tends to blind us to the root “motocentric” basis of conscious phenomena. Animals are movers, and nervous systems earn their keep by servicing movement. Other things being equal (and there are a lot of other things), the better and faster the brain’s predictive capacities relative to the animal’s *modus vivendi*, the better the organism’s behavioral portfolio in the cut-throat competition to survive and reproduce. (Churchland, et al., 1994; Allman, 1999; Damasio, 1999; Llinas, 2001; Guillery, this volume). Hence we do well to keep in mind that moving, planning, deciding, executing plans in behavior, and more generally, keeping the body alive, is the fundamental business of the brain. Cognition and consciousness are what they are, and have the nature they have, because of their role in servicing behavior. Evolution just works that way.

Broadly speaking, the solution found by evolution to the problem of prediction is to modify motor programs (or fixed action patterns) by sensory information. The value of the sensory impact is greater if it can signal *me-relevant* causal regularities between events. To achieve this, the system needs neural populations, interposed between sensory receptors and motor neurons, to find and embody higher-order causal regularities. The richer the interposed

neuronal resources, the more sophisticated the statistical capacities and the greater the isomorphisms achievable between the brain’s categorical/causal maps and the world’s categorical/causal structures (Churchland and Churchland, 2002). Importantly, much of the brain’s input is consequent upon the organism’s own movements, exploratory and otherwise. This dynamical loop extracts vastly more information, in a given time interval, about the causal properties of the external world than could a purely passive system.

To acquire predictive prowess is to acquire skills regarding the causal structure of the world. And the essential thing about causal knowledge, is that *time* is at its heart.² Predicting durations, interception intervals, velocities, and speeds of the agent’s own various body movements are everywhere critical. Hence memory of durations, velocities, and self-movement parameters are everywhere critical. On one construal, the most fundamental problem for a nervous system, is *how to get the timing right*; that is, how to interact with the world so as to succeed in the four Fs: feeding, fleeing, fighting, and reproducing.³

Skill in these functions is not a matter of passive observation but of interaction with the world, and prediction of what will happen next can be greatly improved by having access to the just-issued motor commands. Efference copy, given this perspective, is one of evolution’s clever solutions to tuning the timing. What is inspiring about the Sherman and Guillery approach (2001; Guillery and Sherman, 2002), is their insight that efference copy is essential information relied upon to tune up the synapses to get the timing right. Efference copy is crucial to enabling the brain to get beyond the fixity of “fixed action patterns”, to flexibility and adaptivity in planning, interacting and predicting. This insight entails, among other things, that we seriously rethink efference copy — where the signals go, how they modify sensory processing, how they prepare the nervous system for what’s next, and how they contribute to conscious phenomena.

²I owe much of the discussion on time and temporal properties to the insights of Rick Grush.

³As Paul Maclean puts it.

According to the Guillery and Sherman hypothesis, all messages to the thalamus and cortex, including the ostensibly “pure” sensory signals, carry information about ongoing instructions to motor structures. At first this may sound puzzling, accustomed as we are to the conventional wisdom that “the sensory pathways are *purely* sensory”. However, the hypothesis is profoundly right. Simplified, part of the point is this: axons from the sensory periphery have collateral branches that go to motor structures. Consequently, as a developing organism begins to interact with the world, a sensory signal becomes also a prediction about what movement will happen next; thus, as the animal learns the consequences of *that* movement, it learns about what in the world will probably happen next, and hence what it might do after that.

Loops between thalamic and cortical structures are probably the substrate for embodying (“ensynapsing”, one might say) the temporal/causal properties of the world, and also the temporal/causal portfolio of one’s own body. Eyes move faster than legs, head movements are faster than whole body movements, and some of these parameters can change a little or a lot with maturation, practice, anticipation, and changes in emotional state. The thalamus, given its connectivity to all cortical structures and to other subcortical structures, and given the physiology of purported driver and modulator cells, looks particularly well-suited to negotiate temporality in all its diverse aspects — in learning, ongoing prediction, attentional shifts to different sensory — motor tasks, calling up stored timing information — to getting the timing right.

The temporal is endemic to all conscious phenomena, probably because the temporal is the *sine qua non* of movement. Consequently, it is suspected that the secrets of consciousness are embodied in the thalamocortical, thalamostriatal, thalamo-brainstem loops. (See also Groenewegen and Berendse, 1994; Purpura and Schiff, 1997; Damasio, 1999; Llinas, 2001). Admittedly, this anatomical vista includes a lot of brain territory, but with a revision of the prevailing assumptions concerning “pure vision” and “visual hierarchies,” the deeply inobvious may become experimentally quite accessible. (Guillery and Sherman, 2002).

Concluding remarks

Conscious phenomena are under study at many different levels of brain organization, using many different lines of attack. So far, however, no explanatorily competent theory has yet emerged. One obstacle is that most efforts are essentially visuo-centric. Another is that exploration is typically limited to cortex, disregarding subcortical activity as merely part of the background conditions. This may be like missing the significance of finding fossils on a mountaintop or the difference in color between arterial and venous blood.

Until the *known* anatomy and physiology of thalamocortical connections are better appreciated, and until more of the *unknown* anatomy and physiology is revealed, neuroscientists are unlikely to see the advantages of addressing conscious phenomena from the perspective of the motor organization. In particular, we need to determine the significance of the vast number of projections from cortical layer 5 to the thalamus (Guillery and Sherman, 2002). By shifting perspective from “visuocentricity” to “motor–sensory-centricity”, the singular importance of temporality takes center stage. (See also Casagrande, this volume; Colby, this volume.) In turn, this shift engenders the hunch that “time management”, for want of a better term, is the key to the complex job portfolio of thalamic nuclei, and very probably the key to a range of conscious phenomena as well.

References

- Allman, J.M. (1999) *Evolving Brains*. Scientific American Library, New York.
- Baars B., Banks, W. and Newman, J. (Eds.) (2003) *Essential Sources in the Scientific Study of Consciousness*. MIT Press, Cambridge, MA.
- Casagrande, V. (2005) Constructing visual reality: the impact of attention and motor planning on the lateral geniculate nucleus. *Prog. Br. Res.*, (this volume).
- Chalmers, D. (1996) *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press, New York.
- Churchland, P.S. (2002) *Brain-Wise: Studies in Neurophilosophy*. MIT Press, Cambridge, MA.
- Churchland, P.S. and Churchland, P.M. (2002) Neural worlds and real worlds. *Nature Reviews Neuroscience*, 3: 903–907.

- Churchland, P.S., Ramachandran, V.S. and Sejnowski, T.J. (1994) A critique of pure vision. In: Koch, C. and Davis, J.L. (Eds.), *Large-Scale Neuronal Theories of the Brain*. MIT Press, Cambridge, MA, pp. 23–60.
- Colby, C. (2005) Corollary discharge and spatial updating; when the brain is split is space still unified? *Prog. Br. Res.*, (this volume).
- Damasio, A.R. (1999) *The Feeling of What Happens*. Harcourt Brace, New York.
- Frith, C.D. (1992) *The Cognitive Neuropsychology of Schizophrenia*. Lawrence Erlbaum and Assoc., Hillsdale, N.J.
- Groenewegen, H.J. and Berendse, H.W. (1994) The specificity of the “nonspecific” midline and intralaminar thalamic nuclei. *Trends in Neurosci.*, 17: 52–57.
- Guillery, R.W. and Sherman, S.M. (2002) The thalamus as a monitor of motor outputs. *Philos. Trans. R Soc. Lond. B Biol. Sci.*, 357: 1809–1821.
- Jahnsen, H. and Llinas, R. (1984) Electrophysiological properties of guinea pig thalamic neurones: An in vitro study. *Journal of Physiology (London)*, 349: 205–226.
- Libet, B. (1985) Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Sciences*, 8: 529–566.
- Llinas, R.R. (2001) *I of the Vortex: From Neurons to Self*. MIT Press, Cambridge, MA.
- Massimini, M., Huber, R., Ferrarelli, F., Hill, S. and Tononi, G. (2004) The sleep slow oscillation as a traveling wave. *The Journal of Neuroscience*, 24: 6862–6870.
- Nagel, T. (1974) What is it like to be a bat? *Philosophical Review*, 83: 435–450.
- Penfield, W. and Boldrey, E. (1937) Somatic, Motor and Sensory Representation in the Cerebral Cortex of Man as Studied by Electrical Stimulation. *J. Bale, Sons & Curnow*, London.
- Plum F. and Posner, J. (1981) *The Diagnosis of Stupor and Coma*, 3rd Edition. Oxford University Press, New York.
- Purpura, K.P. and Schiff, N.D. (1997) The thalamic intralaminar nuclei: A role in visual awareness. *The Neuroscientist*, 3: 8–15.
- Sherman, S.M. and Guillery, R.W. (2001) *Exploring the Thalamus*. Academic Press, San Diego, CA.