validity. Outside the laboratory attention sweeps broadly across extensive territories – landscapes, mountain ranges, panoramas of all sorts. Whereas a restricted set of cues may support automatized responding, a panorama rarely does.

Would behavior be affected if consciousness were lacking? The plight of the blindsight victim cited by Velmans implies that voluntary behavior would grind to a halt. The blindsight patient can discriminate between this and that, but under what conditions? When bullied by the experimenter to respond although he sees nothing. Left to himself (in life beyond the laboratory) he initiates no behavior whatever with respect to information in the affected visual half field. Were his difficulty to extend across all sensory modalities, he would not voluntarily respond to anything at all.

When conscious representations are necessary for behavior and when they are not (and why) is a matter for debate and research. But to sweep consciousness aside as epiphenomenal (as irrelevant to behavior as the hiss of escaping steam is to the movement of the locomotive) exemplifies a blinkered and excessively focalized vision.

Is consciousness information processing?

Raymond Klein

Department of Psychology, Dalhousie University, Halifax, Nova Scotia B3H 4J1, Canada

Electronic mail: *klein@ac.dal.ca*

Velmans's target article asks whether human information processing is conscious. He answers that consciousness is not necessary for human information processing of any type to occur and that consciousness can play no causal role in any human information processing. In spite of his claim, he also concludes that consciousness can be studied scientifically and must be to achieve a complete psychology. Surely Velmans cannot eat his cake and have it, too. A mental entity without functions might be interesting to study, but it could not be important to study. Perhaps there is something amiss with Velmans's view of the relationship between information processing and consciousness.

Velmans reviews evidence consistent with the conclusion that consciousness is not necessary for human information processing stages from input to output, and more strongly, that because it receives only the outputs of these stages consciousness can never influence them. One might dispute the way Velmans uses individual studies to support this argument (e.g., I would disagree with his interpretations of Libet, 1985, Nissen & Bullemer, 1987, and other studies), but I'm sure other commentators will develop this line of attack. My main point is apparent if we turn Velmans's question around and ask, "Is consciousness information processing?" The answer is (indeed, unless we resort to dualism, must be): Yes, consciousness is a form of human (and perhaps nonhuman as well) information processing. Although it is not identical to focal attention and primary memory, consciousness is closely related to these functions. In relation to input processing, a relatively unified conscious perception/experience is achieved when selective attention functions as a gateway controlling those aspects of the environment that are likely to enter awareness, likely to be stored in memory, and likely to be responded to. In relation to output processing, voluntary behavior is activated by a selective process which chooses from amongst competing action patterns (cf. Shallice 1972) those that are most likely to achieve a consciously experienced goal state.[1] Velmans cites this pretheoretical notion that identifies conscious awareness with focal attention and he finds reasons (I'm not convinced by the evidence he cites here) for rejecting it. Yet Velmans himself fails to reject it completely: "A complete dissociation of consciousness from focal-attentive processing is difficult to achieve, as the disruption of consciousness is also likely to interfere with at least some aspects of (normal) focal-attentive processing." The converse, that disruption of focal-attentive processing will disturb consciousness, is also true.

What are the functions of consciousness? Velmans's answer, which on evolutionary[2] and other grounds I cannot accept, is: None. My view is more closely aligned with that of Carr (1979), Posner (1980), and Shallice (1972). Shallice (1972), for example, identified the functions of consciousness as setting and storing goals and selecting from amongst competing action patterns (usually to achieve the goal). I would add that consciousness is a representational "workshop" that supports activities such as decision making, imagination, planning, problem solving, hypothesis testing, the novel use of habitual routines (Rozin 1976), and writing this commentary. Velmans points to examples in which such complex mental functions (e.g., creative problem solving) appear to be performed either without or prior to consciousness. Even if one accepts these examples, they fail to demonstrate that consciousness does not normally assist in and usually expedite such activities. (My toaster is no less a toaster because I can toast bread in my fireplace.)

Discussions of the nature and functions of consciousness should distinguish between awareness of ideas, sensations, images, movements, actions, and so on, which I refer to as consciousness_1 and awareness of oneself – or more commonly – self-awareness, which I refer to as consciousness_2 (Puccetti & Klein 1980). It is interesting to note that consciousness_1 is a prerequisite for the development of a self-concept, for self-awareness, and for other metacognitive functions. The development and nature of our individual self-concept is critically dependent upon consciousness_1. Although there may be learning without awareness, only conscious experiences (of sensations, movements, objects, and events) seem to be stored in the form of self-referenced, episodic traces. It is primarily these traces that make up our "self" concept. Because knowledge of oneself can be extrapolated to conspecifics, an individual with the capacity for self-awareness would be able to predict the behavior of others much more accurately than an individual without such a capacity (Humphrey 1984). Thus having self-consciousness of consciousness_2 would confer tremendous selective advantages.

Can the human mind ever achieve a complete understanding of itself? The answer is: "Probably not" (Barresi 1987), but our understanding can improve. By generating discussion of the concept of consciousness, which I assume is a prerequisite for an organism like me to entertain such a question, Velmans's controversial claims will (even though they are basically wrong) help move us in the direction of such improved understanding.

NOTES

1. As Attneave (1961) so aptly pointed out, consciousness can access input information at varying levels of abstraction ranging from sensory attributes to categorical information and the experience of objects and events; similarly, conscious or voluntary control of output can be aimed at individual muscles or more abstract actions.

2. If, as Velmans claims, consciousness does not enter into or influence any human information processing, then it is awfully hard to imagine what the adaptive functions of consciousness might be that would have conferred a selective advantage leading to its evolution.

Understanding awareness at the neuronal level

Christof Koch[a] and Francis Crick[b]

[a]Computation and Neural Systems Program, California Institute of Technology, Pasadena, CA 91125; and [b]The Salk Institute, La Jolla, CA 92037

Electronic mail: *[a]koch@iago.caltech.edu or koch@caltech.bitnet*

Velmans recounts many interesting results, though we should note that some of the more challenging ones (e.g., those of

Lackner & Garrett 1973) seem not to have been studied further, whereas others, such as those concerning hypnosis, are controversial. Rather than discuss these in detail we prefer to comment briefly on the 10 points in his summary and then to describe our own approach to these difficult problems.

In brief, then, we agree with items 1, 3, 4, 5, and 6 in his summary. We do not like the second sentence of item 2. In our view what enters awareness is the result of a special process that builds on unconscious processing.

We totally disagree with item 7 since we have constructed a plausible model that does not have the problems Velmans sees. It is wise to avoid philosophical conclusions of this type until we understand how the brain works. Items 8 and 9 we feel to be largely irrelevant at this stage.

It seems strange to us that when Velmans considers the extent to which information processing involves consciousness he utterly neglects the stuff of which brains are made, that is, neurons. Cognitive science can no doubt give us important insights into this and the related issue of focal attention. A definite answer, however, can ultimately be provided only by having recourse to the neuronal level.

How could such a neuronal theory of consciousness or awareness be structured? We will give a number of plausible suggestions and argue that knowing this answer will go a long way toward understanding what processes are or are not conscious in human information processing.

We assume without further ado that awareness exists and is somehow expressed at the level of neurons and circuits. Only this position takes account of the human subjective experience of awareness on the one hand while being compatible with a reductionist stance on the other.

What form could this neuronal correlate of awareness take? Several possibilities exist. For an external (or internal) event to be perceived, one of at least three possibilities has to be satisfied:

1. A sufficiently large number of neurons somewhere in the "brain" (or, more specifically, within the "cortical" system) have to be activated. For example, any sensory event that activates more than several thousand neurons would lead to awareness of that event (Libet 1989).
2. Special "tagged" or "labeled" neurons have to be activated by this event; for example, cortical pyramidal cells in layer VI projecting to subcortical targets.
3. At least two distinct forms of neuronal firing exist, one being correlated with awareness and the other not. More specifically, we assume that the desynchronized, random firing of neurons can cause their postsynaptic target cells to discharge, leading ultimately to some motor response, but without awareness. Awareness is associated with a particular type of spatiotemporal discharge pattern. We have proposed (Crick & Koch 1990; 1991) that awareness is associated with phase-locked, oscillatory firing behavior in the 35–65-Hz frequency range as observed by Gray et al. (1989) in the visual cortex of the cat. Other possibilities, such as high-frequency bursts (Crick 1984), also exist.

These possibilities are not mutually exclusive. Thus, awareness may be associated with a sufficient number of neurons in layer VI of cortex firing in synchrony. Furthermore, we assume that the neural events associated with awareness activate short-term memory rather strongly (Crick & Koch 1990). Some alterations, however, may be produced by *un*conscious processes, as in priming.

Once we have established how "what type of firing of what types of neurons" correlates with awareness, we can immediately explain in an unambiguous manner such phenomena as blindsight (Cowey & Störig 1991), facial recognition without awareness in prosopagnosia (Tranel & Damasio 1985) and other processes that bypass awareness. For example, in our framework a visual stimulus in a blindsighted patient would cause neurons in cortical and subcortical structures to fire in a desynchronized manner, resulting in the pointing of the hand toward the stimulus. Because of the nature of the lesion, however, this stimulus would not lead to phase-locked neuronal oscillations, nor would it activate short-term memory; it would therefore not be perceived by the patient. The time for these oscillations to become established and phase-locked across cortex could, in principle, explain some of the temporal effects in "lexical priming" (Neeley 1977).

Such a theory, formulated in neuronal terms, would also clarify the function and mode of action of focal attention. A major finding of cortical electrophysiology is that while there do exist on the order of 30–40 cortical areas (per hemisphere) specializing in different aspects of visual processing, no area exists which corresponds to everything we see (Van Essen et al. 1991). The existence of such neurons, coding simultaneously for the position, orientation, depth, color, texture, and so on of objects, also appears unlikely given the resultant combinatorial explosion. Thus, at any given time, a perceived object has to be represented by the firing of a set of neurons all over the cortex. For example, if I am looking at my daughter talking to me, neurons in area MT code for the motion of her face, neurons in area V4 for the hue of her face, neurons in primary and higher auditory cortices code for her voice, and neurons in the inferior temporal lobe may code for the visual template or icon associated with my visual concept of my daughter. In what manner is the firing activity of this highly dispersed group of neurons, however, combined to lead to the unitary subjective concept of my daughter speaking to me? And, furthermore, how is this neuronal activity and the associated percept distinguished from the activity of neurons responding to the face of my son in the background? This problem, faced by any highly distributed connectionist system, is known as the *binding problem.*

One needs to distinguish at least three types of binding.

1. A simple cell in visual cortex has its preferred orientation always perpendicular to its preferred direction of motion. Thus, orientation and direction of motion are combined. Neurons throughout cortex seem to compound a number of variables in this manner.
2. A second type of binding is probably acquired by overlearning. Thus, ecologically important and frequently viewed stimuli, such as grandmothers or letters and words, may be represented by the firing of small groups of neurons.
3. Because the capacity of both types of binding is limited, however, we need to postulate a third, very rapid and transient type of binding mechanism with essentially infinite capacity. It is likely that focal attention instantiates this type of binding mechanism.

Specifically, following von der Malsburg and Schneider (1986), we suggest that this binding is achieved by all neurons that are associated with the perceived object firing in synchrony. Thus, phase-locked oscillations *are* the cellular expression of attention. The time to set up this synchronized timing activity could – at least in principle – explain phenomena such as illusory conjunction (Treisman & Schmidt 1982) and the well-known limitation of primary or short-term memory (Crick & Koch 1991). Furthermore, such a theory would also explain why certain processes can be carried out in parallel, without requiring awareness. We simply assume that repeated synchronous activation of groups of neurons responding to a particular stimulus (e.g., a letter or a face) cause individual neurons or small groups of neurons to become directly activated in response to this stimulus, on the basis of some Hebbian rule. Such a learning mechanism would eventually bypass the need for a focal attention mechanism for the perception of such frequently viewed stimuli (i.e., transforming a type 3 binding problem into a type 2).

Such a neurobiological theory of awareness and consciousness will also help us understand to what extent (and how) anesthetic agents render us really "unaware" or "unconscious" (Kulli & Koch 1991) and will enable us to manipulate the brain in novel ways.

The more general point, however, is that only a theory of awareness formulated in neuronal language will ultimately resolve the issues addressed by Velmans in a clear and unambiguous manner. We do not wish to slight the accomplishments of cognitive science in unraveling the mysteries of the mind. Until the advent of routine recordings in the brains of cats and primates, psychology was the main source of knowledge about such processes. It seems to us, however, that the time is ripe for a neurobiological approach to these problems, relating awareness and consciousness to neurons and their firing patterns, in other words, for explaining the mind on the basis of the brain.

Conscious functions and brain processes

Benjamin Libet

Department of Physiology, School of Medicine, University of California, San Francisco, CA 94143-0444

Velmans skillfully and clearly reviews the evidence that consciousness (i.e., introspective awareness) is not necessary to human information processing of all kinds, including that involved in the control of motor action. I can agree with much of that; we have ourselves produced direct experimental evidence for that proposition (Libet 1965; 1985; 1987; 1989; Libet et al. 1991). But Velmans generalizes the argument so as to conclude that no human information processing involves consciousness entering into or causally influencing the process (sect. 9.1). I argue that that more general conclusion goes beyond the evidence and is unwarranted, except as an item of belief.

In another analysis of the causal status of the conscious mind, Velmans introduces the concept of complementarity between the first-person and third-person perspective. I agree fully that the two perspectives are complementary and mutually irreducible, in the sense that neither is describable in any a priori fashion by the other (Libet 1965; 1966; 1981; 1985; 1987; 1989). However, Velmans's extension of that concept to deny that consciousness interacts with the brain and his analysis of the significance of causality as viewed from each perspective is not convincing or perhaps even tenable. I elaborate my position on both of these general issues below.

1. Does consciousness enter into input analysis and choice?

1.1. Neural time factor in identification of and response to signal. Experimental evidence has indicated that a substantial period of cortical activity, up to about 500 msec or more, is necessary for the production and appearance of a conscious sensory experience (Libet 1966; 1981; 1989; Libet et al. 1979). As we had pointed out, the ability to react meaningfully to a stimulus within as little as 100 msec clearly suggests that detection, processing, and organized decision-making to respond can all take place unconsciously, before awareness of the signal has developed. A direct confirmation of this was recently carried out by Taylor and McCloskey (1990) and was actually already indicated by previous reports (Fehrer & Biederman 1962). These investigators showed that when a visual signal was followed by a second visual stimulus, which retroactively blanked out any reportable awareness of the initial signal, the reaction-time to that first signal was the same as that when the signal was given alone (with reportable awareness).

In recent work (Libet et al. 1991) we have shown directly that a stimulus to the somatosensory thalamus of human subjects can be detected correctly in a forced choice response even when the subjects are completely unaware of any stimulus-induced sensation. This occurred predominantly with stimulus durations of about 250 msec or less. To produce even an uncertain and ambivalent report of sensory awareness (with correct detection) required a substantially larger stimulus duration (an *increase* of >250 msec) than did correct detection with no awareness. Of additional interest is the fact that the subjects were intensely *focusing their attention* on the 1-sec intervals during which they were to receive a possible stimulus; yet they did not achieve conscious awareness for many of the stimuli that were nevertheless correctly detected in their forced choice response. This demonstrates that cerebral information processing can be dissociated from awareness of that processing, and that focal attention and cognitive responses can be separated from the conscious feature of a signal, in accord with Velmans's contention (sect. 5.1, 8).

That "It is important to distinguish the contents of consciousness . . . from the processes which encode information" (sect. 4.2) had indeed been directly and experimentally demonstrated by Libet et al. (1979). In this work, crucial tests confirmed that the subjective conscious time for a sensory event was antedated (referred backwards in time) by up to about 500 msec relative to the time for neural production of the sensory experience.

1.2. Conscious voluntary choice to act: conscious causality. Velmans (sect. 3) refers to our experimental finding that the initiation of a voluntary act appears to develop unconsciously in the brain, about 350 msec before the subject is consciously aware of the urge or intention to act (Libet 1983; Libet et al. 1983). This indicted that conscious volition may not be necessary to the initiation of a decision or choice of when to act; to that extent it adds evidence for the general argument that the neural processes for a volitional choice, as for recognition of a sensory stimulus, can be dissociated from the conscious function related to these events.

The experimental evidence, however, does not justify the broader conclusion that "no human information processing is conscious in (the) sense . . . that consciousness enters into or causally influences the process" (sect. 9.1). If such a conclusion were valid it would indeed have far-reaching theoretical implications, as would the opposite conclusion. But that conclusion would be faulty on at least two grounds.

First, in the case of voluntary action one should distinguish between the *initiation* of the volitional process and the final *outcome* of it in the actual performance of the act (Libet 1985). Although conscious intention appears well after the cerebral process starts, it does precede the actual motor act by about 150 msec. In that final period after the subject has became aware of the intention to act, the subject could consciously control the outcome of vetoing the intention and not moving at all, or by passively or actively promoting its completion. There is no experimental evidence that would deny a causal role for a conscious control function here, although admittedly there is none to demonstrate such a role either. The alternative possibilities remain viable at the level of philosophical outlook, with neither having been experimentally excluded.

Second, Velmans may be making a leap that is unwarranted in principle, from the argument that consciousness is *not necessary* for the development of many human information processes to one that consciousness therefore can *have no role* in affecting these processes. A condition can be a sufficient factor without being a necessary one. Furthermore, it may be premature to generalize that consciousness is even not necessary for any human information processing, although it could be an acceptable working hypothesis. It should be noted that the argument about an active role for consciousness can be independent of that about monistic or dualistic theories of the mind-brain relationship (e.g., Libet 1985).

2. Complementarity for first-person versus third-person perspectives. I take it that the terms "first-person perspective" and "third-person perspective" are equivalent, respectively, to the