A Tri-Level Analysis of Consciousness

Richard Banister
CGS 5900 Cognitive Science Seminar
23 April 2008

Introduction

Despite its old age, the study of consciousness is still a very young science. After being derailed by behaviorism early in the 20th century, it has been guided by philosophers and as such has made little progress towards explanation. In fact, in the debate over consciousness, philosophers seldom get past the first step – deciding whether or not there actually should be a debate (Dennett 1988). It follows that in the past century there has not been a proper philosophical resolution even of what consciousness *is*. Yet the notion of consciousness is of such personal relevance to every human that it is unlikely to shrug off the shackles of scholarly debate any time soon. What then are our prospects for actually coming to a reasonable understanding of consciousness? Luckily, in the past two decades, while philosophers have struggled with high-level analyses and criticisms, many psychologists and neuroscientists have managed to sneak away from the scuffle and begin serious work in the scientific study of consciousness. It now seems reasonable that this bizarre phenomenon that we all know but cannot explain may be well on the way to scientific understanding. Although consciousness appears to be a mysterious intangible, it can be precisely defined and explained through a multi-level approach that employs the strengths of many disciplines.

Such an explanation will likely not settle the issue for philosophers of mind, however, nor is finding this explanation an easy proposition. Consciousness is somewhat unique in the challenges it faces on the path to explanation. Nothing defines our lives so much as the conscious experiences we undergo every day; thus everyone has a very personal stake in the matter. More than any other aspect of the mind it is attacked by enemies of functionalism who assert that a cognitive approach will not suffice in defining personal experience (Jackson 1982). How could the notion of *what it is like to be something* ever be translated into a schematic, even if it is detailed down to the neural level? Furthermore, cognitive approaches have even been criticized by scientists in other fields, including physicists who speculate consciousness must emerge out of quantum interactions (notably Penrose 1989; for a counter argument see Koch and Hepp 2006).

Unfortunately the intensity of the debate has prevented the development of a complete cognitive theory of consciousness to date. Moreover, it is not easy to categorize extant research in classical cognitive science terms. Just as some convenient high-level theories may be compromised by lack of evidence in lower levels (Block 1995), implementation-level accounts may lack the focus that computational and algorithmic accounts would grant. An amalgamation of approaches would foster understanding, as Crick and Koch argued:

> The most effective way to approach the problem of consciousness would be to use the descriptions of psychologists and cognitive scientists and attempt to map different aspects of the models onto what is known about the neuroanatomy and nerophysiology of the brain. (Crick & Koch 1990)

But like axons of one neural group reaching for dendrites of another, the connections between the three levels are not yet fully established. One of the purposes of this paper is to show where such connections exist and suggest where others might be made.

Computational

In tackling the daunting topic of consciousness we will need a place to begin, and Crick & Koch nicely give this guideline: “We make two basic assumptions. The first is that there is something that requires a scientific explanation. [...] The second assumption is tentative: that all the different aspects of consciousness ... employ a basic common mechanism or perhaps a few such mechanisms” (Crick & Koch 1990). While we will leave the second assumption aside for the moment, the first assumption is important in a computational explanation. It allows us to set aside most philosophy and focus on real constraints we can use to define our problem. We cannot consider every aspect of consciousness in the limited scope of this paper, but we can at least get an idea of it as a cognitive process.

The first problem before us lies in simply defining what is meant by the term “consciousness.” One way to approach this is from an evolutionary standpoint: What are the faculties humans have which other animals lack? What function do they, and consciousness in general, serve? Some cognitive systems, e.g. vision, may have fairly obvious advantages, but what adaptive use is there in consciousness? Why has it been a naturally selected trait? Donald (1995) has taken just such an evolutionary approach to distinguishing features of consciousness. Donald shows a general order of evolution for three aspects of consciousness: *state*, *attentional awareness*, and *representational awareness* (Donald 1995). *State* refers simply to conscious state (e.g. awake, asleep, altered) and likely evolved first, as it is the most simple and fundamental aspect and can be seen in many mammals. *Attentional awareness* is a control system that governs conscious behavior; it is complex but appears in other advanced species, thus it must have emerged later. And *representation* was likely the last to arrive, as it is the most difficult to show in other animals (language, for example, appears to be a uniquely human representation system) (Donald 1995). Donald concludes, "the special nature of human consciousness is defined primarily by our special representational skills, and that there are also crucial architectural features that distinguish human consciousness from that of our closest mammalian relatives" (Donald 1995).

One of these features is the human ability to cue and recall memories at will, which animals apparently lack. Donald theorizes that "autocued nonverbal recall in humans ultimately derives from an expansion of primate imitative and manipulative skills" (Donald 1995). Thus an attentional awareness architecture must allow for voluntary recall, and representation is likely tied to sensorimotor systems. There is likely some evolved physical difference that allows humans and not primates to autocue memories (presumably testable in neuroanatomy). In this view,

> The recoding of knowledge for explicit retrieval seems to involve symbolic or quasi-symbolic processing, and produces both a path for explicit memory access ... as well as a reformatting of that knowledge. This recoding process can feed on its own outputs, iteratively, so as to construct increasingly larger and more abstract representations. (Donald 1995)

Donald's computational speculation has clear implications on the algorithmic level. We will explore some *attentional awareness architectures* that may satisfy these constraints when we look at the algorithmic level of consciousness.

Seth, Baars and Edelman (2004) approached the question from the other side, developing a list of criteria for determining consciousness in animals (notably more charitable than Donald). While they classify accurate reporting as the primary indicator of a conscious phenomenon in humans, they give other observations which may be useful to us on a computational level. In particular, they present several important constraints under which consciousness operates: an *informativeness* preference, whereby attention tends to go to stimuli that contain more information; the *fleeting* nature of consciousness; the *consistency constraint* and the *serial constraint*; the *focus-fringe* nature, in which many scenes exist just outside of consciousness; and *facilitation of learning* (Seth, Baars & Edelman 2004). Some of these points merit further discussion here.

Seth et al. note that conscious scenes last less than thirty seconds, and our immediate sensory experience is even shorter. As a result, consciousness can move rapidly from one scene to the next, quite the opposite of long-term memory. They speculate an evolutionary explanation: "for conscious scenes to have adaptive value for an organism, they must have a short lifetime – enough time to recruit a broad network of neural resources to generate appropriate behavior, yet also a tendency to evolve into subsequent scenes" (Seth et al., 2004). Consciousness must have adapted to deal with the fleeting present.

Connected with this are the ideas of *consistency* and *serial* constraints. Consistency indicates that if multiple, conflicting interpretations of a stimulus are possible (for example, in the case of words with more than one meaning), only one interpretation can be active in consciousness at a time. The *seriality* constraint specifies that consciousness can only be devoted to one scene at a time (Seth et al., 2004). So in total, so far we have seen that consciousness has a limited focus and attention can only be given to a fairly rapid stream of individual scenes, although scenes can wait on the fringes of consciousness. This description already can provide the beginnings of an architectural description of consciousness.

Furthermore, consciousness appears to have a connection with memory and learning: "there is very little evidence for long-term learning of unconscious input. In contrast, the evidence of learning of conscious episodes is overwhelming" (Seth et al., 2004). This means there must be a connection with learning structures as well. It seems consciousness is well-connected throughout the brain. Based on the extensive capabilities of consciousness, Seth et al. assert "an integrative concept of consciousness must therefore involve many brain regions as well as the interactions among them, along with the ability to recruit regions such as hippocampus ... and cerebellum" (Seth et al. 2004).

This seems remarkably similar to Donald's attentional awareness architecture. In fact, there is general consensus that consciousness utilizes many unconscious subsystems. Dehaene (1998), for example, points out this strange, ironic nature of consciousness:

> Mental effort is clearly unrelated to objective measures of computational difficulty: We routinely perform vision and motor control tasks without awareness of the complex underlying information processing, whereas elementary tasks, such as subtracting 9 from 37, call for our attention and conscious effort. (Dehaene 1998).

While most automatic functions are relegated to cerebral modules, according to Dehaene, humans have the additional ability to select and connect these modules to perform more deliberate actions (Dehaene 1998). In order to perform conscious subtraction (which involves representations and is more computationally complex than we might think) we need to do just such a recruitment.

Sergent and Dehaene (2004) attempted to clarify the border between conscious and unconscious thinking. Do scenes come gradually into consciousness, or is it a sharp, all-or-nothing affair? To resolve the issue, Sergent and Dehaene employed the attention blink (AB). AB is a phenomenon in which perception of one stimulus can interfere with the ability to perceive another if it is presented 200-500ms after the first. The goal was to determine if the AB somehow deteriorates information available to consciousness, or cuts out awareness completely. In the case of degradation, there should be a gradual curve of awareness over the manipulated length of AB, with least awareness midway at 300ms; an all-or-nothing result graph would be discontinuous at the AB (Sergent & Dehaene 2004).

The participants were asked to evaluate subjective visibility of a number word shortly after its presentation. Participants were shown a succession of frames: random four-letter non-word distracters presented for 43ms followed by a blank for 43ms. The target stimulus T1, either "XOOX" or "OXXO," was shown after either 7 or 10 frames, followed by distracters for 1-8 frames (86ms-688ms), then target stimulus T2, a number word. Some trials presented no T2. Participants were asked about T2 215ms after its presentation. The question asked the subjects to rate the visibility of T2 as "finely as possible" on a continuous scale ranging from "not seen" to "completely visible," under no time constraint. Two types of trials were used: in single-task trials, subjects asked to simply rate T2 visibility; dual-task asked about the central letters of T1 as well (Sergent & Dehaene 2004).

The results show a discrete nature of awareness of T2: “although the subjective visibility scale was designed to be sensitive to continuous changes in perception, participants used it in an all-or-none fashion. The AB did not result in a gradual reduction of T2 visibility, but rather resulted in an increase in the proportion of trials on which T2 was missed and had the same visibility as on target-absent trials” (Sergent & Dehaene 2004). There is likely a threshold for activating consciousness; any amount of stimulus below this level remains in the unconscious realm. Sergent and Dehaene (2004) see this as evidence corroborating a global workspace model of consciousness, which we shall see at the algorithmic level.

Algorithmic

Much work has been done to try to define and isolate the components of the mind that contribute to consciousness. Although these functional decompositions are still very young, some remarkable progress has been made. We will first look at the issue of awareness, and the possibility of defining an *awareness architecture* like the one specified by Donald (1995) earlier. Visual awareness is a good place to start, as the visual pathways are cognitively relatively well understood. Crick and Koch (1995) presented evidence that primates are not directly aware of activity in V1, the primary visual cortex, and extrapolated that humans likely share this property. Although they did not have a good idea of how the awareness system itself worked, this study helped narrow down what modules are involved by eliminating V1. Although this approach is from the neuronal level, it has interesting implications on the algorithmic level.

We know from studies of vision that V1 runs in serial from the lateral geniculate nucleus (LGN) on to higher visual cortical areas. Macaque monkey brains lack a direct connection from V1 to systems in the frontal cortex, which is critical for awareness. It has no connection with the intralaminar nuclei (ILM), which are thought to be critical for consciousness (Bogen 1995) as we will see at the implementation level. V1 instead connects to the superior colliculus (role in eye movement), LGN, and pulvinar. The path through colliculus to pulvinar may be important in visual attention, but is “not sufficiently direct or strong to produce, by itself, vivid visual awareness of the neural activities in V1" (Crick & Koch 1995). The lack of connections to the frontal cortex means V1 is unlikely to act in consciousness. According to Crick and Koch, "unless a visual area has a direct projection to at least one [frontal area], the activities in that particular visual area will not enter visual awareness directly, because the activity of frontal areas is needed to allow a person to report consciousness" (Crick & Koch 1995).

This theory is corroborated by evidence of the blindsight phenomenon in macaque monkeys shown by Cowey & Stoerig (1995). Blindsight is characterized as better-than-random identification of presented objects, despite phenomenal blindness, when forced to guess. Their study shows that "monkeys with lesions of the primary visual cortex can learn to detect, localize, and distinguish between visual stimuli presented within their visual file defects" (Cowey & Stoerig 1995). This is equivalent behavior to blindsighted humans. The subjects of the study were one monkey with normal vision and three with surgical removal of V1 in the left cerebral hemisphere and thus affected vision in the right hemifield. This surgical procedure was performed several years before the experiment, so the monkeys had time to adjust to their condition. In the first experiment, the monkey fixated on the center of a screen, then stimulus was presented in one of four locations: two in left hemifield, two in right. If monkey touched the location of the stimulus, he/she was rewarded with food. The results showed that monkeys with lesioned V1 responded very effectively to stimulus in the affected hemifield, although less effectively than in the unaffected hemifield (Cowey & Stoerig 1995).

To rule out the possibility that the monkeys were experiencing simply degraded vision instead of true blindsight, Cowey & Stoerig (1995) performed a second experiment. This experiment introduced blank trials which required a different response than stimulus trials. Stimulus trials showed a flash of light on the left side, which the monkeys would have to press for reward. On probe trials the stimulus appeared in the right hemifield. Half of the trials were blank trials, in which the monkey had to press a permanent blank square for reward. The normally-sighted monkey performed all tasks at near perfect accuracy. Affected monkeys performed well, although not as well as their counterpart, on left hemifield stimulus and blank trials, but were unable to detect stimulus in the right hemifield (Cowey & Stoerig 1995). These results suggest that the monkeys were experiencing blindsight and not just poor vision in one field. The study shows that the primary visual cortex is likely not involved in visual awareness in monkeys, which have human-like visual systems. Similar methods could hopefully be used to determine which brain components are involved in consciousness and which are simply peripheral to the system.

The modularity of consciousness in general is argued for by Spinazzola, Pia, Folegatti, Marchetti and Berti (2008), who presented possible double dissociations for awareness of some disorders. This shows “potentially different brain localizations for motor and sensory monitoring processes” (Spinazzola, Pia, Folegatti, Marchetti & Berti 2008). The authors point out studies of patients who show anosognosia – unawareness of a disability – for hemiplegia (paralysis of half of the body) and hemianaesthesia (loss of sensation in half the body). A double dissociation between anosognosia for hemiplegia and anosognosia for hemianaesthesia would show that the sensory and motor monitoring systems are localized differently and thus would support modularization.

In fact the data does support this hypothesis. Spinazzola et al. studied four patients with both disorders, and showed dissociations between the two forms. An example is the patient PR, who suffered a stroke which caused total left hemiplegia and severe hemianaesthesia, manifested in tactile deficiency. PR was perfectly aware of the hemiplegia, but was convinced he performed well on sensory tests (Spinazzola et al., 2008). This shows a dissociation: damage to awareness of motor systems does not necessarily affect awareness of sensory systems. Spinazzola et al. also point to the rare cases of anosognosia for hemianaesthesia without anosognosia for hemiplegia, which makes a double dissociation indicating that monitoring processes for sensory and motor systems are located in different areas. These areas “may be involved in execution of primary functions and the emergence of awareness related to the monitoring of the same functions” (Spinazzola et al., 2008).

Modularity has important implications for consciousness, as Spinazzola et al. point out: “The ‘illusion’ of unity of the self, assumed by common-sense theories of consciousness, is evident in the normal experience of correspondence between the actual presence/absence of a stimulus and the presence/absence of a subjective experience of it” (Spinazzola 2008). However, consciousness is not necessarily a single centralized structure – it may instead be the sum of workings of different components distributed across the brain. This means that damage to certain areas would cause impairments in their related functions instead of wholesale damage of consciousness (Spinazzola et al., 2008). These results certainly support the idea of modularity of consciousness.

We will now look at some relatively complete algorithmic models which propose connections of modules to form consciousness. Shallice (1988) provided a good early sense of the flow of consciousness. The principle of Shallice's system is to distribute the conscious control system over several modules. According to Shallice, "any mental operation which gives rise to awareness necessarily involves a number of different subsystems" (Shallice 1988). A basic diagram of Shallice's system is given in figure 1.

Shallice's model contains two primary control systems – contention scheduling and the supervisory system – and two subsidiary ones: the language system and episodic memory process. Shallice describes their interactions in the following manner:

> Quite often, as with the selection of a very overlearned schema in contention scheduling, an operation in one subsystem may have no counterparts in the other subsystems. However, very frequently when one of these subsystems is operating its operation would be complementary to the operation of on or more of the others ... In these situations, there would be a coherent pattern of control over all *other* subsystems which is shared between those control subsystems that are active." (Shallice 1988)

It is the pattern of control then that is important in this model. Contents of consciousness "would correspond to the information flow between the control subsystems and, in turn, to the flow of information and control between them and all the rest of the cognitive system" (Shallice 1988). Note that this corresponds nicely with Spinazzola et

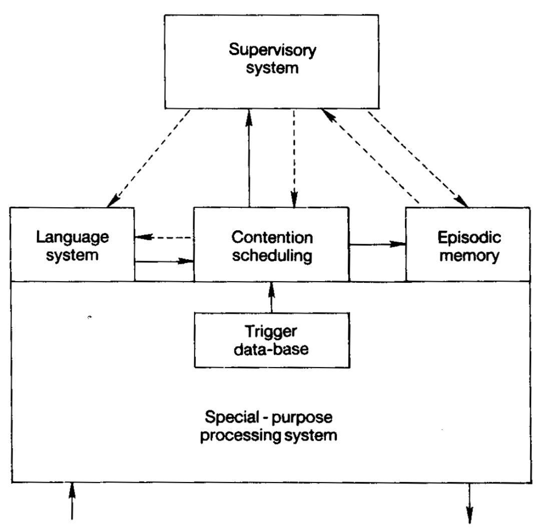

Figure 1. Shallice's model of consciousness. Reprinted from Shallice (1988).

al.'s idea that consciousness is not necessarily one central control system but rather the sum of several different modules.

Global Workspace theory (summarized in Baars & Franklin 2003) is another model of consciousness that attempts to connect disparate unconscious subsystems. It classifies the brain as a set of interconnected modules for different purposes. There is a shared global workspace of short memory, the contents of which are broadcast to the dispersed modules. GW theory also integrates perceptual networks as well as unconscious networks ("contexts"), which can affect consciousness. GW theory views emotions as goal contexts, and overarching executive functions as hierarchies of goal contexts (Baars & Franklin 2003).

Baars and Franklin also theorize that conscious events influence working memory, based on the close association of consciousness with WM: "All active components of classical working memory are conscious: input, rehearsal, visuospatial operations, recall and report," as evidenced by accurate reporting of these components (Baars & Franklin 2003). This sits well with the computational guidelines given by Seth et al. (2004) and may work with Donald (1995) as well.

Crick and Koch (1990), working at a neuronal level, also support the idea of an attentional model tied to working memory, using the visual system again as a guide. Their research has in turn been influenced by such accounts. In particular, "[cognitive models] point to the importance of attention and short-term memory and suggest that consciousness should have easy access to the higher, planning levels of the motor system" (Crick & Koch 1990). Crick and Koch outline two types of visual awareness: a fast one, tied with very short iconic memory, and a slower one which is linked with visual attention and short-term memory. Long-term memory seems unimportant to conscious experience. Crick and Koch proposed that awareness is a serial process sitting on top of the massively parallel process of brain computation. The serial attention mechanism "binds together all those neurons whose activity relates to the relevant features of a single visual object. We suggest this is done by generating coherent semi-synchronous oscillations ... these oscillations then activate a transient short-term (working) memory" (Crick & Koch 1990). They go on to suggest that a kind of unity (perhaps the kind that would satisfy Spinazzola) is created across disparate parts of the brain by sets of neurons oscillating at certain frequencies (Crick & Koch 1990). Although this idea is challenged by Bogen (1995), it is the beginning of an implementational account of the algorithmically specified attentional mechanism.

Implementation

We will see two distinct approaches in the implementation level. First there is an attempt to explain a neural correlate of consciousness (Crick & Koch's term), discovering those parts of the brain which are important for the emergence of consciousness. Second, there have been attempts to create computer simulations of some algorithmic models of consciousness.

Bogen (1995) asserts that activity in the intralaminar nucleus (ILN), a part of the thalamus, gives rise to conscious awareness. Although the thalamus and cortex are both thought to be important for consciousness, Bogen points out that bilateral lesions of ILN cause loss of conscious state, while lesions in cortex deprive consciousness of certain contents but consciousness itself is retained (Bogen 1995). Bogen concludes that ILN is the most likely structure to form the seat of consciousness. Bogen shows that ILN is well-connected to other brain regions, indicating that it is probably related to consciousness in several ways. There is probably a role in primitive perception that does not extend to the cortex. ILN probably contributes to awareness of activity in the cortex, due to its extensive connections from the cortex. Additionally, ILN is heavily connected to the striatum, indicating a likely role in volition, the neocortex, which is important for ideation, and connections through the reticular nucleus may give it a role in attention-selection (Bogen 1995).

This certainly makes ILN a strong candidate for an implementation of consciousness, as it satisfies many of the criteria for defining consciousness which we considered on the computational and algorithmic level. It is well connected to different subsystems, and possibly has awareness of what cortex is doing. Additionally, Bogen notes "the existence of connections to ILN from globus pallidus ... suggests a monitoring of motor systems, as do the cortical projection to ILN from sensorimotor and premotor cortex" (Bogen 1995). Of course, this does not end the quest for a neural correlate of consciousness (it is not likely that ILN *is* consciousness!). But it does begin it.

We will now see an approach to modeling a theory of consciousness in a computer system. Dehaene, Kerszberg, and Changeux created a connectionist model that was used to perform effortful tasks – in other words, tasks in which some conscious control is needed to produce correct results. Dehaene et al. proposed approaching effortful tasks from a GW perspective, noting that a global workspace can "potentially interconnect multiple distributed brain areas in a coordinated, though variable manner, and whose intense mobilization might be associated with a subjective feeling of conscious effort" (Dehaene et al., 1998). In GW, perceptual systems, long-term memory, evaluative systems, and attentional systems feed into a shared global workspace, which is also connected with the motor systems. Dehaene et al. implemented a computer model of this system in order to test if GW can properly handle effortful tasks. While the details of the model are outside the scope of this paper, it basically consisted of input modules that could accept words and colors and a module that output color names. The modules were limitedly connected each other and also to a global workspace of malleable processing units. Using the Stroop task, in which color names must be provided based on (sometimes conflicting) input stimulus, the researchers gave the computer model three tasks. Two were routine: first, given a color, provide the color name and second, provide the color name given a color word with incompatible color interference. There was one effortful task, which was to produce a color name given a color and conflicting word interference (Dehaene et al., 1998). For example, a subject in the Stroop test would be presented with the word "blue" written in green ink and asked to say the name of the color. The conflicting stimulus invites mistakes and thus involves more conscious processing.

For the first routine task, a simple association, the simulation produced correct results using only connections from color units to name units. There was, as would be expected, no activity in the workspace. The network performed the second task with ease as well, avoiding activating the workspace at all. Dehaene explains that "word-to-name connections are stronger than color-to-name connections. Hence the naming response appropriate to the word is activated faster and more strongly than the one appropriate to the color" (Dehaene et al., 1998). The effortful task involved significantly more processing. The first trials were met with errors as the network tried to perform routine task two and failed to produce correct results. Negative reward from these errors then caused network vigilance to go up, and thus push activity into the global workspace. Through several more trials the network discovers activation patterns through the GW that produce correct results. After this discovery period, the network is still highly vigilant for a time and so even routine tasks tend to go through GW. Notably, the network learns to "routinize" the effortful task, and eventually GW activation drops out completely (Dehaene et al., 1998).

Although Dehaene et al. admit this network is incomplete, its activities may be instructive to us. How would we perform in the same task? When presented with the conflicting stimuli we would likely make mistakes at first, then process the problem more closely and produce the correct result. Afterwards we would probably be in "processing mode" and overwork the next few stimuli, even if they are just simple associations. Training with the effortful task would likely make it routine. This study is notable in effectively modeling human behavior while producing a convenient implementation of the global workspace theory. Whether this theory is implemented in human neuroanatomy is another question.

Conclusion

Consciousness is clearly not a fully-explained phenomenon, but the research presented here should give the impression that we are on our way to an understanding. Within the past twenty years significant research has been done and data is being used to reveal pieces of the puzzle. Whether we will ever see the complete picture is a matter which has been debated by philosophers for a long time; unfortunately such debate masks progress in the field. Already we have working models of consciousness which describe it as a serial control and awareness system that recruits and connects unconscious modules to perform high level tasks; it works with short-term and very short-term memory, and is also tied to learning. We know likely places to find it in the neuroanatomy, and we have the beginnings of computer implementations. We also have a sense of the purpose of consciousness to humans, even if every human defines the word differently. With all this knowledge of consciousness, it is clearly premature to write it off as unknowable. Whether we can fully understand consciousness remains to be seen, but we have the methodology available to try.

References

Baars, B. J. (2005). Subjective experience is probably not limited to humans: The evidence from neurobiology and behavior. *Consciousness and Cognition, 14,* 7-21.

Baars, B. J. & Franklin, S. (2003). How conscious experience and working memory interact. *TRENDS in Cognitive Sciences*, Vol.7 No.4. 166-172

Bogen, J. E. (1995). On the neurophysiology of consciousness: I. An overview. *Consciousness and Cognition, 4*(1), 52-62.

Cowey, A. & Stoerig, P. (1995). Blindsight in monkeys. *Nature,* Vol. 373, 19. 247-249.

Crick, F. & Koch, C. (1990). Towards a neurobiological theory of consciousness. *Seminars in the neurosciences*, Vol. 2. 263-275.

Crick, F. & Koch, C. (1995). Are we aware of neural activity in primary visual cortex? *Nature*, Vol. 375, 11. 121-123.

Crick, F. & Koch, C. (1998). Consciousness and neuroscience. *Cerebral Cortex*, Vol. 8, No. 2. 97-107.

Dehaene, S., Kerszberg, M., & Changeux, J. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences (USA)*, Vol. 95. 14529-14534.

Dennett, D. C. (1988). Quining qualia. In A. Marcel and E. Bisiach (Ed.), *Consciousness in contemporary science* (pp. 42-77). Oxford: Oxford University Press.

Donald, M. (1995). The neurobiology of human consciousness: An evolutionary approach. *Neuropsychologia,* Vol. 33, No. 9. 1087-1102.

Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32. 127-136.

Koch, C. & Hepp, K. (2006). Quantum mechanics in the brain. *Nature*, Vol. 440, No. 30. 611-612.

Penrose, R. (1989). *The emperor's new mind*. Oxford: Oxford University Press.

Sahraie, A., Weiskrantz, L., Barbur, J. L., Simmons, A., Williams, S. C. R., & Brammer, M.J. (1997). Pattern of neuronal activity associated with conscious and unconscious processing of visual signals. *Proceedings of the National Academy of Sciences (USA)*, Vol. 94. 9406-9411.

Sergent, C. & Dehaene, S. (2004). Is consciousness a gradual phenomenon? Evidence for an all-or-none bifurcation during the attentional blink. *Psychological Science*, Vol. 15, No. 11. 720-728.

Seth, A. K., Baars, B.J., & Edelman, D.B. (2004). Criteria for Consciousness in humans and other mammals. *Consciousness and Cognition,* 14. 119-139.

Shallice, T. (1988). Information-processing models of consciousness. In A. Marcel and E. Bisiach (Ed.), *Consciousness in contemporary science* (pp. 42-77). Oxford: Oxford University Press.

Spinazzola, L., Pia, L., Folegatti, A., Marchetti, C., & Berti, A. (2008). Modular structure of awareness for sensorimotor disorders: Evidence from anosognosia for hemiplegia and anosognosia for hemianaesthesia. *Neuropsychologia*, 46. 915-926.