The Search for Neural Correlates of Consciousness

Jakob Hohwy

[Published as: Hohwy, J. The Search for Neural Correlates of Consciousness. *Philosophy Compass* 2/3: 461-474, 2007.]

Most consciousness researchers, almost no matter what their views of the metaphysics of consciousness, can agree that the first step in a science of consciousness is the search for the neural correlate of consciousness (the NCC). The reason for this agreement is that the notion of 'correlation' doesn't by itself commit one to any particular metaphysical view about the relation between (neural) matter and consciousness. For example, some might treat the correlates as causally related, while others might view the correlation as evidence for identity between conscious states and brain states. The common ground therefore seems to be that the scientific search for the NCC is largely independent of the metaphysics of consciousness.

Though contributors to this debate generally set aside metaphysical issues, they do acknowledge that one's *conception* of consciousness is important to the search for the NCC. One will not know what one is looking for, or how to look for it, if one's concept of consciousness is confused or otherwise unclear. Researchers with different concepts of consciousness are likely to arrive at quite different NCCs, or to disagree about the prospects of finding an NCC.

In section 1, I briefly introduce the basic idea of the NCC, together with some paradigm NCC-studies in the science of consciousness. In Sections 2 and 3, I review the two above issues, concerning the conception of consciousness and the metaphysics of consciousness.

1. Searching for the neural correlate of consciousness

One of the main researchers in the area, Christof Koch, defines the NCC as "the minimal set of neuronal events and mechanisms jointly sufficient for a specific conscious percept" (Koch 2004: 16). This means that to find the NCC one must isolate the neuronal activity that is specific to the conscious representation of a particular content (such as the experience of a face) (Chalmers 2000). It is important for researchers to identify the 'minimal' set of sufficient conditions because much of the total set of neural activity that is sufficient for a given conscious state is viewed as irrelevant for the occurrence of that state. Therefore, researchers aim to whittle away at the sufficient neural conditions in order to find the minimal set that is most directly relevant for a particular conscious content. For example, activation of the cells in the retina is part of the sufficient conditions for seeing a face, but as the phenomena of binocular rivalry (described below) demonstrates, it doesn't seem to be what triggers the conscious experience of a face.

To identify neuronal activity researchers use a variety of neuroscientific measurement techniques. These include, for example, single cell measurements with microelectrodes (mostly in monkeys) and functional magnetic resonance imaging (fMRI). However, it is difficult to devise experimental paradigms that focus exclusively on neural activity specific to consciousness. When one sees a face there will be much activity (for example, on the retina and in the early visual cortex) that seems explanatorily redundant for the formation of the conscious percept of a face.

Researchers therefore aim to keep as many things as possible constant, including the stimuli, while varying the conscious percept. One widely used design uses binocular rivalry where different stimuli are shown to each eye (e.g., a face to one eye and a house to the other); the subjective experience is not of some kind of blend of the two images but rather alternates between the two images every few seconds. (See demonstrations of this at http://www.psy.vanderbilt.edu/faculty/blake/Rivalry/BR.html, or try it at home by looking through a rolled-up piece of paper with your right eye and at the back of your hand with the left.) (Alais and Blake 2005) Using this design, changes in neural activation will reflect changes in the conscious percept rather than changes in the stimuli (Frith, Perry et al. 1999; Leopold and Logothetis 1999; Koch 2004). One study found that activation in the fusiform "face area" (FFA) correlates with face perception, and activity in parahippocampal "place area" (PPA) with house perception (Tong, Nakayama et al. 1998).

So far, many studies using binocular rivalry (plus studies with many other experimental designs) have been published (for reviews, see Frith, Perry et al. 1999; Leopold and Logothetis 1999; Kanwisher 2001; Koch 2004). However, it has proven very difficult to integrate the experimental findings from these studies with each other. Single cell studies tend to show less activity correlated with consciousness in early visual areas (V1) and more in higher (temporal and prefrontal) areas; fMRI studies have found correlated activity in high (right prefrontal) areas but also in low visual areas, even areas as low as the lateral geniculate nucleus, a visual relay station before the visual cortex (Blake 2001; Tong 2001; Haynes, Deichmann et al. 2005). Moreover, even though many agree that the FFA correlates with face perception, FFA activity is not sufficient for face perception. In unilateral neglect (which is also often studied in consciousness science), patients with lesions to their right parietal lobe are not consciously aware of objects presented in the left side of their visual field. However, there may nevertheless be FFA activity for faces presented to their left even though they don't report seeing the face. Thus, though the NCC search seems promising, it is not producing much in the way of univocal results so far (I review some recent mainly theoretical progress below) (Block 2001; though see also Dehaene and Naccache 2001; Dennett 2001).

One reason for this lack of progress is that there is as yet no firm theory about how the brain represents its environment, let alone how these representations become conscious. It is no wonder that it is hard to interpret activity correlated with consciousness if we still are not sure how the brain's wiring allows it to as much as represent the world. For binocular rivalry specifically, it is not known why a representational system such as the brain should decide to alternate between percepts. This is further complicated by our lack of a firm theory of attention, which makes it difficult to disentangle changes in consciousness from changes in attention. Thus, unilateral neglect is sometimes viewed as an attentional deficit, yet it clearly interacts with consciousness (Driver and Vuilleumier 2001); similarly, attentional factors influence the dominance ratios in binocular rivalry (Leopold and Logothetis 1999).

This situation is linked to larger issues affecting the progress of cognitive neuroscience. For a long time, there has been a focus on localising functions, using imaging techniques like fMRI. This is why we have become used to seeing pictures of brains with a few coloured blobs (a la "here is the feeling of love") in the media. But

there is now a renewed interest in going from such "blob-ology" to functional interconnectivity. It is hoped that this will solve some of the problems about how the brain represents (Friston 2002; Friston, Rothstein et al. 2006) and thus allow easier interpretation of NCC studies.

The move away from blob-ology towards functional interconnectivity has been accompanied by a growing interest in the temporal dynamics of both the brain's activity and of conscious experience. This allows exploration of a more "ecological" conception of conscious content of an embodied interacting subject (Shapiro 2006), rather than the "snapshot" conception often implicit (and sometimes explicit, see Koch 2004: 264-8) in much NCC research. This is an interesting area of research that promises to integrate computational (in particular, probabilistic) theories of the brain with neurophysiological theories of overall brain function and correlate these with more phenomenologically appealing notions of content. There is vigorous debate about how to conceptualise this kind of dynamics research as it relates to consciousness (Cosmelli, Lachaux et al. 2007; Grush forthcoming), and to human brain function (Raichle 2001; Friston 2002), as well as to probabilistic theories (Gregory 1980; Rao, Olshausen et al. 2002; Friston 2005).

These newer developments notwithstanding, many of the classic NCC studies depend quite closely on theories about the nature of representation. This reflects how many (but not all) of the NCC researchers conceive consciousness. Many seem to more or less equate full-blown or explicit representation with the notion of conscious content (Koch 2004). Some of these researchers would probably see the NCC search as completed if we could robustly correlate subjective reports of representational content with explicit representations in neural representational systems. (One may think of this conception of the NCC search in terms of the neural correlates of *explicit representation*.) This suggests that one's conception of consciousness is relevant to how one approaches the NCC search. In the next section we look at some philosophical discussions concerning this issue.

2. How our conception of consciousness may influence the NCC search.

We shall look at three criticisms of the orthodox NCC search that stem from different conceptions of consciousness: (i) that it focuses on access while neglecting phenomenality; (ii) that it is too atomistic; (iii) that it is too narrow.

(i) Ned Block has argued that consciousness is a mongrel concept (Block 1995). On the one hand, it concerns accessibility: how representational content is poised for direct control of reasoning, speech and action. On the other hand, it is tied to phenomenality: the *what it's like*-aspect of experience. Block argues that much of the current NCC work prioritises access over phenomenality (Block 2001). The risk is that the neural substrate of access consciousness is irrelevant for the neural substrate of phenomenality. In that case, we will not discover the NCC for the phenomenal character of conscious experience by focusing only on the neural substrate of what it takes for representational content to be accessible to reasoning, speech and action.

Though the distinction seems intuitive, our idea of consciousness is certainly challenged by the notion that one could be phenomenally conscious of a face yet not be aware of this fact, in the sense that one could not introspect or report it (see the Peer Commentary on Block 1995). However, for the purposes of the search for the

NCC, it seems reasonable to have an open mind. Nothing (short of controversial analyses of the concept of consciousness) seems to exclude *a priori* the possibility that phenomenality is constituted by some neural processes that do their work *before* the content becomes poised for direct control of reasoning, speech and action. This leaves the practical question of how researchers should go about discovering inaccessible phenomenality. The problem is that one cannot isolate phenomenality through subjective reports since reported states are of course accessed states. Block counsels use of indirect evidence for phenomenality (for example, using signal detection theory (Block 2005)), and cites a couple of promising recent studies of so-called reentrant processing, that has to do with neural feedback loops, as examples of an approach that focuses on phenomenality more than access (Lamme, Super et al. 1998; Lamme 2004). This strategy is in line with the common scientific practice of using inference to the best explanation of evidence gained in a variety of ways. Using this approach, whether we should believe in inaccessible phenomenal content will be an empirical matter.

The difficulty in isolating phenomenality connects to a number of questions that pervade the science of consciousness concerning what should qualify as evidence about the presence or nature of conscious states. Many of those critical of the whole NCC enterprise are suspicious of the subjective, seemingly untestable nature of subjective reports. On the other hand, the less sceptical tend to emphasise the trust we have in other subjects' reports of their conscious states, as well as the behavioural means we have of testing the reliability of the reports (see papers in Jack and Roepstorff 2003; 2004). For his part, Block advocates that we use subjective reports but also other sorts of scientific evidence (for criticism of Block, see, e.g., Dennett 2001).

There is lively debate about the trust we should have in subjects' reports of their conscious states. Some argue that progress will come from incorporating more detailed reports about how subjects interpret the experimenter's instructions (Jack and Roepstorff 2002). Some rely more on analysis of the necessary conditions for experience thrown up by traditional (Husserlian) phenomenology (Zahavi 2006); these analyses could be used to front-load experimental designs (Gallagher 2003), or to interpret existing fMRI data (Lloyd 2002). Some argue that phenomenology has the resources to radically alter the field (Thompson and Varela 2001; but see Bayne 2004), whereas others defend the current NCC-methodology (Hohwy and Frith 2004b).

(ii) The standard NCC approach is primarily interested in the neural substrate for having one rather than another *content* represented in consciousness (e.g., a percept of a face rather than of a house). This is distinguished from an interest in what it takes for a *creature* to be conscious rather than not conscious (Chalmers 2000; Koch 2004). This distinction seems initially plausible but John Searle (2000) objects that the content-based approach misses its target because it must presuppose that the creature in question already has a conscious 'unified field' in which the content is represented. If it is *presupposed* that the creature is conscious, then research on content NCC will hardly illuminate the underpinnings of consciousness. On the other hand, if it is not assumed that the creature is conscious, then research on content NCC is unlikely to be relevant to consciousness. Searle therefore advocates a renewed focus on the 'unified conscious field' which seems akin to the notion of creature consciousness, and which he views as in some sense prior to conscious content.

This debate cross-cuts with another common distinction in the field. As mentioned, the standard definition set out above concerns *minimal* content NCCs: the minimal neural substrate whose activity is sufficient for conscious experience of, say, a face. Few would accept that activity in this system on its own – cut out of the brain and put in a jar – would suffice for the experience. It seems clear that the minimal system must be hooked up to most of the rest of the brain to work. There must be some *enabling* conditions, in addition to the minimal conditions that *trigger* the system. So it makes sense to distinguish between the minimal or *core* NCC that triggers face perception and the *total* NCC that enables face perception (Chalmers 2000; Koch 2004; Block 2005).

Searle can be seen as insisting that there cannot be consciousness when a subject is creature *un*conscious, even though a particular core NCC is activated. But Searle's own notion of a unified conscious field seems equally problematic. He seems to say that content is irrelevant for the science of consciousness, which in turn seems to commit him to the idea that there can be creature consciousness without any content at all. But it is very hard to accept that an utterly empty 'unified field' is really conscious in any sense at all. Without a place for content, Searle's quest for consciousness therefore appears misguided too. In combination, this seems to pose a serious problem: both the search for the content NCC and for the creature NCC appear to miss their targets.

Perhaps the mistake here is to view creature and content consciousness as two distinct properties. An analogy can show this. When it comes to the solubility of a sugar cube we know that it depends on both triggering (whether we put the sugar cube in water) and enabling conditions (whether everything is normal about the water). If we don't put it in, then it will not dissolve; and if we put it in supercooled water, then it will not dissolve. But this does not tempt us to ascribe two distinct properties to the sugar cube: triggering-solubility and enabling-solubility. We know there is just one property, and it is the job of science to disentangle the triggering and enabling conditions in a theoretically interesting manner. Similarly, one could argue, we should not be tempted to ascribe two distinct properties to a subject: content consciousness and creature consciousness. There is one property, consciousness, which needs both triggering conditions (relevant to the content NCC) and enabling conditions (relevant to creature NCC). Then the problem does not arise because both types of conditions are necessary for the occurrence of consciousness; while exploring one set of conditions one must then assume that the other set is held constant.

(iii) There is a further way in which one's conception of conscious experience could have consequences for the NCC search, one that derives from phenomenology. To appreciate it, we first need to consider how researchers attempt to go beyond the approach to the NCC discussed thus far.

Straight off, it appears as if the search for the NCC is a bit like a fishing expedition: we have a rough idea of where to look (e.g., the brain, not the foot) and then we see what brain activity shows up in response to various tasks that engage consciousness. We do learn something from this: we would be very surprised if no neural activity at all correlated with the perceptual shifts in binocular rivalry. But such a 'catch' of evidence of neural activity is not going to be much help for furthering the study of

consciousness: it will just be a long, unsystematic list of areas of activity that can hardly be used in explanation, theorising or prediction concerning consciousness. To go beyond this type of "raw" correlation study, the neural correlate must therefore itself have some kind of systematicity (Chalmers 2000).

Such a general notion of systematicity is familiar from research of causal relations in science at large. For example, a number of events may correlate with a storm: preceding low pressure and low barometer readings. We then intervene on these prior events to see if there is any stable relation to the effect we wish to explain, the storm. We find that manipulating the barometer has no stable effect on the weather whereas low pressure changes do. In general, we go beyond mere correlation to causation by testing for such stable, or invariant, relations between the supposed cause and effect under various types of interventions. (This accords with an influential recent account of causation and causal explanation; see Woodward 2002; 2003). This type of testing requires some systematicity in the data; for without systematicity, it is impossible to identify factors on which to intervene. Analogously, the NCC search requires systematicity in the correlation data, so that researchers can intervene meaningfully in order to find the neural states that relate invariantly with the conscious states (see Cleeremans and Haynes 1999; Atkinson, Thomas et al. 2000; Fell 2004). The analogy is, however, limited, since the NCC search may concern constitution rather than causation. (Most researchers think brain states constitute conscious states rather than cause them.)

Stated in terms of intervention and invariance, there is no specific way in which systematicity should be cashed out. Thus invariance can be in on-off terms (suitable for study of creature consciousness), or it could be in terms of degrees of intensity (suitable for pain studies, say). However, as we notice earlier, in much NCC research systematicity is cashed out in terms of invariance in the modulation of the representational *content* presented in consciousness (Chalmers 2000). On this *content* NCC approach, one must independently identify the systematic neural *representational* systems for the conscious contents one is interested in. This often involves mapping the receptive field of certain populations of neurons (e.g., recording neural activity that occurs specifically when face stimuli are shown) and then subsequently looking for the most minimal part of the neural system whose activity correlates with reports of conscious experience of that content. The hope is then that theoretically motivated interventions on the neural system will relate invariantly to changes in the conscious state, which would allow a deeper understanding of the mechanisms that constitute consciousness than mere "fishing expeditions" do.

With these aspects of the NCC search in place, we can now turn to the phenomenological objection.

We associate the notion of 'representation' closely with our own conscious experiences that often represent the world around us. Therefore it is a little hard to understand what it is for an unconscious neural system, defined in the more mechanistic terms of neural receptive fields, to be 'representational'. Some researchers have seized on this to criticise the content NCC approach. They claim that on a successful content NCC approach there must be a perfect *match* between the experiential content and the neural content. Then they use phenomenological analysis to identify some central properties of conscious experience, and argue that these

properties cannot be perfectly matched with neural content properties. One such property is that experience is *perceptually coherent* (e.g., that things are never experienced on their own but always as being in a figure-ground context). Thus, for example, one will change the experience of a face if one alters the context in which it is seen. Then, the argument goes, it is a mistake to think the neural representational system for faces is responsible for face perception since it is atomistic and cannot capture the context-dependence of face perception. Put differently, it is difficult to see how one systematically could capture the neural receptive field for figure-ground modulation, so it is difficult to see how a neural match for conscious figure-ground perception could be found. If this is right then the search for the correlates of consciousness should not be limited to the brain. In fact, some of these researchers advocate a much broader approach, according to which the correlate of consciousness includes the body and parts of the environment (Noë 2004; Noë and Thompson 2004). This seems like a radical departure from the current paradigm in the NCC search.

It would be natural for the orthodox NCC researcher to answer by accepting that there are these intricate properties of experience but insisting that the fledgling science of consciousness is as yet incapable of studying them. Instead, one could add, there is much useful study of more simple and partial percepts, such as faces vs. houses, or vertical vs. horizontal gratings (see, e.g., Koch 2004). But this response is not quite adequate, for the phenomenological objection denies that there are any such simple or partial percepts; according to the objection, the only percepts are perceptually coherent ones. Moreover, the objection stems from the orthodox approach's own adherence to the systematicity requirement, which the objectors interpret in terms of precise matching of neural and conscious contents. It would therefore be desirable to find a different interpretation of the requirement, one which does not invite that problematic interpretation. One possibility here is to stress the analogy to causal research: systematicity is achieved by identifying invariant relations under interventions on neural representational systems, rather than by establishing identities between contents. In this way, identifying a content NCC for faces would involve identifying an invariant relation between neural representations and different experiences of faces under interventions on the neural system where these interventions could very well involve changing the context such as to modulate the perceptual coherence of faces (see Hohwy and Frith 2004a). And indeed there is a renewed neurocomputational focus on how context in general modulates what and how things are represented (e.g., Rao, Olshausen et al. 2002).

Summing up on these three issues, it is clear that different conceptions of consciousness can have different consequences for the NCC search. It transpires, however, that the nature of these consequences, and the seriousness of the problems they may generate for the orthodox NCC search, depend heavily on issues in philosophy of science concerning evidence for unobservables, the nature of triggering and enabling conditions, and how one passes from "raw" correlation searches to more systematic searches.

3. Theories of consciousness, the NCC, and metaphysics

Most people view the NCC search as neutral with respect to the metaphysics of consciousness. This is remarkable given the deep differences in theories of consciousness and in estimations of the prospects for arriving at a reductive theory of phenomenality at all (for a review, see Kriegel 2006). Thus, some reductive

physicalists would simply *identify* the phenomenal with its neural correlate; functionalists think the NCC *realises* phenomenal functional roles; dualists think it somehow co-varies or *causally* interacts with the phenomenal. (Perhaps idealists, who claim that reality is fundamentally mental, would deny that there are non-conscious neural correlates of consciousness).

It is hard to think of other areas of science where people agree on the fundamental methodology for empirical inquiry yet have wildly different opinions about the metaphysics of the phenomenon under scrutiny (setting aside differences in conception of the phenomenon, discussed above). However, it seems correct that the NCC strategy is not wedded to any particular theory of consciousness. A number of factors may explain this.

There is overwhelming evidence that the brain is somehow involved in maintaining conscious states. When the brain is lesioned, changes in conscious state may immediately arise; similarly, psychotropic drugs like LSD operate through interaction with the brain. Every remotely reasonable theory of consciousness must therefore find a role for the brain. That is, no remotely reasonable theory can afford to claim that correlating conscious states with brain states is completely irrelevant to a full understanding of consciousness.

Moreover, it is very clear that the role of the brain for consciousness is systematic: the relations between neural structures and conscious states are not erratic or random. (Otherwise it is difficult to see, for example, how drugs like Prozac reliably work to lift depression.). So all reasonable theories take systematicity on board and expect at least the kind of invariance under interventions mentioned above. It seems fair to say that this sensitivity to empirical findings comes back to haunt the majority of the metaphysical positions, namely when we consider how they can really be consistent with, e.g., mental causation (for excellent discussion of this, see Kim 1998).

It seems, then, that one's prior metaphysical theory of consciousness isn't terribly important for how one embarks on NCC research. In contrast, one's empirical hypotheses about consciousness are very relevant: one's studies of the brain in search of the NCC will be guided by some kind of hypothesis about where and how one should look to 'see' consciousness. (E.g., does one use fMRI, or EEG (Frith 2001; Revonsuo 2001)? Does one focus on neural assemblies where the 'winner takes all' (Koch 2004), or on re-entrant processing (Lamme 2004)?). It seems natural to suppose here that gradual discovery of the NCC will go hand in hand with the gradual development of theories about consciousness and the brain. This is a common pattern in the development of science (Lipton 1991; Hohwy 2004). Given the immense, and widely agreed upon, difficulty in developing theories of consciousness, this might mean that the discovery of the NCC will be very slow. But notice that discovering the NCC does not require something as dramatic as bridging the explanatory gap between neuroscience and consciousness (Chalmers 1996; Levine 2001). It simply requires arriving at a good neuroscientific theory about the workings of the neural systems that underpin phenomenality.

Though the various metaphysical theories appear neutral with respect to the NCC, it seems that full discovery of the NCC could have consequences for our view of the metaphysics of consciousness. For consider what we should say if highly plausible,

strongly systematic NCCs were found for most types of conscious states in most humans (and perhaps monkeys). Those of a reductionist bent would probably be tempted to say that consciousness is then *identical* with those brain states – that is, we would be tempted to accept the identity theory of consciousness (Smart 1959; Hill 1991). This is not an unreasonable inference – it seems not too different from what happened when we discovered the identity of water and H_2O, or of Lawrence of Arabia and private T. E. Shaw. For the identity theorists, identities cannot be explained (Papineau 1998), so for them, the systematic NCC could be the only empirical step required in solving the problem of consciousness. For those attracted to other metaphysical theories, there would be further empirical questions to ask.

Would our appetite for explanation really be satisfied with discovery of the NCC in the manner suggested by the identity theorists? Some think not, because an unexplained identity would have to treat physical and phenomenal properties as *distinct* fundamental properties and that would leave the metaphysical status of the phenomenal an open question. Something extra would be needed to persuade us that this fundamental property is actually a *physical* property. That is to say, it is left an open question whether they are, in fact, identical (Chalmers 2002; Kim 2005).

There is a further problem for the identity theorist, which is more closely related to the NCC methodology. Consider a creature functionally similar to us humans yet whose physical realization is fundamentally different (e.g., the android Commander Data from Star Trek). Nothing seems to prohibit that we employ NCC methodology to this creature to find the android correlate of its reports of being in conscious states. And if we find it, then we should, by parity of reasoning, identify the android states with its consciousness. But when we acknowledge that the physical realization is so different from ours we should hesitate because it seems we cannot be sure the creature really has any phenomenal states at all (hence the dramatic effect of Data in the TV show). Indeed we may have no conception of how to even begin finding out if it is conscious or not. Block calls this the 'harder problem of consciousness' (2002; for discussion, see McLaughlin 2003; Hohwy 2004). For creatures that are physically similar to us (other humans, monkeys, dogs, squids…), we don't have the same problem, for there we can rely on the argument from analogy to establish that these creatures have minds like ours. But this problem suggests that the NCC approach is not on its own enough to warrant the inference to identity and thus to resolve the problem of consciousness. In the case of Data, the NCC approach would allow us to make inferences about how he represents the world; but we would have serious doubts about whether those inferences are relevant for questions concerning android consciousness. Transposing this back to the human case, it is as if we say "these findings support inferences about how the brain represents – oh, and given that *we* are conscious it supports inferences about consciousness too". In other words, it seems as if, in addition to doing the scientific NCC work, the identity theorist must secure consciousness a place by making the additional, unscientific assumption that the system under scrutiny is in fact conscious. The idea that consciousness has to be somehow added to the scientific data is unpleasantly reminiscent of the traditional problems concerning consciousness and suggests that the NCC on its own may be insufficient for justifying metaphysical inferences about consciousness.

Concluding remarks

Though researchers of many different kinds of metaphysical outlook agree that discovery of the NCC should be a prime goal for a science of consciousness, there is much disagreement about which conception of consciousness should be prioritised in this search. Still, the science of consciousness has progressed so much that the debates about conflicting conceptions may profit from some standard lessons from the philosophy of science. Identification of the NCC may influence the debate about the metaphysics of consciousness, but it will not establish the identity theory of consciousness on its own. It is not clear therefore what the next step will be, once the NCC has been identified.

Alais, D. and R. Blake, Eds. (2005). *Binocular Rivalry*. Cambridge, Mass., MIT Press.

Atkinson, A. P., M. S. C. Thomas, et al. (2000). "Consciousness: mapping the theoretical landscape." *Trends in Cognitive Sciences* **4**(10): 372.

Bayne, T. (2004). "Closing the Gap? Some Questions for Neurophenomenology " *Phenomenology and the Cognitive Sciences* **3/4**: 349-364.

Blake, R. (2001). "A Primer on Binocular Rivalry, Including Current Controversies." *Brain and Mind* **2**: 5-38.

Block, N. (1995). "On a confusion about a function of consciousness." *Behavioral and Brain Sciences* **18**: 227-287.

Block, N. (2001). "Paradox and cross purposes in recent work on consciousness." *Cognition* **79**: 197-219.

Block, N. (2002). " The Harder Problem of Consciousness." *The Journal of Philosophy* **XCIX**(8): 1-35.

Block, N. (2005). "Two neural correlates of consciousness." *Trends in Cognitive Sciences* **9**(2): 46.

Chalmers, D. (1996). *The Conscious Mind*. Harvard, Oxford University Press.

Chalmers, D. (2000). What is a neural correlate of consciousness? *Neural Correlates of Consciousness: Empirical and Conceptual Issues*. T. Metzinger. Cambridge, Mass., MIT Press.

Chalmers, D. (2002). Consciousness and its place in nature. *Philosophy of Mind*. D. J. Chalmers. Oxford, Oxford University Press.

Cleeremans, A. and J.-D. Haynes (1999). "Correlating Consciousness: A view from empirical science." *Revue Internationale de Philosophie* **3**: 387-420.

Cosmelli, D., J.-P. Lachaux, et al. (2007). Neurodynamics of Consciousness. *The Cambridge Handbook of Consciousness*. P. D. Zelazo, M. Moscovitch and E. Evan Thompson. Cambridge, Cambridge University Press.

Dehaene, S. and L. Naccache (2001). "Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework." *Cognition* **79**: 1–37.

Dennett, D. (2001). "Are we explaining consciousness yet?" *Cognition* **79**(1-2): 221.

Driver, J. and P. Vuilleumier (2001). "Perceptual awareness and its loss in unilateral neglect and extinction." *Cognition* **79**(1-2): 39-88.

Fell, J. (2004). "Identifying neural correlates of consciousness: The state space approach." *Consciousness and Cognition* **13**(4): 709-729.

Friston, K. (2002). "Functional integration and inference in the brain." *Progress in Neurobiology* **68**: 113-143.

Friston, K. J. (2005). "A theory of cortical responses." *Philosophical Transactions: Biological Sciences* **369**(1456): 815 - 836.

Friston, K. J., P. Rothstein, et al. (2006). "A critique of functional localisers." *NeuroImage* **30**(4): 1077.

Frith, C., R. Perry, et al. (1999). "The neural correlates of conscious experience: an experimental framework." *Trends in Cognitive Sciences* **3**(3): 105.

Frith, C. D. (2001). "Commentary on Revonsuo's 'Can functional brain imaging discover consciousness in the brain?" *Journal of Consciousness Studies* **8**(3): 30-30.

Gallagher, S. (2003). "Phenomenology and Experimental Design Toward a Phenomenologically Enlightened Experimental Science " *Journal of Consciousness Studies* **10**(9-10): 85-99.

Gregory, R. L. (1980). "Perceptions as hypotheses." *Phil. Trans. R. Soc. Lond., Series B, Biological Sciences* **290**(1038): 181-197.

Grush, R. (forthcoming). "How to, and how not to, bridge computational cognitive neuroscience and Husserlian phenomenology of time consciousness."

Haynes, J. D., R. Deichmann, et al. (2005). "Eye-specific effects of binocular rivalry in the human lateral geniculate nucleus." *Nature* **438**: 496.

Hill, C. (1991). *Sensations: A Defense of Type Materialism*. Cambridge, Cambridge University Press.

Hohwy, J. (2004). "Evidence, explanation, and experience: On the harder problem of consciousness." *The Journal of Philosophy* **CI**(5): 242-254.

Hohwy, J. and C. Frith (2004a). "The Neural Correlates of Consciousness: Room for improvement, but on the right track." *Journal of Consciousness Studies*.

Hohwy, J. and C. Frith (2004b). "Can neuroscience explain consciousness?" *Journal of Consciousness Studies* **11**(7-8): 180-198.

Jack, A. I. and A. Roepstorff (2002). " Introspection and cognitive brain mapping: from stimulus-response to script-report." *Trends in Cognitive Sciences* **6**(8): 333-339.

Jack, A. I. and A. Roepstorff, Eds. (2003). *Trusting the Subject*. Exeter, Imprint Academic.

Jack, A. I. and A. Roepstorff, Eds. (2004). *Trusting the Subject*. Exeter, Imprint Academic.

Kanwisher, N. (2001). "Neural events and perceptual awareness." *Cognition* **79**: 89-113.

Kim, J. (1998). *Mind in a Physical World*. Cambridge, Mass., MIT Press.

Kim, J. (2005). *Physicalism; or Something Near Enough*. Princeton, N. J., Princeton University Press.

Koch, C. (2004). The Quest for Consciousness: A Neurobiological Approach.

Kriegel, U. (2006). "Consciousness, Theories of." *Philosophy Compass* **1**(1): 58-64.

Lamme, V. A. F. (2004). "Separate neural definitions of visual consciousness and visual attention; a case for phenomenal awareness." *Neural Networks* **17**(5-6): 861.

Lamme, V. A. F., H. Super, et al. (1998). "Feedforward, horizontal, and feedback processing in the visual cortex." *Current Opinion in Neurobiology* **8**(4): 529.

Leopold, D. A. and N. K. Logothetis (1999). "Multistable phenomena: changing views in perception." *Trends Cogn. Sci.* **3**: 254-264.

Levine, J. (2001). *Purple Haze: The Puzzle of Consciousness*. Oxford, Oxford University Press.

Lipton, P. (1991). *Inference to the Best Explanation* London, Routledge.

Lloyd, D. (2002). "Functional MRI and the Study of Human Consciousness." *J. Cogn. Neurosci.* **14**(6): 818-831.

McLaughlin, B. (2003). "A Naturalist-Phenomenal Realist Response To Block's Harder Problem." *Philosophical Issues* **13**: 163-204.

Noë, A. (2004). *Action in Perception*. Cambridge, Mass., MIT Press.

Noë, A. and E. Thompson (2004). "Are there neural correlates of consciousness?" *Journal of Consciousness Studies* **11**(1): 3-28.

Papineau, D. (1998). "Mind the Gap." *Nous* **32**(S12): 373-388.

Raichle, M. E. (2001). "A default mode of brain function." *Proc Natl Acad Sci U S A* **98**(2): 676-82.

Rao, R. P. N., B. A. Olshausen, et al. (2002). *Probabilistic Models of the Brain*. Cambridge, Mass., MIT Press.

Revonsuo, A. (2001). "Can functional brain imaging discover consciousness in the brain?" *journal of Consciousness Studies* **8**(3): 3-23.

Searle, J. R. (2000). "Consciousness." *Annual Review of Neuroscience* **23**(1): 557-578.

Shapiro, L. (2006). "The Embodied Cognition Research Program." *Philosophy Compass* **1**(1).

Smart, J. J. C. (1959). "Sensations and Brain Processes." *Philosophical Review* **68**: 141-156.

Thompson, E. and F. J. Varela (2001). "Radical embodiment: neural dynamics and consciousness." *Trends in Cognitive Sciences* **5**(10): 418.

Tong, F. (2001). "Competing theories of binocular rivalry: a possible resolution." *Brain Mind* **2**: 55-83.

Tong, F., K. Nakayama, et al. (1998). "Binocular rivalry and visual awareness in human extrastriate cortex." *Neuron* **21**: 753-759.

Woodward, J. (2002). "What is a mechanism? A counterfactual account." *Philosophy of Science* **69**: S366-S377.

Woodward, J. (2003). *Making Things Happen*. New York, Oxford University Press.

Zahavi, D. (2006). *Subjectivity and Selfhood: Investigating the First-Person Perspective*. Cambridge, Mass., MIT Press.