


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Consciousness and Cognition

journal homepage: www.elsevier.com/locate/concog

The neural correlates of consciousness: New experimental approaches needed?

Jakob Hohwy*

School of Philosophy & Bioethics, Monash University, Australia

ARTICLE INFO

Article history:

Received 30 July 2008

Available online 3 April 2009

Keywords:

Neural correlates of consciousness

Content consciousness

State consciousness

Binocular rivalry

Vegetative state

Philosophy of neuroscience

ABSTRACT

It appears that consciousness science is progressing soundly, in particular in its search for the neural correlates of consciousness. There are two main approaches to this search, one is content-based (focusing on the contrast between conscious perception of, e.g., faces vs. houses), the other is state-based (focusing on overall conscious states, e.g., the contrast between dreamless sleep vs. the awake state). Methodological and conceptual considerations of a number of concrete studies show that both approaches are problematic: the content-based approach seems to set aside crucial aspects of consciousness; and the state-based approach seems over-inclusive in a way that is hard to rectify without losing sight of the crucial conscious–unconscious contrast. Consequently, the search for the neural correlates of consciousness is in need of new experimental paradigms.

© 2009 Elsevier Inc. All rights reserved.

1. Introduction

Finding the neural correlates of consciousness (the NCC) is a major focus of cognitive neuroscience. There are two dominant methodological approaches to finding the NCC (set out early by Baars, 1988). One approach is directed at finding the neural correlates of conscious *content* (such as the conscious perception of a face rather than a house). Another approach is directed at finding the neural correlates of a creature's overall *state* of consciousness (such as being awake rather than in dreamless sleep or a vegetative state). It is shown here that, from the point of view of philosophy of science, these approaches are problematic both conceptually and methodologically. In Section 3, I support and develop the view that the first, content-based approach is methodologically compromised because it is unable to target a crucial aspect of consciousness, and I show that reflection on this problem casts doubt on whether the content-based approach is truly directed at a property of being conscious at all. In Section 4, I show how the second, state-based approach has oddly contradictory findings, it risks being over-inclusive, and attempts to evade these problems seem to make the contrast between conscious and unconscious states slip away or make the approach collapse to the problematic content-based approach. I conclude that the field is in need of new experimental approaches.

2. The two approaches to the neural correlates of consciousness

A discussion by John Searle is useful to introduce the two different approaches and the initial problem for the content-based approach (notice that Searle uses the term “unified conscious field” as a cognate term for state of consciousness; later in this introductory Section I elaborate on some of the terminology common in the debate).

* Fax: +61 3 9905 3206.

E-mail address: jakob.hohwy@arts.monash.edu.au

URL: <http://www.arts.monash.edu.au/philosophy/staff/jhohwy.php>

Searle is arguing against the approach taken by for example Christof Koch, who favours the content-based approach to the NCC. In a review of Koch's influential book *The Quest for Consciousness*, Searle says:

I think [Koch's] overall conception of consciousness and how it fits in with the rest of the world is flawed (Searle, 2005: Section 4) [...] Another approach, much less commonly pursued, I call the unified field approach. Here the aim is not to find the NCC for this or that conscious experience but to find out how the brain creates the unified conscious field in the first place (Searle, 2005: Section 1). [...] I am very skeptical [about Koch's approach] because the subjects on whom these [binocular rivalry] experiments are performed are already conscious. They already have a unified field of qualitative, subjective intentionality. So the most we can reasonably expect from this research is an explanation of how, within a brain that is already conscious, we can cause this or that perceptual experience. But that does not give us the NCC for consciousness as such. Perception of the sort that Koch is investigating does not create consciousness but modifies a preexisting conscious field. Koch is pursuing an important line of investigation, but so far there is no reason to suppose it will explain how the brain creates the conscious field (Searle, 2005: Section 2).

Clearly, Searle believes the searches for the correlates of content and state consciousness can be kept distinct and that the discoveries from the content NCC search are irrelevant to the study of how consciousness arises (see also Searle, 2000, pp. 272–274; Searle, 2004, chap. 5).

Koch, on his side, is fully aware that something is left out of the picture by pursuing a content-NCC approach:

In considering the NCC, it is important to distinguish *enabling* factors and *specific* factors. Enabling factors are tonic conditions and systems that are needed for any form of consciousness to occur at all, while specific factors are required for any one specific conscious percept. [...] By definition, no NCC can form without the [enabling-] NCC (Koch, 2004, p. 88).

He doesn't exclude the possibility that there could be content-less unified conscious fields underpinned by an enabling NCC (he considers meditation as one context where it may arise). Koch's reasons for focusing on conscious content seem mainly methodological; it is what we can study best at the moment, though he does call the specific NCC the "NCC proper" (Koch, 2004, p. 97). However, though Koch recognises the potential validity of states of consciousness, he does not discuss how the specific and enabling NCCs may fit together.

The intuitions driving Searle seem right: there is something odd about the content-based NCC approach if you have to presuppose and set aside consciousness in order to study it. But Koch also seems right to insist that it matters to consciousness which contents are represented "in it". They are both suggesting distinct research strategies: we can meaningfully set aside state consciousness whilst solving for content consciousness, and *vice versa*.

Here is a characterisation of the searches for the neural correlates of consciousness (for background, see Chalmers, 2000):

NCC for content consciousness: the minimally sufficient neural conditions for a specific (mostly representational) content being conscious rather than not being conscious.

The basic methodology was set out early by Baars (1988) and a core example of this approach is Tong and colleagues (1998) in a fMRI study of binocular rivalry. In binocular rivalry, one stimulus (a picture of a face, say) is presented to one eye and another (a picture of a house) to the other eye of a conscious subject. Conscious experience shifts between the two every few seconds (for a recent review, see Tong, Meng, Blake, 2006). Binocular rivalry is used in paradigm studies in consciousness science since the stimuli are held constant and so brain activity that correlates with reported perceptual shifts is likely to reveal the neural substrate of consciousness. This particular study showed that activity in the fusiform face area correlates with reports of conscious face perception. The picture is complicated, however, since there can be activity in that area even if there is no face perception (Rees et al., 2000); furthermore, activity in many areas, including V1 and LGN also seem to be correlated with conscious face perception making interpretation of these results difficult (Haynes, Deichmann, & Rees, 2005; Polonsky, Blake, Braun, & Heeger, 2000; Tong & Engel, 2001). Other examples of this kind of approach include paradigms using inattention blindness, neglect, blindsight, and various sorts of masking (see Frith, Perry, & Lumer, 1999; Rees & Frith, 2007 for reviews and discussion). I assess these kinds of studies, in the light of the Searlian complaint, in Section 3.

Since the idea of finding the correlates, neural or otherwise, of something seems quite general the state NCC search is characterised in a parallel way:

NCC for states of consciousness: the minimally sufficient conditions for a creature's being in an overall conscious state rather than an overall unconscious state.

As defined here, a study counts as state-based if it is directed at the contrast between being conscious rather than being unconscious. This clearly distinguishes it from the content-based approach, which cannot be directed at this contrast since it presupposes that the subjects in question are in a conscious state. It is relatively rare for researchers to explicitly identify their research as directly relevant to the state-based approach to the NCC. However, there are many studies, particularly those concerned with global disorders of consciousness, that seem to concern the neural substrate for the contrast between being in a conscious state rather than an unconscious state (see, e.g., Laureys (2005) for a suggestive review; see also Schiff, 2007). This makes such studies candidates for being state-based even though their primary concerns in many cases may be clinical and diagnostic, often having to do with patients in coma or vegetative state, or with anaesthesia (e.g., Laureys & Boly, 2007, p. 610). Section 4 discusses this approach.

A core example of the state-based approach is Laureys et al. (2000, see also Laureys, 1999). In this study, a patient in a vegetative state (VS) was contrasted with healthy controls and with himself after recovery. Analysis of effective connectivity showed that, in contrast to the VS patients, the healthy controls and the recovered patients had a specific pattern of corticothalamic activity. This may then be part of the neural correlate of state consciousness.

Other types of studies are less straightforward to categorise as state-based even though they differ from the content-based approach by not presupposing that the subjects in question are already conscious. Here are some examples. Children with hydranencephaly are born with almost no cortex but an intact brain stem. These children are commonly diagnosed as being in a vegetative state but Merker (2007, see also Shewmon, Holmes, & Byrne, 1999) argues that they display behaviour that suggests they are in fact in a conscious state. Therefore the brain stem may be the neural correlate of state consciousness. Schiff et al. (2007) used deep brain electric stimulation to the central thalamus in a patient who had been in a minimally conscious state for 6 years. He subsequently developed behaviour suggesting he was in a more fully conscious state. This supports the role of thalamocortical connectivity for state consciousness, as does a range of studies of anaesthesia (see Alkire, Miller, & Laureys, 2005 for review). Owen et al. (2006) found similarities in the brain responses in healthy controls and a VS patient to verbal instructions and concluded that the patient must have been in a conscious state.

It is natural to think of the state NCC plus the (union of) content NCCs as the *total* NCC. And, hence, to think of the total NCC minus content NCCs as the state NCC. Certain other sufficient conditions will be left out as not minimally sufficient for either phenomenon (for a more general introduction and motivation for the appeal to 'minimal sufficiency', see Chalmers, 2000; Hohwy, 2007). This seems to be an appropriate characterisation of the NCC searches, if we are to take the debate between Searle and Koch seriously. The issue between Koch and Searle concerns which of the two is primary – the upshot of my discussion is that, on their own, each is problematic.

The terminology in this area is unfortunately not uniform. Luckily, this is not so important to the approach I take here. The starting-point for my discussion is simply the fact that NCC studies cluster around those that focus on the contrast between conscious states and unconscious states, and those that focus on the contrast between different conscious contents in creatures already in a conscious state. It is hard to deny this fact about the field, irrespective of where one stands on the various terminological issues. At this stage, we can take these two different experimental contrasts to define operationally our notions of state and content consciousness. This is sufficient to generate a variety of puzzles and problems. (In Appendix A, I situate my debate in relation to some of the terminological and conceptual distinctions in the general debate about consciousness. Readers who are comfortable with the terminology so far may want to skip this).

3. The content-based approach to the NCC

The notion of minimal sufficiency is a pragmatic one. It concerns what, in certain background conditions, *makes a particular difference* in some target phenomenon (cf. Lipton, 2004). What we decide to leave in or out of the minimally sufficient conditions therefore depends on what the particular difference and target phenomena are, as well as on what the background conditions are taken to be. The particular difference will often be a contrast such as "did the target phenomenon occur rather than not occur?", "did the target phenomenon occur in a certain way (e.g., with a certain intensity) rather than in a different way (e.g., with a different intensity)?"

The content NCC search, as conceived by Koch and many others, is thought to have as its contrastive target finding what makes the neural difference between a content's being conscious and it not being conscious, on the background assumption that the creature is already in an overall conscious state (see, e.g., the rivalry study by Tong et al., 1998 described above).

3.1. Is the content-based approach methodologically clear? An analogy

This is the dominant paradigm in the field but – as Searle observes – it is very puzzling. For what role does the background assumption of the creature's being in an overall conscious state play here? Here is an analogy to illustrate the puzzle. Assume a researcher wants to find out what it is for an organism to have a certain kind of cancer. One group of mice is exposed to a carcinogen and another is not. If the researcher assumes that all the mice already have cancer, then the experimental variable that is manipulated, exposure to the carcinogen, is insufficient to show what it is to have this type of cancer. It would be justified to complain that the cancer may very well be due to something entirely different than this carcinogen or that without the existing cancer the trajectory could have been entirely different. This holds even if there is an effect of the exposure, such as making the cancer spread more rapidly in the first group of mice. Similarly, if a researcher wants to find out what it is for a content to be conscious, then she may induce this content in one set of conditions and not in another closely matched set of conditions. If the researcher assumes that subjects in both sets of conditions are conscious throughout, then the experimental variable that is manipulated, inducing a conscious content rather than another, is insufficient to show what it is for that content to be conscious rather than not. It would be justified to complain that the content's being *conscious* rather than not could be due to something entirely different than inducing this particular content. This is so even if there is an effect of inducing the content, such as making subjects in this condition conscious of a face rather than of a house.

An objection to this line of argument is that the analogy is flawed. In the carcinogen case, the researcher's study design is flawed because she is interested in the property of being a certain kind of cancer, which is the very property whose instantiation is presupposed in all the mice. In the content NCC case, researchers are interested in the property of a *content's* being

conscious but they don't fail because what they presuppose is the different property of the creature's being in an overall conscious *state*. The analogy can be elaborated to deal with this objection. Assume it is found in the carcinogen case that the mice share cancer of type *A* at the outset but that the mice exposed to the supposed carcinogen end up with a cancer of type *B*. Assume moreover that the researchers are clear that they are interested in type *B* cancers. Now the analogy is closer to the content NCC case. However, if the researchers want an answer to whether the substance is a type *B* carcinogen, then they still have an ill-designed study. It might still be the case that the substance is a type *B* carcinogen only because the mice already had the type *A* cancer. The pre-existing type *A* cancer could very well be the difference-maker for type *B* cancer because, without the *A* cancer, the substance might be benign. Similarly, if the content NCC researchers want an answer to what neural activity is minimally sufficient for a content's being conscious, then presupposing state consciousness introduces a serious doubt: it may be that the content is conscious only because the subjects were already in a conscious state. The pre-existing conscious state could very well be the difference-maker for the content's being conscious. Given that the NCC programme is focused on minimal sufficiency, and that we explain minimal sufficiency in the pragmatic terms of being difference-makers for a particular contrast in the target phenomenon, this shows that presupposing conscious states at least potentially introduces a confounding factor in the standard content-based NCC approach.

3.2. *Being selected for consciousness vs. being conscious*

It would be a mistake to conclude from this puzzle that the content NCC search is wholly confused. But it is crucial to formulate its contrastive target carefully. The target seems to be to find the neural correlate of a content's being *selected* for conscious presentation, rather than not selected, on the background assumption that the creature is already in an overall conscious state. This conception makes sense of the idea that we believe there are both conscious and unconscious contents in an otherwise conscious creature, and something makes the difference between a content's being selected for conscious presentation rather than not. In the analogy, it is like asking the legitimate question of why in cancerous mice, some cancers become type *B* cancers, and finding that it is linked to exposure to a certain carcinogen.

This shows there are two types of questions to ask in the state NCC search. One question concerns a content's *being* conscious; we have seen that this question cannot be answered if the creature in question is in an overall unconscious state throughout but likewise that it cannot be answered if the creature is presupposed to be in a conscious state throughout – in those cases the right experimental variable cannot be manipulated. The other type of question concerns a content's being *selected* for consciousness, and for this question it makes very good sense to assume that the creature is in the overall state of being conscious throughout. However, if it is the second type of question then the content NCC search will not be relevant for revelation of the neural systems that constitute the property of a content's being conscious. It will rather reveal the criterion for selecting a content for consciousness, which is something different altogether. To illustrate, you may want to know what it is like to be a student at a prestigious high school but justifiably refrain from taking a list of the selection criteria for admittance to the school as an informative answer. Searle is right that the content NCC is irrelevant to consciousness, but – crucially – only when the target is the correlate of a content's being conscious. It is relevant when the question concerns selection for consciousness.

One could complain about the content-based approach that it doesn't deliver sufficient conditions for content consciousness. But this would be a poor complaint. Its proponents are clear enough that they want to uncover only the *minimally* sufficient conditions. The Searlian complaint is different, however. It says that the approach is not likely to find even minimally sufficient conditions for a content's being conscious because what has been excluded from the minimally sufficient conditions, and relegated to the status of background condition, is the correlate for being conscious in the first place. At best, I now say, it will find the minimally sufficient conditions for selection into consciousness. This is reflected in the quote from Searle above that the content-based NCC approach will show at most “how, within a brain that is already conscious, we can cause this or that perceptual experience”. This may tell us interesting things, and it may be necessary for understanding consciousness in general (such as why content processed in a certain way tends to be selected for conscious presentation), but it will not help in arriving at the neural correlates of phenomenality, of *what it is likeness* as such.

In summary, the content-based NCC approach asks what it is for this content rather than that content to be conscious rather than not conscious. There are two contrasts in this question, indicated by the two occurrences of ‘rather’. The approach is better suited for addressing the first contrast, concerning why one content rather than another is represented, and is not very well suited for addressing the second contrast, concerning why the chosen content is conscious rather than not conscious. The basic issue is that, by focusing on the first contrast, any contribution from the neural correlates of state consciousness to the constitution of a content's being conscious is effectively screened off.

3.3. *What is the content-based approach really directed at?*

With this result about the content NCC, the worry next arises in what sense the content-based NCC approach is in fact about consciousness. Imagine a team of researchers that are interested in finding the neural correlates of some creatures' ability to *represent* the world but who have no interest in whether its study subjects are phenomenally conscious or not. These researchers might work on the assumption that in adaptive systems, only one amongst a number of candidates for representational content will win or will be processed up to the highest level in the cortex. They might, that is, be interested in the neural correlates of explicit (or invariant) representation (Koch, 2004, pp. 25–28; see also O'Brien & Opie, 1999). The

researchers will be well advised to adopt the content-based NCC approach to find the neural correlates of explicit representation since that may pinpoint the neural structures responsible for the winning representation. They can be completely neutral about whether the creature is in fact in an overall conscious state or not and yet gain very fruitful results using this method. In fact, the results would likely be very much like the results gained by those that claim that this method is appropriate for the study of conscious content. But if this is the case then the assumption in question, that the creature is already in an overall conscious state and that these contents are selected into *consciousness*, plays no instrumental role for the execution of this method on any given creature. The question that arises is why we should think that this method is then specifically directed at the property of a content's being conscious, if it is just as well directed at a different phenomenon, such as representation (relevant examples may be, Quiroga, Reddy, Kreiman, Koch, & Fried, 2005; Quiroga, Koch, & Fried, 2007).

Of course, in creatures that are in fact conscious, like us, the content-based NCC approach might indirectly tell us something, namely that in already conscious creatures only explicit content makes it to conscious presentation. This is very valuable but the actual neural structures that are directly revealed are not the correlates of what it is for those contents to be conscious, rather they are correlates of what it takes for them to be explicitly represented. This suggests that all the relevant work for consciousness is done by other mechanisms, presumably, one should think, having to do with the enabling conditions for state consciousness and its causal interactions with representational content.

4. The state-based approach to the NCC

It could be thought that if, as Searle suggests and I have just elaborated and developed at length, the content-based NCC approach is not directed at what it is for a content to be conscious, then researchers should discard the content-based approach and adopt the state-based NCC approach on which the crucial contrast is between creatures in overall conscious and in unconscious states. But, contra Searle, this is not what I shall advocate here. It turns out that the state-based approach faces a number of problems: it can be over-inclusive and it has oddly contradictory findings. This situation engages some venerable philosophical problems and makes it difficult to operationalise its basic contrast between conscious and unconscious states without changing the topic away from consciousness or appealing to the already problematic content-based approach.

The state-based NCC approach aims to uncover the minimally sufficient conditions for a creature's being in an overall conscious state rather than an unconscious state. A background assumption here is, for example, that the creature is alive. Sufficient conditions for state consciousness that are more relevant for whether the creature is alive or not than for whether the creature is in a conscious state or not are excluded from the minimally sufficient conditions for state consciousness. Importantly, as we have just reviewed, Searle considers that particular *content* NCCs are not difference-makers for the contrastive target concerning state consciousness or not – they are to be factored out of the NCC for state consciousness. A further background assumption for this approach is therefore that, for the purposes of answering its contrastive question, one should target the property of being in some overall conscious state whilst excluding any contribution from particular conscious contents.

There seems to be two ways to conceptualise this approach. On the first conceptualisation, the idea is to intervene on a creature's overall conscious state in conditions where there is *no* conscious content at all. This could be what Searle has in mind, given his insistence on the conscious field being “pre-existing”.¹ I shall not discuss this conceptualisation much. I myself doubt that such a phenomenon as an utterly empty conscious field even exists (the best bet may be the kinds of states reported by master meditators); and even if it did exist it seems an excessively difficult topic to investigate (and some also have problems with as much as forming a clear conception of this kind of state, see Dainton, 2000, pp. 53–55). In fact, most of the research in this area operates with the other conceptualisation of this approach according to which the contribution of content to the overall conscious state is maintained but kept constant across conditions, whilst the overall conscious state is intervened on.

This second conceptualisation is evident in the kinds of studies that would most naturally be taken as examples of the state-based NCC search. In Laureys et al. (2000) the contrast is between VS patients before and after recovery, where both groups are subjected to noxious and auditory stimuli. Similarly, in Merker's (2007) discussion of hydranencephaly, the presence of being in a conscious state is tested by inducing an emotional content (joy about a baby brother) that elicits a behavioural response we associate with being conscious (expressing joy) in normal subjects too. In Schiff et al., 2007, the presence of consciousness is measured by behavioural response to verbal, contentful prompts. In Owen et al. (2006) a VS patient was asked to imagine a tennis game and it was found that activity corresponds well to the brain activity of healthy conscious volunteers under similar stimulus conditions. In none of these cases can it be concluded that there was an utterly empty overall conscious state since particular contents were present across conditions (e.g., thinking about a baby brother, responding to a verbal prompt, imagining a tennis game). I shall therefore examine the problems that the state-based NCC approach faces under this conceptualisation.

4.1. Controlling for content and losing sight of unconscious states

Consider first Laureys et al. (2000, see also Laureys et al., 1999). This study is clearly distinct from the content-based approach because it did not presuppose that the subjects in question were already conscious. It compared a VS patient before

¹ Consider also Block's interpretation of Merker's (2007) study that the children with hydranencephaly have “the conscious field with little or nothing in the way of any conscious contents: that is, they have the total without much in the way of core [NCC]” (Block 2007).

and after recovery and with healthy controls. A distinct corticothalamic pattern of interconnectivity was present in the recovered patient and healthy controls but not in the VS patient. This pattern of activity is therefore a candidate for inclusion in the state NCC.

Subjects were subjected to auditory (clicks at 95 Hz) and somatosensory (electrical impulse to the wrist) stimuli both in the VS condition and the non-VS condition. We can see this as an attempt to control for the contribution of representational content to consciousness so that the revealed activity would pertain to the conscious–unconscious contrast rather than a contrast between processing of different contents in already conscious creatures. This is therefore an impressive attempt to reveal the state NCC.

Nevertheless, there is some reason to think that the study falls short of revealing minimally sufficient conditions for state consciousness. It seems likely that there are differences in the neural processing of representational content between VS and non-VS states, which has little to do directly with consciousness. In particular, when the subject is conscious, content may be more globally available to cognitive consuming systems throughout the brain (e.g., one may be planning how to avoid the electrical stimulus to the wrist where such planning has more to do with processing of this specific content than with being in an overall conscious state). So, some of the activity seen in the non-VS condition presumably stems from additional processing of the presented stimuli rather than from activity subserving state consciousness as such. The revealed activity is therefore falling short of giving the state NCC in the sense of its minimally sufficient conditions. These considerations are parallel to the considerations that led to binocular rivalry becoming an attractive paradigm on the content-based approach.

The question is then how such studies can be designed to get closer to the state NCC. The problem I point to is that as one passes from a VS to a non-VS state plausibly one gains both state consciousness *and* more global content processing capabilities. One possible strategy is to focus on unconscious states in which content processing is likely to be more closely matched with what content processing is like in the normal conscious state: candidates could be absence seizures, somnambulism, and perhaps the minimally conscious state.

However, at this point a complexity must be noted: as content is more carefully matched, it becomes harder to uphold the contrast between conscious and unconscious states, which is crucial for the state-based approach. This is dramatically illustrated in the literature. In Owen, Coleman et al.'s recent fMRI study of a VS patient (2006, see also Owen, 2008), content is kept constant in the sense that a VS patient is compared with healthy controls whilst asked to imagine playing tennis. The patterns of brain activity (e.g., in the supplementary motor area, SMA) are very similar across subjects, and the authors argue that this patient is therefore likely to be in a conscious state in spite of being diagnosed as vegetative.

This conclusion is obviously at odds with the Laureys et al. studies. One possibility is that the tennis imagery study shows that VS patients in general, and probably including those in the Laureys et al. studies, are in fact conscious. This would undermine our confidence that the thalamocortical activity in the Laureys et al. studies pertains to the contrast between conscious state vs. unconscious state after all. It could not then be argued that it forms part of the state NCC.

One way to avoid this conclusion is to hypothesise that the VS patient in the Owen et al. study is somewhat atypical and has the thalamocortical activity in contrast to the VS patient in the Laureys et al. study. It goes with this hypothesis that patients without the corticothalamic activity should also lack the SMA activity induced by imagining tennis. These are testable hypotheses that if confirmed would support the belief that the tennis-imagining patient was in fact conscious.

Whatever the outcome of this situation, it highlights a complicating lesson for the state-based approach to the NCC. The lesson is that the better content is matched across conditions, the less certain we become that the conditions contrast the unconscious state with the conscious state rather than merely contrasting conscious states with conscious states. I argued above that in order to arrive at the minimally sufficient conditions for state consciousness, processing of representational content needed to be carefully matched across the conscious–unconscious state conditions. Now we see that if content processing is in fact matched (such as processing of tennis imagery which is in fact very similar across the states), then we cease to strongly believe that the experimental conditions still manages to contrast being in a conscious state with being in an unconscious state. The reason is, of course, that we normally take processing of content to a high level as a sign of being in a conscious state. But, and this is the Searlian lesson of the content-based approach, whether a content is conscious or not fails to engage the distinction between overall states of consciousness.

It therefore seems that, when the aim of uncovering the state NCC is kept firmly in focus, attempts at better controlling experimental designs to reveal the minimally sufficient conditions for state consciousness has a tendency to make the distinction between conscious and unconscious states slip away. It thus becomes unclear how the state-based approach to the NCC can be pursued methodologically.

4.2. An independent measure of state consciousness?

It might be thought that this kind of problem to some degree can be evaded if the presence or absence of conscious states could be gauged *independently* of attributions of content to the subject in question. Thus, if only we had an independent measure of conscious state, then we could pick for our state NCC studies the subjects that have the most content processing yet are still unconscious. This proposal comes up against the tough, classic philosophical Other Minds problem – captured well by Chalmers' early point that if only we had a "consciousness-meter" the NCC project would not have to rely on first-person methodologies (Chalmers, 1998).

On the one hand, if a subject does not display any of the behaviour that we normally associate with being conscious, it cannot easily be argued that there is an *absence* of consciousness. This is why researchers like Owen et al. investigate VS

patients for consciousness, and why it was important to discover the class of locked-in patients. The researchers in this field of study are also very clear that the absence of thalamocortical activity, or of SMA activity in response to requests to imagine tennis, cannot be taken as evidence of absence of consciousness (see Laureys, 2005; Owen, 2008). Even if those activity patterns are indicative of being conscious, alternative patterns of neural activity may also be. The case of hydranencephaly is a case in point because there is no cortex or SMA to activate and yet patients may be conscious (Merker, 2007). This is neatly enshrined in the definition of the NCC, which avoids the notion of 'necessary conditions' in favour of 'minimally sufficient conditions', thus allowing for differing realisations of the same phenomenon.

On the other hand, neither is it easy to argue for the *presence* of an overall conscious state given a certain pattern of neural activity. Consider for example Owen et al.'s tennis imagery study, where it is argued that their VS patient is in a conscious state, given the observed SMA activity. There is a criticism of their argument, which can be answered but at a cost to those interested in pursuing the state-based approach.

The criticism is that the brain activity observed in the VS patients is just an *automatic* response that the patients could have even though they are not conscious. Owen (2008) makes a strong case that the response is voluntary and not merely automatic because the activity levels are sustained for very long periods of time (approx 30 s) and change upon requests to rest or imagine something else. This is not like the automatically evoked activity normally seen, which is more transient.

This response is compelling, though it is slightly puzzling that this logic removes the need for comparison with a control group: consideration solely of the VS patient's SMA activity is enough since any pattern of activity anywhere in the brain that is time-locked to the requests could demonstrate the voluntariness.² More importantly, with Owen's argument in favour of state consciousness, an important shift of focus is introduced. Now the mark of consciousness is *voluntary control* rather than what philosophers refer to as *phenomenality* (e.g., what it is like to experience pain or imagine tennis or to be conscious at all). Volition might very well be a mark of consciousness – and Owen is certainly right that it carries extremely important clinical and ethical implications – but in respect of the NCC debate it changes the topic. The aim of the state NCC search is to find the minimally sufficient neural conditions of what it is like to be in a conscious state, not of whether a subject is capable of voluntary action. Finding the neural correlates of volition, the NCV, is a valid and important but different project from finding the NCC.

To illustrate the problem, consider what would be involved in using the tennis imagery study in the state-based NCC approach. It is defining of this approach that it contrasts conscious with unconscious subjects. This rules out comparisons with healthy, conscious subjects since the crucial point is that the VS patient is also conscious. To operationalise the contrast one would presumably need to identify the significant neural differences between VS patients that display the tennis imagery response and patients that do not display this response. These could perhaps be VS patients in a worse condition, or patients in coma. It could then be argued that the activity that is spared in the responsive patients constitutes part of the state NCC. But this would be turning a blind eye to the fact that this neural activity would be most relevant to whether the patient is performing a voluntary action, rather than whether there is something it is like for the patient to do this. That is, it may be that the non-responsive patient is phenomenally conscious but unable to act voluntarily.

There is however another line of argument, which does not rely on a notion of volition, to the conclusion that the VS patient is *phenomenally* conscious. This argument begins with considering the brain response in the healthy, conscious control subjects to a request to imagine tennis. It is very likely that this pattern of brain activity includes, amongst other things, the sufficient conditions for *consciously* imagining tennis. Because this brain activity is so similar to that of the patient when given the same instruction, this argument concludes that the patient is also consciously (in the phenomenal sense) imagining tennis being played.

This argument is not without strength, though the study in question for practical reasons only looked at overlap of activity in SMA and other task-specific areas, which are not likely to be parts of the state NCC. But, even with a broader focus, it appeals centrally to the possession of the *content* NCC as a mark of state consciousness. In effect the argument says that, if the VS patient has the neural substrate of the conscious content of imagined tennis playing, then she is conscious. Of course this only follows if the neural activity includes the *minimally* sufficient conditions for such a conscious content. So identification of the content NCC is essential to running this argument. Section 3 showed that the content-based approach presupposes that the subjects in question are already in an overall conscious state and it therefore follows that it cannot be used to establish independently whether a subject is in fact conscious. This line of argument therefore would fail in this context.

These complexities in the state-based NCC approach show that the most obvious candidates for being a study of state consciousness tends to miss the crucial contrast between conscious and unconscious states as it attempts to be better controlled, and that attempts to resolve this, by finding independent measures of conscious state, risk turning the study away from phenomenal consciousness and towards issues of voluntariness, or returns to the problematic content-based approach.

² I think a sceptic could still persist this line of argument by referring to recent studies of apparent voluntary action, which reveals sustained unconscious automaticity in response to requests of various sorts (Bargh & Chartrand, 1999; Wegner, 2002). The sceptic could argue that since it is not known how mechanisms underlying this kind of automaticity are affected in VS patients, it could be that their seemingly voluntary responses are wholly automatic (it might be that there is more automaticity the less conscious control there is; this sceptic could be sidelined if the brain response itself could be used as a communicative system but this did not seem to happen in the study in question).

5. Discussion

The argument has been that the *content*-based approach to identifying the NCC presupposes consciousness in a methodologically compromised manner, and that it may in fact fail to be a methodology uniquely suited to reveal the content NCC, at least when the focus is on what it takes for contents to *be* conscious rather than for contents to be *selected* for consciousness. And it has been argued that the *state*-based approach risks being over-inclusive in a way that is difficult to resolve without losing sight of the crucial contrast between conscious and unconscious states, or changing the topic from phenomenal consciousness to volition, or collapsing to the content-based approach.

It seems then that neither the content- nor the state-based approach taken in isolation from one another will help us discover the NCC. Though the present arguments by no means show that an NCC can never be found, they do show that the current paradigms are in need of re-evaluation and re-development in the light of methodological and conceptual considerations.³

A natural proposal is then that the two approaches should be pursued in tandem such that studies somehow target conscious contents and conscious states simultaneously. A number of theories of consciousness seem consistent with a view of consciousness on which content and state go together (e.g., Baars, 1997, 2002; Dennett, 1991; Kinsbourne, 1988; Kriegel, 2007; Rosenthal, 1997; Tononi, 2005; Tye, 1995; as well as representational vehicle views such as O'Brien & Opie, 1999). The arguments presented here are neutral on these theories but do indeed suggest that the current NCC methodologies taken separately most likely cannot gain much direct evidence for or against them. For example, on a global workspace theory (e.g., Baars, 1997), one could hold that being state conscious is having at least one content in the global workspace. This is not an implausible view but it does not sit well with pursuing the content-based NCC approach. On that approach, as we have seen, it is presupposed that the creature is already in a conscious state, so on the global workspace theory it must be presupposed that there is already a content in the global workspace. But then what studies such as Tong et al.'s (1998) binocular rivalry study can reveal are the neural structures responsible for selecting one or other new content for the workspace, not what it takes for the original content to be conscious in the first place.

This suggests that we should not respond to these problems for the NCC approaches by simply conjoining the various neural structures that the different approaches reveal and nominate that as the “joint NCC” or some such. This would perhaps seem attractive because it would allow a neat division of experimental labour where different groups of researchers can stick with either the content- or the state-based approach and thus concentrate on different aspects of the NCC and simply conjoin their findings in the end.⁴ But given the critical examination of the NCC approaches above, it is difficult to see how this simple measure could overcome their problems: (i) that the content-based approach in practice can proceed without any assumptions about whether the creature is in an overall conscious state or not, and (ii) that the state-based approach is over-inclusive when taken in isolation from issues of content and difficult to stabilise when content is then taken into consideration.

It therefore appears likely that further progress in the search for the neural correlates of consciousness requires that something new be brought to the study of consciousness. Rather than merely conjoining the approaches, it may be that they must be integrated in a new type of experimental approach that targets the presumably causal, mechanistic interplay (cf. Bechtel 2007) between content processing and overall conscious state across different contents and across different types of conscious and unconscious states.

Acknowledgments

Thanks to audiences at ANU, Melbourne University's Brain and Mind Club, AAP08, and ASSC12. Special thanks to Tim Bayne. This research was supported by Australian Research Council grant DP0984572.

Appendix A

A.1. Terminological and conceptual distinctions in the debate about consciousness

There are a number of further terminological and substantial distinctions that relate to the discussion. Since these often influence one's interpretation of the issues I include an appendix on them here.

³ Bayne, in a recent and congenial paper (Bayne, 2007), argues that our conception of consciousness, in particular concerning the distinction between conscious creatures and conscious states in various ways constrain our attempts to explain consciousness. My approach here is more driven by methodological considerations.

⁴ Baars advocates a combination of creature and state consciousness. He argues that deep structures in the brain including the intralaminar nucleus (see also Bogen, 1995; Kinsbourne, 1995) are necessary but not sufficient for consciousness, and that cortical areas may likewise be necessary but not sufficient for consciousness (Baars, 1988; Baars, 1995). Churchland (1995, see also 2007) advocates a similar line. The first neural structures would be related primarily to state consciousness and the second to content consciousness. So Baars suggests that we cannot really have either of these properties, state or creature consciousness, on their own. We must expect consciousness to arise when both kinds of structures are activated and, presumably, interacting. Baars doesn't argue this on conceptual or methodological grounds, as I have above. He cites, e.g., blindsight studies as evidence that cortical state NCC is necessary for consciousness and uses this to show that the brainstem creature NCC is not sufficient.

The notion of having a conscious content is relatively intuitive and is often explained in terms of representational content: it is for example what happens when someone consciously perceives a house. Philosophers find it natural to speak here of being in a conscious ‘state with a certain content’. But neuroscientists have appropriated the notion of ‘state’ for the creature’s overall level of consciousness. For that reason I shall use ‘content’ rather than follow philosophical usage. Philosophers have tended to use the term ‘creature consciousness’ roughly for what the neuroscientists call ‘state of consciousness’ though again I shall use the notion ‘state’ rather than ‘creature’.

It is common to see philosophers use the term ‘creature consciousness’ simply for the property of being awake or sentient (Dretske, 1993; Rosenthal, 1986, 1997, pp. 729–730). Rosenthal intends this to contrast with ‘state consciousness’ (i.e., what the neuroscientists call ‘content consciousness’), which is a property of a mental state, on the understanding that some mental states can be unconscious. This is similar to those neuroscientists who argue that being in a conscious ‘state’ is a matter of being awake, sentient or aroused (e.g., Baars, 1995). However, this characterisation of state consciousness is not satisfactory. In one sense, *being awake*, *sentient* or *aroused* are properties that have nothing much to do with being conscious at all (for example, you can be awake and unconscious, as in epileptic absence seizures; or asleep and conscious, as in dreaming). In another sense, *being awake* or *sentient* or *aroused* is just the same as *being conscious* in which case there is no real distinction indicated by appealing to them (awakening is in fact a complex process, which is often explained in terms of consciousness, see, e.g., Balkin et al., 2002).⁵

There is therefore a need for a term like ‘state consciousness’ that is distinct from the ‘awake state’ or ‘arousal’. It is, however, difficult to express succinctly the sense in which ‘state consciousness’ should then be used – a fact which is a core motivation for writing this paper. To begin, the notion can be operationally defined simply as the kind of property that is, as Searle says, presupposed in studies of content NCCs, that is, in those studies it is presupposed that the creature is in an overall conscious state (so for example, we wouldn’t perform such a content NCC study of epileptic patients during an absence seizure, and we would perform them on dreaming subjects). This notion of state consciousness makes enough good sense for numerous empirical studies to have been conducted on its basis. In the main text I argue that things are not so simple after all.

Most discuss the Koch-like content NCC search in terms of ‘content’ rather than ‘states’ (Baars, 1995; Bogen, 1995; Chalmers, 2000; Noë & Thompson, 2004). Not much hangs on this: for practical purposes either term would do since most NCC research is focused on states with representational content. However, it would be most satisfying if the term ‘state’ were in fact dominant because the scientific methodology seems neutral on whether there is in fact a mental state with a representational content. For example, it seems reasonable (practical issues notwithstanding) to execute the content NCC approach on a conscious sensation such as orgasm, even if most people would say that though orgasm is a conscious sensation it has no representational content.

Searle’s use of ‘unified conscious field’ is supposedly close to what I talk about as ‘state consciousness’ and what others talk about as ‘creature consciousness’ but as argued in the main text of this paper, the use that Searle makes of his notion is not plausible. In any case, the intuitive content of the notion of state or creature consciousness, or unified conscious field, is captured by the idea that it is the background condition in studies of conscious contents.

There is another piece of terminology that is often used. It is natural to think that if we find the content NCC for a given conscious content, then we will not be able to reproduce that state in this tissue on its own – cut out of the brain and stimulated appropriately in a jar. So it is common to say (following usage in Shoemaker, 1981) that the content NCC is the “core” NCC and the remainder belong to the background conditions, or enabling factors, that are needed for the tissue to work in the right way but is explanatorily redundant for the production of the conscious content. In fact, this harks back to the debate between Searle and Koch but the terminology is somewhat misleading. The problem is that (i) it immediately relegates the non-core NCC (what I have called state NCC) to the backseat as explanatorily redundant for the production of consciousness; it treats all the non-core factors as on a par with blood supply. But this is precisely what is at issue in the debate in question: Searle insists that it is not redundant. It is in fact possible to distinguish between non-core causal factors and non-core constitutive factors (cf. Block, 2007). This improves matters a bit but immediately returns us to the very problem highlighted in the debate between Searle and Koch for it is difficult to see how some factors can be constitutive yet non-core. (ii) The core/non-core terminology masks the fact that the state NCC just as much needs the distinction between core and non-core: if the state NCC is found, then it will not produce creature consciousness either, if cut out and put in a jar. It too needs enabling conditions, this just follows from treating the search for state consciousness as a general search for the neural correlates of something.

Ned Block has championed a distinction between ‘cognitive accessibility’ and ‘phenomenally conscious states’ (Block, 1995, 2005). He argues that we should allow the possibility of conscious states that are not cognitively accessible (or, lately, that cognitive accessibility is not part of the constitutive conditions of phenomenality (Block, 2007)). The problem I discuss here is orthogonal to this debate. I view the supposed distinction between content and state consciousness to be falling under the notion of phenomenality. Of course, it may be objected that it scarcely makes sense to think of someone being unaware of being in an overall conscious state. However that may be, part of my overall conclusion here is that we should not think of state consciousness as a property distinct from the property of being content conscious. So to the extent Block’s distinction applies to conscious contents it will apply to state consciousness.

⁵ Rosenthal and others have also discussed the extent to which creature and state consciousness (to the neuroscientists: state and content consciousness) may depend on each other, and whether they are distinct or not (Manson, 2000; McBride, 1999). I side with those that argue that they are not distinct.

Finally, many will associate Searle's talk of a 'unified conscious field' with debates about the *unity* of consciousness, as well as the spatial and temporal binding of conscious representational content (e.g., Noë & Thompson, 2004). I think it is true that Searle's notion has these connotations but they will not be the focus in this paper. From his quotes given in Section 2 it is clear that Searle claims the unified conscious field as something a subject could have independently of having any particular conscious state, and thus without there being any contents that are bound or unified. It is in order to capture this aspect of the debate that I prefer to use the notion of an overall 'conscious state'.

References

- Alkire, M. T., Miller, J., & Laureys, S. (2005). *General anesthesia and the neural correlates of consciousness*. *Progress in brain research* (Vol. 150). Elsevier (pp. 229–244, 596–597).
- Baars, B. J. (1988). *A cognitive theory of consciousness*. Cambridge University Press.
- Baars, B. J. (1995). Tutorial commentary: Surprisingly small subcortical structures are needed for the state of waking consciousness, while cortical projection areas seem to provide perceptual contents of consciousness. *Consciousness and Cognition*, 4(2), 159–162.
- Baars, B. J. (1997). *In the theater of consciousness: The workspace of the mind*. Oxford: Oxford University Press.
- Baars, B. J. (2002). The conscious access hypothesis: Origins and recent evidence. *Trends in Cognitive Sciences*, 6, 47.
- Balkin, T. J., Braun, A. R., Wesensten, N. J., Jeffries, K., Varga, M., & Baldwin, P. (2002). The process of awakening: A PET study of regional brain activity patterns mediating the re-establishment of alertness and consciousness. *Brain*, 125(10), 2308–2319.
- Bargh, J., & Chartrand, T. L. (1999). The unbearable automaticity of being. *American Psychologist*, 54, 462–479.
- Bayne, T. (2007). Conscious states and conscious creatures: Explanation in the scientific study of consciousness. *Philosophical Perspectives*, 21(1), 1–22.
- Bechtel, W. (2007). *Mental mechanisms: Philosophical perspectives on cognitive neuroscience*. NY: Routledge.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18, 227–287.
- Block, N. (2005). Two neural correlates of consciousness. *Trends in Cognitive Sciences*, 9(2), 46.
- Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, 30(5–6), 481–499.
- Bogen, J. E. (1995). On the neurophysiology of consciousness: 1. An Overview. *Consciousness and Cognition*, 4(1), 52–62.
- Chalmers, D. (1998). On the search for the neural correlate of consciousness. In S. Hameroff, A. Kaszniak, & A. Scott (Eds.), *Toward a science of consciousness II: The second Tucson discussions and debates*. Cambridge, Massachusetts: MIT Press.
- Chalmers, D. (2000). What is a neural correlate of consciousness? In T. Metzinger (Ed.), *Neural correlates of consciousness: Empirical and conceptual issues*. Cambridge, Massachusetts: MIT Press.
- Churchland, P. (1995). *The engine of reason, the seat of the soul*. Cambridge, Massachusetts: MIT Press.
- Churchland, P. (2007). *Neurophilosophy at work*. Cambridge: Cambridge University Press.
- Dainton, B. (2000). *Stream of consciousness: Unity and continuity in conscious experience*. NY: Routledge.
- Dennett, D. C. (1991). *Consciousness explained*. Little, Brown & Co.
- Dretske, F. I. (1993). Conscious experience. *Mind*, 102(406), 263–283.
- Frith, C., Perry, R., & Lumer, E. (1999). The neural correlates of conscious experience: An experimental framework. *Trends in Cognitive Sciences*, 3, 105–114.
- Haynes, J. D., Deichmann, R., & Rees, G. (2005). Eye-specific effects of binocular rivalry in the human lateral geniculate nucleus. *Nature*, 438, 496.
- Hohwy, J. (2007). The search for neural correlates of consciousness. *Philosophy Compass*, 2(3), 461–474.
- Kinsbourne, M. (1988). Integrated field theory of consciousness. In A. J. Marcel & E. Bisiach (Eds.), *Consciousness in contemporary science* (pp. 239–256). Oxford: Clarendon Press.
- Kinsbourne, M. (1995). The intralaminar thalamic nuclei: Subjectivity pumps or attention-action coordinators? *Consciousness and Cognition*, 4(2), 167–171.
- Koch, C. (2004). *The quest for consciousness: A neurobiological approach*. Englewood, Colorado, Robert and Company Publishers.
- Kriegel, U. (2007). A cross-order integration hypothesis for the neural correlate of consciousness. *Consciousness and Cognition*, 16(4), 897–912.
- Laureys, S. (1999). Impaired effective cortical connectivity in vegetative state: Preliminary investigation using PET. *Neuroimage*, 9, 377–382.
- Laureys, S. (2005). The neural correlate of (un)awareness: Lessons from the vegetative state. *Trends in Cognitive Sciences*, 9(12), 556–559.
- Laureys, S., & Boly, M. (2007). What is it like to be vegetative or minimally conscious? *Current Opinion in Neurology*, 20, 609–613.
- Laureys, S., Faymonville, M. E., Luxen, A., Lamy, M., Franck, G., & Maquet, P. (2000). Restoration of thalamocortical connectivity after recovery from persistent vegetative state. *Lancet*, 355, 1790–1791.
- Lipton, P. (2004). *Inference to the best explanation*. London: Routledge.
- Manson, N. (2000). State consciousness and creature consciousness: A real distinction. *Philosophical Psychology*, 13(3), 405–410.
- McBride, R. (1999). Consciousness and the state/transitive/creature distinction. *Philosophical Psychology*, 12(2), 181–196.
- Merker, B. (2007). Consciousness without a cerebral cortex: A challenge for neuroscience and medicine. *Behavioral and Brain Sciences*, 30(1), 63–81.
- Noë, A., & Thompson, E. (2004). Are there neural correlates of consciousness? *Journal of Consciousness Studies*, 11(1), 3–28.
- O'Brien, G., & Opie, J. (1999). A connectionist theory of phenomenal experience. *Behavioural and Brain Sciences*, 22, 127–148.
- Owen, A. (2008). Disorders of consciousness. *Annals of the New York Academy of Sciences*, 1124, 225–238.
- Owen, A. M., Coleman, M. R., Boly, M., Davis, M. H., Laureys, S., & Pickard, J. D. (2006). Detecting awareness in the vegetative state. *Science*, 313(5792), 1402.
- Polonsky, A., Blake, R., Braun, J., & Heeger, D. J. (2000). Neuronal activity in human primary visual cortex correlates with perception during binocular rivalry. *Nature Neuroscience*, 3, 1153.
- Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C., & Fried, I. (2005). Invariant visual representation by single neurons in the human brain. *Nature*, 435(7045), 1102–1107.
- Quiroga, R. Q., Reddy, L., Koch, C., & Fried, I. (2007). Decoding visual inputs from multiple neurons in the human temporal lobe. *Journal of Neurophysiology*, 98(4), 1997–2007.
- Rees, G., & Frith, C. (2007). Methodologies for identifying the neural correlates of consciousness. In M. Velmans & S. Schneider (Eds.), *The Blackwell companion to consciousness*. Oxford: Blackwell Publishing.
- Rees, G., Wojciulik, E., Clarke, K., Husain, M., Frith, C., & Driver, J. (2000). Unconscious activation of visual cortex in the damaged right hemisphere of a parietal patient with extinction. *Brain*, 123(8), 1624–1633.
- Rosenthal, D. M. (1986). Two concepts of consciousness. *Philosophical Studies*, 94, 329–359.
- Rosenthal, D. M. (1997). A theory of consciousness. In N. Block, O. Flanagan, & G. Güzeldere (Eds.), *The nature of consciousness: Philosophical debates* (pp. 729–753). Cambridge, Massachusetts: MIT Press.
- Schiff, N. D. (2007). Global disorders of consciousness. In M. Velmans & S. Schneider (Eds.), *The Blackwell companion to consciousness*. Oxford: Blackwell Publishing.
- Schiff, N. D., Giacino, J. T., Kalmar, K., Victor, J. D., Baker, K., & Gerber, M. (2007). Behavioural improvements with thalamic stimulation after severe traumatic brain injury. *Nature*, 448(7153), 600–603.
- Searle, J. R. (2000). Consciousness. *Annual Review of Neuroscience*, 23(1), 557–578.
- Searle, J. R. (2004). *Mind – A brief introduction*. Oxford: Oxford University Press.
- Searle, J. R. (2005). Consciousness: What we still don't know. *The New York Review of Books*, 52(1).
- Shewmon, D. A., Holmes, G. L., & Byrne, P. A. (1999). Consciousness in congenitally decorticate children: Developmental vegetative state as self-fulfilling prophecy. *Developmental Medicine and Child Neurology*, 41, 364.

- Shoemaker, S. (1981). Some varieties of functionalism. *Philosophical Topics*, 12, 93–119.
- Tong, F., & Engel, S. A. (2001). Interocular rivalry revealed in the human cortical blind-spot representation. *Nature*, 411, 195–199.
- Tong, F., Meng, M., & Blake, R. (2006). Neural bases of binocular rivalry. *Trends in Cognitive Sciences*, 10(11), 502.
- Tong, F., Nakayama, K., Thomas Vaughn, J., & Kanwisher, N. (1998). Binocular rivalry and visual awareness in human extrastriate cortex. *Neuron*, 21, 753–759.
- Tononi, G. (2005). Consciousness, information integration, and the brain. In S. Laureys (Ed.), *Progress in brain research* (Vol. 150, pp. 109). Elsevier.
- Tye, M. (1995). *Ten problems of consciousness*. Cambridge, Massachusetts: MIT Press.
- Wegner, D. M. (2002). *The illusion of conscious will*. Cambridge, Massachusetts: MIT Press.