MORTEN OVERGAARD

CONFOUNDING FACTORS IN CONTRASTIVE ANALYSIS

ABSTRACT. Several authors within psychology, neuroscience and philosophy take for granted that standard empirical research techniques are applicable when studying consciousness. In this article, it is discussed whether one of the key methods in cognitive neuroscience – the contrastive analysis – suffers from any serious confounding when applied to the field of consciousness studies; that is to say, if there are any systematic difficulties when studying consciousness with this method that make the results untrustworthy. Through an analysis of theoretical arguments in favour of using contrastive analysis, combined with analyses of empirical findings, I conclude by arguing for three factors that currently are confounding of research using contrastive analysis. These are (1) unconscious processes, (2) introspective reports, and (3) attention.

1. INTRODUCTION

In the field of consciousness studies it has become quite common to consult the neurobiological research techniques of PET, fMRI, EEG, and MEG to learn about the activity of the brain while being in a conscious state. No matter which theoretical position one has in understanding the relation of mind to body, it seems inescapable that knowledge about what happens in the brain while one is conscious must be a part of a full explanation of consciousness. That is to say, while quite a large number of modern theorists would deny that knowledge about the neural correlates of consciousness is sufficient to explain consciousness, few would deny that it seems to be a necessary part (Vaas 1999).

The raw data output of any such method, no matter which of the research tools that are taken in use, must undergo a so-called contrastive analysis in order to be informative to the scientist. If one were to create a PET-image, say, of a person looking at a red triangle, one would find activations all over the brain. This would of course not be satisfactory to neuroscientists, in that they want to know which specific brain properties underlie seeing the red triangle, yet the observed activations represent everything that the brain is

Synthese **141:** 217–231, 2004.

doing at that moment,[1] whereof seeing the red triangle is just one aspect. Therefore, the neuroscientist needs to contrast the condition of seeing the red triangle with a condition in which the subject is in the exact same state *except* that he or she is now not looking at the red triangle. So, contrastive analysis is a matter of having two conditions in which a, say, mental phenomenon is present in one and not in the other. Then, one proceeds to subtract the activations of the last condition (that I shall refer to as the "control condition") from the first (the "experimental condition").[2] The basic reasoning is, clearly, that this will show the neural correlates that are different in the two conditions – that is, it will show the neural correlates of the mental phenomenon in isolation. In this sense, contrastive analysis makes the assumption that mental states can be treated as variables.

The subject could be asked to lie still and try to think of nothing specifically in the control condition. Thus, as the reasoning goes, the neural correlate of seeing the red triangle will be the only thing different in the two conditions. Of course, there will be some further differences in the two conditions in that the subjects unavoidably will think of different things during the experiment, but this is not considered to be confounding of the results. Over the course of several trials and with the use of several subjects, there is general agreement that such differences are eliminated.

Modern cognitive neuroscience is well aware that a number of methodological problems arise when conducting experiments as simple as this one. Notably, Friston et al. (1996) have discussed problems related to the notion of "pure insertion" (the idea that a given mental phenomenon can exist without affecting other mental processes). There has been some discussion about appropriate control conditions as well (which will be discussed below).

These methods are long established in experimental psychology. It has however been argued that they can be applied to consciousness studies as well without problems of any sort. One of the strongest advocates of this view is Bernard Baars who writes, "In the case of consciousness, we can do this [specify its conditions of occurrence and non-occurrence] by carrying out a "contrastive analysis" of comparable conscious and unconscious events" (Baars 1994, 4). Of course, the subject is not completely unconscious in the control condition in an experiment on consciousness. In the perception experiment, the subject would even still be conscious of his visual perceptions when not perceiving the red triangles. So,

first of all, we could distinguish between what the philosopher David Rosenthal refers to as "creature consciousness" and "state consciousness" (1997). Being "creature conscious" is a question of being dead, in a comatose state, or in dreamless sleep contrasted with being awake and having conscious states of some sort. Being state conscious is a question of being in a specific conscious state with a specific content, contrasted with being in a different state, in which one is not conscious of the content present in the first state. Thus, experiments using contrastive analysis will in general concern state consciousness, where a state of being conscious of the red triangle is contrasted with a state of not being conscious of it.

To tell when subjects are conscious or not conscious of something, a scientist would generally use reports given by the subject. A conscious state is only directly accessible to the one having the conscious state. Therefore, one has to take some kind of behaviour, observable to the third person, as an expression of the conscious state in question. Elsewhere, I have argued in more detail that the strongest possible experiment on consciousness would be one that used explicit reports about conscious states as part of its methodology (Overgaard 2003). That is, when a subject is giving the report that right now he is having an experience of the colour blue, one would have the burden of explanation if one were to argue against him. However, if the scientist were to say that the subject is visually conscious every time he is able to perform correctly in a visual discrimination task, the scientist would still have the burden of explanation to justify why it is so. That is, being conscious is a logical part of a causal explanation of the utterance "I am conscious of X", whereas being conscious is not a logical part of a causal explanation of the ability to functionally discriminate between two objects. This is apparent in the entire research field of unconscious perceptions, memories, etc. that is based upon experiments in which it is investigated what people can do (for instance, if they can discriminate) while reporting to be unconscious. So an apt methodology for the neuroscientific study of consciousness, I think many scientists in the field would agree, is an experiment in which the neural activations in a subject who gives a clear report about his conscious state are contrasted with the neural activations when the subject is not conscious. The ideal experiment would of course also imply even more developed neuroscientific tools to get precise representations of the working of the brain.

A number of authors have argued, mostly from a theoretical position, that with the use of contrastive analysis, one can achieve a scientific approach to consciousness with the possibility of finding quite precise neural correlates of consciousness within its reach (Baars 1994; Frith et al. 1999). Clearly, this is a very straightforward approach that potentially could have extremely important consequences, when contemplating the impact on any theory of consciousness, such a discovery would have.

For this reason alone, I find it necessary and highly timely for scientists to look into the methodology of contrastive analysis to search for any possible confounding factors. A confounding factor, in this context, would be any *systematic* difference between the experimental (conscious) and control (non-conscious) conditions other than the consciousness itself. This would be confounding in that one would not have succeeded in isolating consciousness as a variable. Just as in the case of having no control condition at all, the scientist would in such a case not know which activations that were due to consciousness and which that were due to "the something else". In cases where the scientist is aware of this problem, he would simply find himself unable to point out which neural activations that are correlates of consciousness and which that are correlates of "something else". On the other hand, if the scientist were unaware of confounding factors, he would simply give false claims about which neural substrates that subserve consciousness. To illustrate the theoretical impact of this problem, one might imagine a scholar in consciousness studies who argued in favour of a position stating that there is no neural correlate of consciousness at all. He would be perfectly able to explain the data yielded from a confounded experiment, in that he could argue that all the brain activations we find, are correlates of the "something else" and not of consciousness. If one could even defend such an extreme position based on confounded data, the relevance of avoiding confoundedness should be very clear.

In the following, I would like to point to three such possible confounding factors: (1) differences in unconscious processes between the two conditions, (2) differences in introspective attitudes, and (3) differences in attention processes.

2. UNCONSCIOUS PROCESSES

In order to make a direct comparison between a conscious and an unconscious state to find the neural substrate of consciousness, one

must, as mentioned, be able to argue that the only thing different in the two states is the consciousness. So, clearly, it will not be sufficient to contrast a state in which subjects have a conscious perception with a state in which they have no perception. There will definitely be a lot of unconscious processing going on when having the perception that will not be present when just resting. So, how is one to know which activations, yielded by the experiment, are correlates of the conscious perception and which are correlates of the unconscious processes? The solution to this problem has been to have a "resting state" that is "more in resemblance" of the active state. Let us take a look at a few suggestions given by Bernard Baars (1994, 5) of which mental phenomena to contrast with a conscious perception.

1. Stimuli lacking in intensity or duration, and masked stimuli.
2. Pre-perceptual stimulus processes.
3. Post-perceptual (habituated) processing.
4. Unaccessed interpretations of ambiguous stimuli.
5. Contextual constraints on the interpretation of conscious percepts.
6. Expectations of specific stimuli

All of these suggestions indicate that one should use *perceptual* states as contrasting conditions in which there however is no *conscious* perception. Would this be sufficient? Obviously, it would only be sufficient if one would be able to argue that there are no unconscious processes different when one is, say, watching a stimulus for 10 ms (and thus have no conscious perception of it) and, say, for 300 ms (thus having a quite clear conscious perception of it). Can such an assumption be defended?

Evidence points in the opposite direction: the cognitive processes underlying experiencing a good, steady image of something and being exposed to a very short glimpse are quite different (Merikle and Daneman 1998). Blindsight patients, for instance, increase their performance based on longer stimulus exposure, even though they do not report changes in experience (Weiskrantz 1986). Also, in normal subjects, there is evidence that subjects can have increased performance as a function of stimulus duration time below subjective threshold (Kolb and Braun 1995). All in all, I believe there is a good amount of experimental evidence to say that there is a change in unconscious processes when varying stimulus duration. In other words, it is not so that consciousness is "just added" when subjective threshold is achieved. Thus, the argument is in accordance with Friston et al.'s critique of "pure insertion" (1996).

Some of the experiments that lately have received the most interest in regards to their relevance for consciousness studies have been the works on binocular rivalry. In binocular rivalry, dissimilar images are shown to the two eyes. Because the images cannot be fused by the visual system, conscious perception alternates between each monocular view. Thus, the neural correlates of the non-conscious view can be used as a control condition for the conscious view. Using the technique coupled with MEG, Tononi et al. find that the neural responses associated with the "conscious condition" "are not global but distributed to a subset of brain regions" (1998, 3202). The problem of confounding influences from unconscious processes applies to the Tononi et al. experiment, described above. In order to claim that the data displays the neural differences between conscious and unconscious vision only, one must make a further claim. That is, one must claim that there are no unconscious cognitive differences between the neural system when it is in the state corresponding to seeing the stimulus and the system when in the state corresponding to a suppressing of the stimulus (when seeing the other stimulus).

In another study, we find an almost opposite example of how unconscious processes can account for a confounding factor. Nicholas Schiff and his co-workers (1999) have described the case of a brain-injured patient, LR, who has been in a vegetative state in 20 years – without any awareness of her self or of the environment. However, sometimes she utters words in both English and Spanish without any external prompting. LR suffered from three successive hemorrhages, leaving only few, isolated islands of LR's brain relatively unharmed, some of which were Broca's and Wernicke's area in the left hemisphere. They have long been known as the neural basis of the articulation of words and the understanding of words, respectively. But even those areas, though less affected than the rest of the brain, had a metabolism of 66% of the normal value in the case of Wernicke's area, and only 50% for Broca's area. Furthermore, Schiff and colleagues describe the word utterances of LR as results of the isolated fragments of her brain alone, since she shows no evidence of having any conscious mentality (in that she does not respond to external cueing). Schiff et al. (1999) go much further than saying that she is unconscious, arguing that LR responds based on brain processes alone, without any kind of mental processes present. However, an alternative explanation of LR's utterances saying that unconscious mentality is the real cause of them seems just as fit.

Based on this, one could argue that it is impossible to find a control condition in which the only thing different from the active state is the "consciousness part". I find that such a position is a quite strong one. At least, it seems that if one were to say otherwise, one would have to believe that consciousness is something that is "just added" to a cognitive process.

Is there any way around this problem? One could argue that if one made a large number of experiments on a number of different mental phenomena – comparing conscious thoughts, feelings, perceptions etc. to unconscious ones – one could look for a single brain property that is only present in the experimental (conscious) conditions across all these different cases. It is however also an unverified ontological assumption that such a brain property should exist. If it turns out to be the case that there is one neural correlate of conscious perception, an entirely different correlate of conscious thoughts, etc., we might not find a method to deal with the problem of unconscious processes in contrastive analysis. So, *given* that there is one neural correlate of consciousness for all or a large number of states of consciousness, contrastive analysis might find a solution to the problem by comparing the contrasted data from experiments on several different mental phenomena.

3. INTROSPECTIVE REPORTS

If one is to accept that the best methodology in an experiment on consciousness will make use of subjective reports explicitly about the conscious state in question, one must introduce the way in which one has access to the conscious state as a factor in the experiment. That is, the brain activity involved in giving reports about consciousness would count as a confounding factor.

One could also argue that not only the report, but also the very introspective attitude in it self could be a confounding factor. At least, many philosophers of mind would use a conceptual distinction between being conscious and being introspective. There is a (conceptual) difference between being in a phenomenal state, and being in a state in which one attends to one's own conscious state – that is, being aware of one's being in a conscious state.

Some philosophers of mind believe that the difference is solely conceptual (e.g., Dretske 1995). If this is correct, introspection would hardly in itself account for a confounding factor. However, there is

experimental evidence that shows otherwise. In a recent experiment, it is shown that quite dramatic differences can be measured between conscious and introspective conscious states. In a perceptual task, subjects were to tell what was presented on a computer screen and thus attend to the represented object of the perceptual state (Overgaard et al. 2001; Overgaard and Sørensen, in press). In another condition, the same subjects were instructed to attend to their own experiences of the very same presentations, using the same stimuli and the same mode of responding (pointing at tables of different shapes, colours, and locations). That is, one condition where subjects were asked to describe the stimulus was contrasted with another, in which they were to describe their experience of the stimulus. The responses of the subjects were treated as being either ''correct'', ''incorrect'' or ''near correct'' (e.g., when they pointed at the same colour as the one presented, but in a brighter or darker tone). The results showed significant differences between the two conditions regarding the correctness of the subjects' responses: In the ''guessing'' condition, the subjects had a tendency to give more of the ''correct'' and ''incorrect'' responses, while there were more ''near correct'' responses in the introspective condition. When such behavioural differences can be obtained, it would seem a fair hypothesis that there are neurological differences as well.

Discussions of introspection as psychological methodology have re-appeared after it being abandoned as untrustworthy (Nisbett and Wilson 1977). Jack and Roepstorff (2002) have opened the debate, arguing that psychology can reunite with introspection and gain new insights. They argue that there are two kinds of reports often used in introspection. One kind is the so-called first-order reports with which subjects can report about events in consciousness. Another kind is the second-order reports. Second-order reports reflect an awareness of the subject's own conscious states. Also from this hold, it could be further argued that introspective states being different from non-introspective states must have a neural correlate different from that of first-order conscious states.

One experiment where introspective reports may account for a confounding factor is in Zeki and ffytche (1998). They found that the blindsight patient G.Y. and eight normal subjects have increased activations in V5 when given the task of discriminating the direction of moving targets. For G.Y., the targets were presented in his blind hemifield, for the normal subjects, the targets were reduced in luminance until their responses were at the same level of correctness as

G.Y. The subjects were to report whether they made a ''total guess'' about the movement of the stimulus (they were to report so verbally or press the button 1) or to scale their awareness of it (between 2 and 4). The neuroimaging results were obtained by contrasting fMRI scans where awareness of movement is reported to scans where subjects guessed about the movement. It could now be argued in accordance with the definition of introspection above that the two conditions have not only differed in the subjects' awareness of the stimulus but also in introspection: when reporting themselves aware of the stimulus, they must consult their conscious states in order to live up to the instructions. This is however not necessary for guessing. Given that one accepts that introspection must have a neural correlate separate from that of first-order conscious states, and given that one wishes to investigate the neural correlate of first-order states specifically, introspection could be seen as confounding in an experimental design as that of Zeki and ffytche (1998).

Another line in NCC research more based upon the technical possibilities of EEG and TMS has found that consciousness of physical movements is related to premotor and primary motor areas (Haggard and Eimer 1999; Haggard and Magno 1999). These experiments rest upon a replication of the now famous experiments by Benjamin Libet, taken to suggest that conscious experience of voluntary actions is delayed around 500 ms. compared to the onset of a neural response potential in premotor cortex (Libet 1985). In these experiments, the subjects were to watch a round clockface that included a moving dot. At some point they were to initiate a voluntary movement and to note the location of the dot on the clock at that point in time. In this task, the subjects were explicitly asked to monitor their own mental state, which, quite clearly, is an introspective task. As Libet himself writes: ''In the Libet study subjects were asked to report the time of their first *awareness* of the wish (W) or urge to act.'' (Libet 2002, 292). This kind of instruction was replicated in Haggard and Eimer (1999). So one might ask: are the observed relations between the premotor and motor areas and reports of awareness of movements a relation between those areas and conscious experiences or a relation between them and introspection? Likewise, for Libet's results, is it so that conscious experience is delayed 500- ms. compared to the onset of the response potential, or does the experience occurs, say, simultaneously with the response potential while introspection is delayed 500 ms?

As a solution to the problem of introspection, one might try to counterbalance the effect by using introspective reports in the control condition as well as in the experimental condition (as it is done by Tononi et al. (1998). One should however not use this methodology without knowledge of the influence on behaviour and first-order experiences that introspection may have. Some evidence suggests that the preparedness to respond in a specific manner in itself gives rise to behavioural differences. This evidence comes from an important experiment by Marcel (1993), in which subjects were to signal when they detected a light. It was shown that the results were highly dependent on the type of behaviour the subjects would engage in when giving their responses. When subjects used eye blinks as reports, they reported to be consciously aware of a significantly larger percent of the stimulus material compared to when they used finger movements as reports. Verbal reports proved even less precise in this respect. It is not clear how to interpret these results, however they seem to suggest that preparedness to give a (specific kind of) response has an influence on the way in which one perceives the stimulus. That is to show that already in the experimental situation, the subject will have a "preparedness to respond", which, hypothetically at least, will have a neural correlate. We seem to need much more knowledge about the effects of introspection on cognitive, behavioural and conscious phenomena in order to understand the results when using an introspective methodology.

It could be argued that this problem seems manageable with much more ease by using only non-introspective reports when looking for neural correlates of consciousness. However, if we were to use this line of approach, one would have to give up using reports that are explicitly about conscious states as a part of one's methodology. This would simply make people introspect. This would definitely be a serious loss in that nothing seems to be as reliable an indication of conscious states than reports about them given by the subjects having them. As an alternative, one might attempt to find the neural correlate of introspection and subtract this from the neural correlate of being introspectively conscious in order to find the "clean" neural correlate of consciousness. Yet, this view presupposes a specific understanding of introspection, namely that introspection is something that is added to the conscious state and remains a separate property. That is, if one were to believe that being introspective would have a direct influence on what the conscious state is like and

(thus also) on the neural correlate of the conscious state in question, one would not accept such a procedure.

All in all, it seems the best solution to use introspective reports only, and to investigate experimentally which activations introspection generates, so that this can be taken into account.

4. ATTENTION

In experiments on consciousness it seems more than plausible that the subjects are attending to the stimulus that they are supposed to report themselves being conscious or unconscious of. Clearly, the moment you ask subjects if they are conscious of something, they will shift their attention accordingly to investigate if that is the case. In experiments on attention, it seems just as plausible that the subjects have been conscious of the objects to which they were attending.

One could reply to this that consciousness and attention are not intertwined by accident. In fact, consciousness is more or less identical to what we mean by attention. I believe that such a standpoint is confused. Attention refers to a focusing of one's mental processes on a specific object – that is, an object of a perceptual act, the object of one's thoughts, etc. One needs not refer to any phenomenal states to describe attention; one can do just fine with a simple functional description, saying that attention is a process that maximises e.g., the perceptual process (Lamme 2003). Meanwhile, a state of consciousness is a state, one is directly aware of being in. A mental state, there is something it is like to be in, is a conscious state (Nagel 1974). One can also show the differences between attention and consciousness by pointing to dissociation cases. First of all, we all have the phenomenal impression of a centre in our visual field, in which we see objects very clearly and with a richness of details. This is what we are currently attending to visually. At the same time, however, we experience a "fringe" – objects in our visual field that we do not focus on directly. A number of experiments on visual perception have revealed that our feeling of having experiences outside of our central focus is a result of "filling-in" (Mack and Rock 1998). These experiences are in fact not due to conscious sensations but to top-down mechanisms in our visual system that create the illusion that we actually perceive these objects. Then, clearly, the experiences of objects in the fringe of our visual field are experiences of objects that we do not attend to. One could also point to blindsight cases, where the patients attend to

the blind spot in their visual field even though they have no experiences of the contents of this particular area (Weiskrantz 1986; Kentridge et al., 1999; Ward and Jackson 2002).

But even though attention and consciousness are separable, in theory as well as in real life, they seem very hard to separate methodologically. In binocular rivalry experiments, as in the Tononi et al. (1998) study mentioned above, it is clear that attention is involved in the task. It is, however, not clear how. Can the subjects for instance be said to be attending to the suppressed stimulus? And, if so, are they attending in the same way and to the same degree as they are attending to the experienced stimulus?

In another experiment by ffytche, he finds support for a Zeki-inspired modular view of the brain based on patients' reporting of their having conscious hallucinations: conscious visual hallucinations correlated with brain areas specific for their content, though they failed to correlate with activity outside the specialised areas (ffytche et al. 1998; ffytche 2000). ffytche suggests that one can either interpret these results as suggesting that the neural correlate of consciousness is distributed over a large amount of specialised brain units. Alternatively, he suggests that the findings do not represent neural correlates of consciousness, but correlates of unconscious functional units instead.

Here, the subjects are attending to their hallucinations, their onset and offset. This task involves a change in attention compared to the control state: the patients could not have been attending to their hallucinations when they were not present. One obvious argument would now be that it is fairly easy to control for attention as factor, since attention is one of the most thoroughly investigated mental functions. We have a good amount of knowledge about the neural correlates of attention from several experiments, so those could simply be subtracted from the activations found in an experiment on consciousness. This idea would not work, however in that consciousness quite likely could be a confounding factor in many experiments on attention. At least, subjects have most likely been conscious in the experimental condition in those experiments that have supplied us with knowledge about the neural correlates of attention. Thus, a subtracting of the neural correlates of attention (as we know them) from the neural correlates of consciousness would run the risk of filtering out parts of the neural correlates of consciousness.

Another way to solve the problem of attention is to somehow make sure that the subjects are "just as attentive" in the control

condition as in the experimental condition. This is of course quite difficult to control for, because how do we know exactly "how attentive" a subject is? On the other hand, this seems to be the best we can do. As mentioned, it would be almost impossible to reverse the strategy and make sure the subjects were unattentive during the experimental and control conditions. It might prove to be one of the biggest challenges to contrastive analysis to try and separate out the neural correlates of consciousness from the neural correlates of attention.

5. CONCLUSION

In this paper, I have tried to summarise some of the problems that face the scientist who sets out to investigate consciousness with the use of contrastive analysis. Clearly, this list of problems is not exhaustive, but it might point out some of the hardest ones. My account is not pessimistic in the sense that I intend to suggest that the use of contrastive analysis is hopeless. However, it underlines that for methodological reasons, current research within cognitive neuroscience is not yet able to give conclusive evidence about the neural correlates of consciousness.

ACKNOWLEDGEMENTS

I am highly grateful to David J. Chalmers and Petra Stoerig for their comments and ideas. Also thanks to Tom Teasdale for proofreading my English. Morten Overgaard was supported by a grant from The Carlsberg Foundation.

NOTES

1. That is, everything within the reach of the applied methodology.
2. Other techniques such as averaging and alignment are also crucial parts in the creation of PET/fMRI images, and may be said to represent possible sources for confounding as well. These possible confoundings, I shall claim, are however not specific for consciousness studies.

REFERENCES

Baars, B. J.: 1994, 'A Thoroughly Empirical Approach to Consciousness', *Psyche* **1**(6), 1–18.

Dretske, F. I.: 1995, *Naturalizing the Mind*, MIT Press.
ffytche, D. H.: 2000, Imaging Conscious Vision, in T. Metzinger (ed.), *Neural Correlates of Consciousness*, MIT Press.
ffytche, D. H., R. J. Howard, M. J. Brammer, A. David, P. Woodruff and S. Williams: 1998, 'The Anatomy of Conscious Vision: An fMRI Study of Visual Hallucinations', *Nature Neuroscience* **1**, 738–742.
Friston, K. J., C. J. Price, P. Fletcher, C. Moore, R. S. Frackowiak, and R. J. Dolan: 1996, 'The Trouble with Cognitive Subtraction', *Neuroimage* **4**, 97–104.
Frith, C. D., R. Perry and E. Lumer: 1999: 'The Neural Correlates of Conscious Experience: An Experimental Framework', *Trends in Cognitive Sciences* **3** (3), 105–114.
Haggard, P. and M. Eimer: 1999, 'On the Relation Between Brain Potentials and the Awareness of Voluntary Movements', *Experimental Brain Research* **126**, 128–133.
Haggard, P. and E. Magno: 1999, 'Localising Awareness of Action with Transcranial Magnetic Stimulation', *Experimental Brain Research* **127**, 102–107.
Jack, A. and A. Roepstorff: 2002, 'Retrospection and Cognitive Brain Mapping: From Stimulus–Response to Script-Report', *Trends in Cognitive Science,* **6** (8), 333–339.
Kentridge, R. W., C. A. Heywood and L. Weiskrantz: 1999, 'Effects of Temporal Cueing on Residual Visual Discrimination in Blindsight', *Neuropsychologia* **37**, 479–483.
Kolb, F. C. and J. Braun: 1995, 'Blindsight in Normal Observers', *Nature* **377**, 336–338.
Lamme, V. A. F.: 2003, 'Why Visual Attention and Awareness are Different', *Trends in Cognitive Sciences* **7**(1), 12–18.
Libet, B.: 1985, 'Unconscious Cerebral Initiative and the Role of Conscious Will in Voluntary Action', *Behavioral and Brain Sciences* **8**(4), 529–566.
Libet, B.: 2002, 'The Timing of Mental Events: Libet's Experimental Findings and Their Implications', *Consciousness & Cognition* **11**, 291–299.
Mack, A. and I. Rock: 1998, *Inattentional Blindness*, MIT Press.
Marcel, A.: 1993, 'Slippage in the Unity of Consciousness', in R. Bock and J. Marsh (eds), *Experimental and Theoretical Studies of Consciousness*, John Wiley & Sons.
Merikle, P. M. and M. Daneman: 1998, 'Psychological investigations of unconscious perception', *Journal of Consciousness Studies* **5**, 5–18.
Nagel, T.: 1974, 'What Is It Like to Be a Bat?' *The Philosophical Review* **LXXXIII**, 435–451.
Nisbett, R. C. and T. D. Wilson: 1977, 'Telling More Than we Can Know: Verbal Reports on Mental Processes', *Psychological Review* **84**, 231–259.
Overgaard, M.: 2003, 'Theoretical and Experimental Studies of Consciousness', *Ph.D. thesis*, University of Aarhus.
Overgaard, M., O. Kauffmann and T. Z. Ramsøy 2001, 'Consciousness and Introspection', Toward a Science of Consciousness, Consciousness Research Abstracts, *Journal of Consciousness Studies*.
Overgaard, M. and T. A. Sørensen: in press, 'Introspection Distinct from First Order Experiences', *Journal of Consciousness Studies*.
Rosenthal, D. M.: 1997, 'A theory of Consciousness', in N. Block, O. Flanagan and G. Güzeldere (eds), *The Nature of Consciousness: Philosophical Debates*, MIT Press.

Schiff, D., U. Ribary, F. Plum and R. Llinas: 1999, 'Words Without Mind', *Journal of Cognitive Neuroscience* **6**, 650–656.
Tononi, G., R. Srinivasan, P. Russell and G. M. Edelman: 1998, 'Investigating Neural Correlates of Conscious Perception by Frequency-Tagged Neuromagnetic Responses', *Proceedings of the National Academy of Sciences, USA* **95**, 3198–3203.
Vaas, R.: (1999), 'Why Neural Correlates of Consciousness Are Fine But Not Enough', *Anthropology & Philosophy* **2**(2), 121–141.
Ward, R. and S. R. Jackson: 2002, 'Visual Attention in Blindsight: Sensitivity in the Blind Field Increased by Targets in the Sighted Field', *Neuroreport* **13**, 301–304.
Weiskrantz, L.: 1986, '*Blindsight*', Oxford University Press.

Danish National Research Foundation's
Center for Functionally Integrative Neuroscience
Aarhus University Hospital & University of Aarhus
Nørrebrogade 44
8000 Aarhus C
Denmark
E-mail: overgard@pet.au.dk