Alva Noë and Evan Thompson

Are There Neural Correlates of Consciousness?[1]

In the past decade, the notion of a neural correlate of consciousness (or NCC) has become a focal point for scientific research on consciousness (Metzinger, 2000a). A growing number of investigators believe that the first step toward a science of consciousness is to discover the neural correlates of consciousness. Indeed, Francis Crick has gone so far as to proclaim that 'we … need to discover the neural correlates of consciousness.… For this task the primate visual system seems especially attractive.… No longer need one spend time attempting … to endure the tedium of philosophers perpetually disagreeing with each other. Consciousness is now largely a scientific problem' (Crick, 1996, p. 486).[2] Yet the question of what it means to be a neural correlate of consciousness is actually far from straightforward, for it involves fundamental empirical, methodological, and *philosophical* issues about the nature of consciousness and its relationship to the brain. Even if one assumes, as we do, that states of consciousness causally depend on states of the brain, one can nevertheless wonder in what sense there is, or could be, such a thing as a neural correlate of consciousness.

Our focus in this paper is one particular way of thinking about the neural correlates of visual consciousness that has become widespread among philosophers of mind and cognitive neuroscientists alike. According to this way of thinking, which we call the *matching-content doctrine,* the first task of the neuroscience of

Correspondence: Alva Noë, Department of Philosophy, University of California, Berkeley CA 94720-2390, USA. *Email: noe@socrates.berkeley.edu*
Evan Thompson, Philosophy Department, York University, 4700 Keele Street, North York, Ontario, M3J 1P3, Canada. *Email: evant@yorku.ca*

---

[1] For comments on earlier drafts of this paper, we are grateful to Kathleen Akins, Anthony Atkinson, Tim Bayne, David Chalmers, Robert Hanna, Valerie Hardcastle, Susan Hurley, Thomas Metzinger, Erik Myin, and Francisco Varela. AN gratefully acknowledges the support of faculty research funds granted by the University of California, Santa Cruz, as well as the support of a UC President's Fellowship in the Humanities. ET is supported by the Social Sciences and Humanities Research Council of Canada, the McDonnell Project in Philosophy and the Neurosciences, and the Fetzer Institute through the Center for Consciousness Studies, University of Arizona, Tucson.

[2] Crick goes on to say: 'It is not impossible that, with a little luck, we may glimpse the outline of a solution before the end of the [twentieth] century' [!].

Journal of Consciousness Studies, **11**, No. 1, 2004, pp. 3–28

consciousness is to uncover the neural representational systems whose contents systematically match the contents of consciousness. We believe there are good empirical and philosophical reasons for being suspicious of this matching-content notion of neural correlates of consciousness. There is no reason to think that the neural states that have been shown experimentally to be correlated with conscious visual experiences match those experiences in content; therefore, the experiments to date do not support the matching-content doctrine. In addition, we argue below that there is reason to doubt that neural representational systems, at least as standardly conceived in much of the NCC literature, could match visual experiences in content, and therefore that the matching-content doctrine ought to be rejected. If we are right, then neuroscientists and philosophers ought to pursue a different approach to understanding the brain basis of consciousness from that of the matching-content doctrine.

I: From Isomorphism To The Matching-Content Doctrine

According to the matching-content doctrine, for every visual experience *E*, there is a neural representational system *N*, such that (i) *N* is the minimal neural representational system whose activation is sufficient for the occurrence of *E*, and (ii) there is a match between the content of *E* and the content of *N* (Chalmers, 2000).

As background to this doctrine, consider first that most neuroscientists believe that for every conscious state there is a minimal neural substrate that is nomically sufficient (as a matter of natural law) for its occurrence. Let us call this commitment the thesis of the *minimally sufficient neural substrate* (or 'the minimal substrate thesis' for short).[3] As an example, consider a visual experience, *E*, as of two uniform regions of different brightness. According to the substrate thesis, there must be some minimal neural complex *N* that is sufficient for *E*. This minimal substrate thesis is a ground-level metaphysical and/or methodological commitment of many scientists and philosophers interested in the neural basis of consciousness.

The visual scientists Davida Y. Teller and E.N. Pugh, Jr. make this commitment explicit in their discussion of what they call the 'bridge locus':

> Most visual scientists probably believe that there exists a set of neurons with visual system input, whose activities form the immediate substrate of visual perception. We single out this one particular neural stage, with a name: the *bridge locus.* The occurrence of a particular activity pattern in these bridge locus neurons is necessary for the occurrence of a particular perceptual state; neural activity elsewhere in the visual system is not necessary. The physical location of these neurons in the brain is of course unknown. However, we feel that most visual scientists would agree that they are certainly not in the retina. For if one could set up conditions for properly stimulating them in the absence of the retina, the correlated perceptual state would presumably occur (Teller and Pugh, 1983, p. 581).[4]

---

[3] Chalmers (2000, p. 25) defines minimal sufficiency as follows: '(1) the states of N suffice for the corresponding states of consciousness, and (2) no proper part M of N is such that the states of M suffice for the corresponding states of consciousness.'

[4] The term 'bridge locus' is adopted in order to parallel the usage of 'bridge law' in the philosophy of science. Just as bridge laws are supposed to connect the predicates of two different scientific theories,

The suggestion at the end of this passage that a particular pattern of activity at the bridge locus would be sufficient for the occurrence of a particular perceptual state, despite the absence of the sort of activity elsewhere in the visual system (e.g., in the retina) that would normally give rise to such a perception, expresses what we are calling the minimal substrate thesis (see also Kanwisher, 2001, pp. 98–100; Rees *et al.*, 2002). Teller and Pugh also assert, however, that the bridge locus activity pattern is *necessary* for the occurrence of the perceptual state. There is reason to think that this requirement of necessity is too strong, for it seems possible, given the complexity and plasticity of the brain, for there to be more than one bridge locus or minimal neural substrate of a given conscious visual state, either in one individual over time, or from one individual to the next (see Pessoa *et al.*, 1998, pp. 742–3, 787-88; Van Gulick, 1998; Chalmers, 2000, p. 24).

A second prominent commitment of investigators concerns the nature of the minimal neural substrate. It is widely believed that there will be — indeed, that there *must* be — a one–one mapping (under some description) from features of conscious experience onto features of the minimal neural substrate. For example, if *N* is sufficient for the occurrence of *E*, an experience as of two uniform regions with a brightness step, then *N* must consist (for example) at least in patterns of activity corresponding to the fields of uniform brightness and in a pattern of activity that explains the perceived step-wise difference in brightness (as well as perhaps other features of the percept). Let us call this second commitment the *isomorphism constraint*.

The isomorphism constraint is best understood as a criterion of explanatory adequacy. The thought is that if there is to be an explanatory link between the minimal neural substrate and the perceptual experience, then there must be some way to establish a relation of sameness of structure (at some appropriate level of description) between elements of the substrate and elements of the experience. To suppose that there were no isomorphism in this sense would make it an utter mystery how *N* could give rise to *E* (or how the pattern of activity at *N* could *be E*, according to one view). Put another way, the thought behind the constraint is that one would never have sufficient reason to believe, of any given neural substrate, that it *is* the minimally sufficient substrate of a given experience, unless one could establish such a one–one mapping from features of the experience onto features of the substrate. Thus, to suppose that there were no isomorphism would be tantamount to the supposition that there was no intelligible connection (beyond brute correlation) between the experience and the neural locus in question.[5]

---

thereby correlating each property in the domain of one theory with a nomologically coextensive property in the domain of the other theory, so the bridge locus is supposed to correspond to the site or stage in visual processing at which each property of conscious visual experience can be correlated with a nomologically coextensive physiological (or abstract functional) property (see Teller and Pugh, 1983, p. 588, note 3).

[5] See Van Gulick (1998). This piece by Van Gulick is a commentary on our paper (written with Luiz Pessoa) 'Finding out about filling-in' (Pessoa *et al.*, 1998; for our rejoinder to Van Gulick, see p. 788). As we discuss in that article, the term 'isomorphism' in this context goes back to the doctrine of psycho-neural isomorphism introduced by the Gestalt psychologist Wolfgang Köhler (Köhler, 1947,

Most defenders of the isomorphism constraint — most neuroscientists interested in consciousness — are not committed to the further assumption that the isomorphism be *spatial* or *topographic*.[6] For example, from the fact that we experience the brighter patch as to the left of the darker patch, it cannot be deduced — even given the assumptions of substrate isomorphism — that the neural activity responsible for the greater brightness of the leftmost patch must occur in the brain to the left of the activity responsible for that of the less bright patch. Nor does the fact that we experience a continuous region as of a uniform brightness entail that the pattern of neural activity responsible for that experience must itself be continuous. To assume that the isomorphism must in this way be spatial or, as it were, pictorial, would be to commit the error, first noticed by Kant (in the Second Analogy of Book Two of the Transcendental Analytic of his *Critique of Pure Reason*), and frequently warned against by Daniel Dennett, of confusing *vehicles* and *contents* (Dennett, 1991; Dennett and Kinsbourne, 1991; see also Pessoa *et al*., 1998; see also Hurley 1998, Chapter 1). In general, what is represented offers precious little information about the way it is represented. At most what can be inferred is that, in so far as we seek to explain the occurrence of the experience and its features, there must be some account relating elements of the neural substrate to these various features.

Enter the matching-content doctrine. One way to accommodate the requirement of isomorphism at the minimal neural substrate without committing vehicle/content confusions is to suppose that the isomorphism obtains only at the level of content. A neural representational system *N* can count as the minimally sufficient neural substrate of an experience *E*, according to this way of thinking, only if there is a match between the content of *N* and the content of *E*. David Chalmers, in a recent paper, has formulated this thesis in terms of the notion of a neural correlate of the *content* of consciousness or a *content NCC,* which he takes to be the 'central case' of an NCC:

---

p. 301), who held that mental states and processes have the same structure as brain states and processes (for discussion see Sheerer, 1994). According to Köhler's principle of isomorphism, the phenomenal properties of perception (the properties of the 'phenomenal field') are isomorphic to the topological-functional properties of cortical processes. Similar views can be found today among proponents of the so-called dynamical systems approach in cognitive science (see Kelso, 1995, pp. 275–6, 289). What we are calling the 'isomorphism constraint', however, is best seen not as a substantive hypothesis about the nature of the psycho-neural relation (as was Köhler's principle, at least in part), but rather as a general thesis about explanation: successful explanation requires that there be a one–one correspondence, at some level of description, between elements of the minimal neural substrate and elements of perceptual experience. Whereas Köhler hypothesized that the right level of description to characterize the isomorphism was dynamic and topological-functional, others have proposed that it is computational. For example, Ray Jackendoff, in *Consciousness and the Computational Mind*, expresses the isomorphism constraint when he states that 'Every phenomenological distinction is caused by/supported by/projected from a corresponding computational distinction' (Jackendoff, 1987, p. 24), and he proposes that the relevant computational distinctions reside at an intermediate level of representation between the peripheral level of sensory processing and the central level of cognition or thought. For recent versions of the isomorphism notion in the NCC literature, see Metzinger (2000b, p. 1), and Revonsuo (2000, pp. 57–76).

[6] For a defence of spatially 'isomorphistic' accounts of perception in visual science, see Todorovic (1987). For critical discussion, see Pessoa *et al.* (1998).

> An NCC (for content) is a minimal neural representational system N such that representation of a content in N is sufficient, under conditions C, for representation of that content in consciousness (Chalmers, 2000, p. 31).

A number of neuroscientists seem to be working with this sort of conception. Francis Crick and Christof Koch, for example, in presenting their conception of an NCC, state that 'Whenever some information is represented in the NCC it is represented in consciousness' (Crick and Koch, 1998, p. 98).

Let us return to our earlier example of brightness perception. Here the content NCC of the visual experience (as of two uniform regions and a brightness step) would be the minimal neural system *N* such that (i) *N* is nomically sufficient for the occurrence of that experience and (ii) *N* has as its representational content *that there are two adjacent uniform regions each of a given brightness and that they differ in their brightnesses*.[7]

Note that content-matching satisfies the explanatory requirements of the isomorphism constraint. The intuitive idea is this: Suppose one discovered a neural representational system *N* such that (i) *N* represents that *p*, and (ii) *N*'s activity is correlated with the occurrence of a perceptual experience with the content that *p*. If one discovered such a neural representational system, it might seem reasonable to think that one had reached explanatory bedrock, that one had discovered the place in the brain where the conscious experience happens.

Is there good reason to believe that there are neural representational systems whose contents systematically match the contents of consciousness? We think there is good reason to answer this question negatively. Despite claims to the contrary found in the scientific and philosophical literature on consciousness, there are no known examples of the right sort of content match. Moreover, there is reason to doubt that such matches will be found, for there is reason to doubt — or so we argue below — that neural representational systems, at least as standardly conceived in the NCC literature, could match perceptual experiences in content. Although our main target below is the matching-content doctrine, the arguments we advance will also cast doubt on the minimal substrate thesis (that every conscious state has a minimal neural substrate sufficient for its occurrence).

II: Binocular Rivalry — A Best Case for a Content NCC?

Optimism about the NCC research programme stems in large part from excitement about recent work on the neural basis of binocular rivalry (see Leopold and Logothetis, 1999; Blake, 2001). This research is widely thought to provide one of the most important experimental paradigms for determining the neural states whose contents match the contents of visual consciousness, and its results have prompted a fair amount of speculation about the location of a content NCC for visual consciousness (see Crick, 1996; Crick and Koch, 1995; 1998; Logothetis,

---

[7] Although we have expressed the content of the visual experience in the form of a propositional 'that' clause, we do not mean to imply that the intentional content of visual experience is always equivalent to a whole proposition (that all seeing is 'seeing that').

1999; Chalmers, 2000; Kanwisher, 2001; Rees *et al.*, 2002). For these reasons, we need to take a close look at this research.

What happens when two different visual patterns are presented simultaneously, one to each eye? One might think that one would see both patterns, one superimposed on the other. Although one may experience this sort of superimposition at first, one comes rapidly to experience the two patterns as alternating back and forth, competing, as it were, for perceptual dominance. For example, when one eye is presented with the image of a butterfly and the other eye with the image of a sunburst, one will see the butterfly-image alternating irregularly every few seconds with the sunburst-image. Or, to take another example, when one eye views bars moving upwards and the other eye views bars moving downwards, one will see bars that periodically reverse their direction of movement.[8] This phenomenon is known as binocular rivalry.

One reason that binocular rivalry is important for studies of the neural basis of visual perception is that it would seem to provide a tool for dissociating, on the one hand, neural activity that is driven, as it were, by the stimulus sitting before the eye (e.g., the butterfly, the sunburst), and the neural activity, on the other hand, that corresponds to 'subjective perception' itself, that is, to what the subject experiences visually (the percept of a butterfly, the percept of a sunburst). Although the stimulus is constant, the percept changes dramatically every few seconds. The question is: which of the neural activities subserving visual perception correlates with the stimulus, and which with the percept? The thought is that the neural activity most closely associated with the percept is likely to be the neural activity that represents what the subject is seeing (or that encodes the content of the visual experience), and therefore either is, or is close to in the processing stream, the neural correlate of the perceptual experience for content.

This use of binocular rivalry as a probe for the neural correlates of visual consciousness has been pioneered by Nikos Logothetis and his colleagues in a series of experiments using macaque monkeys as perceptual subjects (Logothetis and Schall, 1989; Leopold and Logothetis, 1996; Sheinberg and Logothetis, 1997). (Macaques and humans are known to have similar visual systems, including similarities in binocular rivalry; see Myerson *et al.*, 1981.) What they found, in general terms, was that neural activity at early stages of the visual pathway — primary visual cortex or V1, and V2 — was better correlated, on the whole, with the stimulus, whereas at later stages the proportion of neurons whose activity correlated with the animal's percept increased, with the highest degree of correlation found in the inferotemporal cortex (IT). To assess whether these studies support the matching-content doctrine and the notion of a content NCC, however, we first need to describe them in greater detail.

In the first study, Logothetis and Schall (1989) investigated neurons in area MT (medial temporal, also known as V5), which are known to respond

---

[8] During the transitions from one pattern to the other (and depending on various spatial aspects of the stimulus), the dominant pattern may break into patches intermingled with patches of the other pattern (Blake, 2001, pp. 9–12). The transitions can also 'often appear as spatially coherent waves of dominance wherein one stimulus appears to sweep the other from conscious awareness' (p. 11).

preferentially to movement. They trained a monkey to report its perception by executing a saccade to the right or the left, depending on which of two images it saw when it was presented with rivalrous upward-moving and downward-moving horizontal gratings. The preferred direction of motion in the receptive fields of a number of cells was determined (i.e., whether the cell responded above its base firing rate for an upward-moving or a downward-moving stimulus in its receptive field), and the activity of these cells was recorded while the monkey reported its percept, both in trials with congruent gratings and in trials with rivalrous gratings. Many neurons showed patterns of activity correlated with the stimulus: their activity was modulated by non-rivalrous gratings according to their direction preference, but their response during rivalry was independent of the monkey's reported perception. Other neurons (22% of all recorded) showed patterns of activity during rivalry that were modulated in accordance with the percept. Of these, half responded during rivalry when the perceived direction of motion reported by the monkey corresponded to the cell's preferred direction exhibited during non-rivalry. Thus these cells 'reflected the monkey's reported perception of motion direction' (Logothetis and Schall, 1989, p. 761). The other half, however, responded in the opposite direction, that is, when their preferred direction of motion was not perceptually dominant, but suppressed.

In a subsequent study, Leopold and Logothetis (1996) studied neurons in V1/V2 and V4, using rivalrous stationary gratings, one oriented optimally for the neuron investigated and the other oriented orthogonally to it. Whereas the previous study recorded neuronal activity only for the first perceptual report (e.g., upward motion), this study recorded activity during several alternating phases of the percept (e.g., left-right, left-right, etc.). In total, about one-third of the neurons they examined showed patterns of activity during rivalry that were modulated in accordance with the percept, the percentage being greater in V4 than in V1/V2 (and only in V4 were there neurons that responded during the perceptual suppression of their preferred stimulus).

In another study, Sheinberg and Logothetis (1997) recorded the activity of neurons in the temporal cortex, specifically IT and the visual areas of STS (superior temporal sulcus). First, they trained the monkeys to perform a categorization task, in which the monkey pulled one of two levers — the left lever whenever a sunburst-like pattern was displayed, and the right lever whenever other images were displayed, such as images of humans, monkeys, apes, wild animals, butterflies, reptiles, and various human-made objects. Once the monkey had learned to perform the task, rivalrous periods of stimulation were introduced, in which the sunburst pattern was presented to one eye and another of the other visual images to the other eye. The experimenters then isolated a number of neurons in IT and STS that responded to such visual-image stimuli, and tested their selectivity in order to find one or more 'effective' stimuli. For example, some IT cells fired upon the presentation of a butterfly-image, but responded minimally to all other tested images, including the sunburst pattern. Finally, the experimenters recorded the activity of such selective cells while the monkey performed the categorization task in both rivalrous and non-rivalrous conditions. The rivalrous stimuli consisted of the effective stimulus presented to one eye and the ineffective sunburst pattern presented to the other eye. It was found that 90% of the recorded neurons in IT and STS fired only before and during the periods in which the monkey reported seeing the effective stimulus: the neurons fired vigorously when the monkey reported seeing the cell's preferred pattern (e.g., the butterfly) in both the rivalrous and non-rivalrous conditions, but the cell responses were almost eliminated when the monkey reported seeing the ineffective stimulus (the sunburst), even when the effective stimulus was present during binocular rivalry. Thus the activity of these temporal cortical cells 'reliably predict the perceptual state of the animal' (Sheinberg and Logothetis 1997, p. 3413).

Do these studies provide evidence in favour of the existence of a neural correlate for content of a given visual experience? A *prima facie* case can be made, along the following lines. First, these experiments show that certain cortical neurons respond preferentially to the presence of (say) a butterfly-image stimulus in their receptive fields (as determined according to standard procedures of mapping receptive fields). Second, in non-rivalrous conditions, these cells respond when the subject is presented with a butterfly-image stimulus. The cells respond equivalently in rivalrous conditions when the subject reports seeing a butterfly, and they cease to respond when the subject reports seeing (say) a sunburst-pattern (even though the butterfly-image is still present, i.e., the stimulus is unchanged). Third, the subject has a perceptual experience as of a butterfly-image when (and only when) these neurons fire. From the foregoing it seems just a short further step to the following stronger claims: there is a match between the representational content of these neurons and the perceptual content of the subject's visual experience; and it is this content match that explains the content of the subject's perceptual experience, and that informs us that we may have hit upon or be close to the content NCC for the experience. Something very close to this *prima facie* case has been presented by Chalmers (2000).[9]

III: The Lessons of the Binocular Rivalry Studies

The binocular rivalry studies of Logothetis and colleagues establish the existence of important correlations between neural systems (the activities of single

---

[9] Thus he writes: 'Interestingly, it seems that in doing this [finding a neural representational system from whose activity we might determine the contents of visual consciousness] we are crucially concerned with the representational contents of the neural systems themselves. In the Logothetis work, for example, it is important to determine the receptive fields of the cells … in order to see whether the receptive fields of active cells match up with the apparent contents of visual consciousness. In essence, the receptive field is acting at least as a heuristic way of getting at representational content in the neurons in question. Then the crucial question is whether the representational content in the neural system matches up with the representational content in visual consciousness' (2000, p. 20). And earlier on the same page: 'These results [of Logothetis and colleagues] naturally lend themselves to speculation about the location of a visual NCC. It seems that V1 is unlikely to be or to involve an NCC, for example, due to the failure of V1 cells to correlate with the contents of consciousness…. On the other hand, IT seems to be a natural candidate for the location of an NCC, due to the strong correlation of its cells with the content of consciousness. At least it is natural to suppose that IT is a "lower bound" on the location of a visual NCC (due to the failure of strong correlation before then), though the NCC itself may be farther in. None of this evidence is conclusive … but it is at least suggestive. '

neurons) and perceptual experience. But do these studies also establish the existence of a content match between the two?

Let us note, first, that this question is ambiguous. It could mean: *do the experience, on the one hand, and the neural system, on the other, have one and the same content?* Understood this way the question is: *does the experience have as its content whatever the neural system represents?* But the question can also be construed as asking whether the neural system and the experience merely *agree* in content in this or that respect. Two content-bearing states can be thought to agree in a relevant respect even though their content is, as it were, not one and the same. For example, a photograph might depict several birds flying across a sunny blue sky, and this content may agree with a verbal report to the effect that there are several birds flying across a sunny blue sky. These two items — the picture and the verbal report — agree in their content even though they do not have one and the same content. The picture, for example, represents the specific shade of blue of the sky, and the birds as in a specific spatial relation to each other. The words, clearly, say nothing about the particular shade or the arrangement of birds. Let us mark this difference between the contents of the picture and the verbal reports by saying that the two *agree* in content, but do not *match* in content. (We can make this point without first settling on a general theory of kinds of content.)

With this distinction in place, it is immediately clear that the binocular rivalry studies provide evidence of striking content *agreement* between perceptual experiences and neural systems, but provide no evidence of any content *match* (i.e., of identity of content). The reason is that the content of the neurons in these studies is receptive-field content (what we will call RF-content), but the content of perceptual experience is not. Exactly what RF-content is and what it tells us about how the brain works are matters of dispute in neuroscience. In general, it seems fair to say that most neuroscientists now believe that perceptual attributes, such as the orientation of a line in a grating, are not encoded by single cells, but by populations or clusters of cells, and that individual cells can participate at different times in the representation of different perceptual attributes or objects. Nevertheless, in the Logothetis studies the RF-contents are delineated in terms of patterns of causal covariation between single-cell responses and various stimulus parameters (orientation, direction of motion, and object category). A neuron is taken to represent the presence of a vertical line, for instance, in so far as that neuron fires preferentially given the presence of such a line in its receptive field. Such a correlation is supposed to serve as a 'heuristic way of getting at representational content in the neurons in question' (Chalmers, 2000, p. 20), that is, of getting at the representational content of the populations, clusters, or assemblies of cells that encode a given perceptual attribute. As Rees *et al.* (2002, p. 262) state: 'Recording the spike activity of neurons is considered to be the gold standard that is necessary for a quantitative and reductionist explanation of perception and behaviour in terms of its underlying constitutive elements.' But the perceptual experience as of a vertical line is never just a matter of the registration of the presence of a vertical line in this way. The perceptual experience as of a vertical line will represent the line as against a background, and as occupying a certain position in egocentric space, i.e., as occupying a certain spatial relation to *you*, the embodied perceiver. Thus RF-content and the content of perceptual experience are two different kinds of content.

This point bears directly on how to understand the significance of the Logothetis experiments. What these experiments show is not that the representational contents of single neurons or cell populations in IT and STS *match* (are the same as) the contents of visual consciousness, but rather that there are strong correlations (i.e., *agreement*) between two different sorts of content found at two different levels — RF-content at the level of single neurons (and by hypothesis neural representational content at the level of cell assemblies), and experiential content at the level of the monkey's perceptual report. Accordingly, the implication of this important finding cannot be that IT and STS, or some location 'higher up in the visual hierarchy' with similar representational contents, are candidates for the location of a visual NCC, in the precise sense of a neural representational system *N* such that representation of a content in *N* is sufficient, under conditions *C*, for representation of *that very same content* in consciousness.

It is not clear to us whether Logothetis himself advances any such matching-content interpretation of his results.[10] In his 1997 study with Sheinberg, he suggests that the responses in V1/V2, V4, and MT to the perceptual dominance of the preferred stimulus, and the responses in V4 and MT to its perceptual suppression, 'may be the result of the feedforward and feedback cortical activity that underlies the processes of grouping and segmentation — processing that are probably perturbed when ambiguous figures are viewed' (Sheinberg and Logothetis, 1997, p. 3413). On the other hand, he states that the responses of cells in IT and STS 'may represent a stage of processing beyond the resolution of ambiguities [in the stimulus representations], where neural activity reflects the integration of constructed visual percepts into those subsystems responsible for object recognition and visually guided action' (ibid.). The term 'reflects' in this sentence is unclear. Our view is that it should be taken to mean only that the responses of cells in IT and STS agree (in our sense) with the content of the monkey's report, but not that they match it.

Finally, there is a point that needs to be made about the notion of RF-content. As we have seen, Chalmers suggests that RF-content corresponds to the 'content of the neural representational systems themselves' or provides a 'heuristic way of getting at representational content in the neurons in question' (Chalmers, 2000, p. 20). The qualification 'themselves' might be taken to mean that the representational content can be established as belonging to the 'core' NCC neurons apart from the background-setting of other brain activity and the awake,

---

[10] With respect to the neural basis of visual consciousness, Logothetis seems not to be of the opinion that the NCC resides at some level 'higher up in the visual hierarchy,' for he writes: 'the findings to date already strongly suggest that visual awareness cannot be thought of as the end product of such a hierarchical series of processing stages. Instead it involves the entire visual pathway as well as frontal parietal areas, which are involved in higher cognitive processing. The activity of a significant minority of neurons reflects what is consciously seen even in the lowest levels we looked at, V1 and V2; it is only the proportion of active neurons that increases at higher levels in the pathway' (Logothetis, 1999, p. 74).

contextually situated animal as a whole. In fact, matters seem not so straightforward, for there is a good deal of evidence to suggest that the receptive-field properties of single neurons depend on what the animal as a whole is doing. The receptive fields in the Logothetis experiments are mapped while the monkey maintains fixation on a spot; its head is also restrained by a surgically placed head restraint post. But it has been shown that many kinds of response in visual cells are highly dependent on behavioural factors, such as the bodily tilt of the animal (Horn and Hill, 1969), the animal's posture (Abeles, 1984; Abeles and Prut, 1996), and auditory stimulation (Fishman and Michael, 1973; Morell, 1974). Furthermore, studies in alert monkeys have shown that attention and the relevance of a stimulus for the performance of a behavioural task can considerably modulate the responses of visual neurons (Chelazze *et al*., 1983; Moran and Desimone, 1985; Haenny *et al*., 1988;. Treue and Maunsell, 1996). In fact, context-dependent phenomena like these are present in the Logothetis experiments. In the 1989 study by Logothetis and Schall, and the 1996 study by Leopold and Logothetis, it was found that the response properties of neurons modulated by the rivalrous percepts 'were influenced by the perceptual requirements of the task' (Leopold and Logothetis, 1996, p. 551). Some of the cells that responded preferentially to the direction of motion or the orientation of a grating, when the monkey's task was to discriminate these features, showed no such preferences when their receptive fields were mapped conventionally during a fixation task (in which the monkey had to fixate a spot). Other cells responded preferentially to these features during rivalrous but not congruent stimulation. Leopold and Logothetis state that 'Such cells may be mediating the effects of selective attention on perception' (ibid.). Sheinberg and Logothetis (1997) also address this issue; they come to the conclusion that 'the phenomenon of binocular rivalry is also a form of selection, but that this selection occurs between competing visual patterns even in the absence of explicit instructions to attend to one stimulus or the other' (p. 3413).[11] The moral we would draw from these considerations is one insisted upon by Francisco Varela: there would seem to be no way to establish the RF-contents of the neurons themselves (and by extension the contents of cell assemblies) independent of the sensorimotor context of the animal as a whole (Varela, 1984; Varela *et al*., 1991, p. 91).

We have argued, in this section, that recent studies of the neural basis of binocular rivalry provide no evidence to support the claim of a content match between neural activity and perceptual experience (as opposed to the weaker claim of content agreement in this or that respect), and hence no evidence for the existence of a content NCC for visual perception. In the next section, we defend the stronger claim that there is reason to doubt that neural representational systems of the sort we have been considering could match perceptual experiences in content.

---

[11] Further evidence that perceptual rivalry is a form of selection comes from two studies by Erik Lumer and colleagues, who used functional magnetic resonance imaging (fMRI) to measure brain activity in humans during perceptual rivalry. They showed that rivalry engages right frontoparietal areas, which have been implicated in tasks requiring visuospatial attention and working memory (Lumer *et al*., 1998; Lumer and Rees, 1999).

IV: The Incommensurability of Content

Our argument in this section depends on distinguishing a number of interrelated features of the content of perceptual experience. Perceptual content is, as we shall say, *structurally coherent*, *intrinsically experiential*, and *active and attentional*. RF-content, on the other hand, seems to be none of these. For this reason, it is doubtful that there could be content matches between perceptual experiences and RF neural representational systems (whether other kinds of neural representational systems could match these characteristics of perceptual content is considered below).

The structural coherence of perceptual experience can be brought out by a consideration of figure–ground structure. It is difficult to make sense of the idea of a visual experience as of (say) a butterfly that was not also the experience as of a butterfly against a background of some sort or another. It seems to be a basic fact about our visual phenomenology that our visual experiences have figure–ground structure in this sense (see Merleau-Ponty, 1962, pp. 4, 22; see also Gurwitsch, 1964). If this is true, then the presence of the butterfly and the presence of the background cannot be separable ingredients of content, as they would seem to be if each answered to the RF-contents of distinct groups of neurons. Figure–ground structure seems to be a global, non-atomistic property of visual experience.[12] It is difficult to see how it could be built up, as it were, out of RF-content atoms.

One might object that this argument is too fast. There is no need to accept that RF-content is decomposable in the way that (we claim) the subject's perceptual experience of figure and ground is not. As an empirical matter, it can be expected that the neural populations subserving figure and ground in a given visual experience will always be linked.

Readers familiar with neuroscience will recognize that, with this point, we are in the vicinity of the so-called *binding problem,* the problem of understanding how distinct attributes (shape, colour, motion, and so on), as 'processed' by single neurons or cell populations in spatially distributed and functionally segregated neural pathways, can be bound together to give rise to a perceptual experience with a unified content. In the words of one study:

> Neurons in most areas of the visual cortex process information only from a limited part of the visual field and respond only to a restricted range of feature constellations. Thus, the outputs of numerous cells must be integrated to create a complete representation of a particular object.… In the visual system, the binding of features pertaining to individual objects appears to be a prerequisite for figure–ground

---

[12] One might object to this point about the structural coherence of perceptual experience that it ignores the neurological evidence of the remarkable structural and qualitative dissociations that occur in visual experience as a result of brain damage (as discussed in, e.g., Milner and Goodale, 1995, and Weiskrantz, 1997). The point to be made in reply is that such pathologies are never equivalent to the shattering of normal perceptual experience into isolable atoms, for they always involve global reorganizations of experience (and also behaviour), thereby maintaining holisms of structural coherence, albeit pathological ones. This point goes back to Kurt Goldstein's classic 1934 work, *Aufbau des Organismus* (Goldstein, 1995).

> segregation and scene segmentation, i.e. for the distinction between several objects present in the visual field (Engel *et al.*, 1992, p. 218).

One currently popular approach to this problem is to hypothesize that feature-detecting neurons are bound together into neural assemblies by firing in synchrony at various frequencies over a limited period of time on the order of a fraction of a second (see Singer *et al.*, 1997). Experimental results have shown that the probability and strength of synchrony in neural responses are correlated with perceptual Gestalt or elementary grouping criteria, such as continuity, proximity, similarity, and common fate. For example, neurons with non-overlapping receptive fields have been shown to oscillate in synchrony in the gamma frequency range (around 40 hertz) if they are activated by a single continuous stimulus, whereas their responses are uncorrelated if they are presented with two independent stimuli moving in opposite directions, one in each receptive field (see Engel *et al.*, 1992, and Singer, 1995 for reviews). Although mainly low-level grouping criteria have been studied so far, it is reasonable to expect that response synchronization subserves scene segmentation, perceptual grouping (Singer, 1995), and indeed the overall integration of all dimensions of a cognitive act, including associative memory, emotional tone, and motor planning (Thompson and Varela, 2001; Varela *et al.*, 2001).

In response to this line of criticism, let us note that the finding that the synchronization probability of visual neurons reflects Gestalt criteria does not in itself succeed in establishing *content matching,* although it does point to some measure of *content agreement*.[13] Consider that in visual perception the kind of structurally coherent content with which we are concerned is that as of a figure-on-a-ground *located in egocentric space* (see Merleau-Ponty 1962, p. 101). Egocentric space is defined by one's whole body and the possible ways it can move (or so we think it reasonable to assume). But what sense can be made of the notion of a neural representation having as part of its content that a figure (more precisely: a stimulus with which it causally covaries) is located in egocentric space? Although neural systems causally enable the animal as a situated agent to orient itself in its egocentric space, they themselves do not inhabit this space, nor do they have any access to it as such. If this line of thought is right, then the binding problem, as standardly articulated, is malformed, because it presupposes a faulty conception of content. Content, at the level of the perceptual experience of the animal, is not composed of (does not have as its proper parts) neural registrations of features (e.g., RF-contents), and therefore is not the result of binding together such features.[14]

---

[13] In this connection it is worth noting that some of these studies were carried out in anaesthetized animals, in which, therefore, there was no visual experiential content at all. Although the findings have been replicated in awake animals, the fact that they can be obtained under anaesthesia would seem to call into question any content-matching claim made for the neuronal responses.

[14] Not all neuroscientists believe that the binding problem as standardly conceived is a genuine problem, rather than a reflection of a certain way of thinking about perception. See, for instance, Kelso, 1995, pp. 250–5).

This point brings us to the next distinctive feature of perceptual content we wish to distinguish. Perceptual content is intrinsically *experiential*, in the sense that the content of an experience is always the content as represented from a point of view. Thus, to have a visual experience as of a butterfly is to have an experience as of the butterfly standing in a certain egocentric spatial relation to you, and as standing out against a background relative to you. But to suppose that a neural system could have this sort of content is to suppose that the neural system *experiences* things from a point of view (e.g., that it experientially represents a butterfly as occupying a position in *its* egocentric space). Animals and persons experience the world as laid out before them, but the neurons do not.

It should be noted that the sense in which we experience the figure-as-against-a-ground in egocentric space need not be the same as that in which we experience one item as simply in front of another. In particular, it is frequently the case that we experience the ground as present, but as present, as it were, only in the background, yet available to attention should we actively turn to it.

With this point we arrive at another distinctive feature of perceptual experience, its active and attentional character. If you see a stationary cat through a picket fence, to take another example, you may have a visual experience as of a cat, even though, in fact, you cannot see the parts of the cat that are occluded by the fence. This example puts us in the vicinity of an important phenomenological point. Strictly speaking, you do not *see* the occluded portion of the cat, but nor do you merely *think* of that portion of the cat as present (see Kanizsa and Gerbino, 1982). The cat, as a whole, is perceptually present. The unseen portions of the cat are present, phenomenologically speaking, in much the way that the background is always present when you see a figure. The presence in question is perceptual or conscious, without being strictly *visual*. We believe that the relevant sense of presence is not so much a visual one as an attentional one. The figure is foregrounded not because it is experienced as in front of the background, but because it is foregrounded in attention. Similarly, the unseen portions of the cat are experienced as present because one implicitly understands (as with the case of the background) that one can actively bring them into primary focus through movements of the head and redirections of attention (see O'Regan and Noë, 2001; Noë and O'Regan, 2000; Noë, 2002; 2004).

Another feature of visual experience that exemplifies the structural coherence and active, attentional character of perceptual content is the unboundedness of the visual field. It is not possible to draw a sharp line where the visual field drops off at its perimeter. Our parafoveal powers of discrimination are quite poor; nevertheless, there is no line that can be drawn, in experience, between what is central to vision and what is, as it were, at its peripheral limits. It is well known that the distribution of receptors across the retina is not homogeneous: receptor density falls off sharply as one moves outward from the fovea on the retina. Why, one might wonder, do we not experience the visual environment as vague and indeterminate at the boundaries of the visual field? The answer is that we do not inhabit the sort of visual experience we encounter on fixation tasks (i.e., the content of experience is not merely the result of integrating successive fixations).

Normally, we move our eyes, as well as our heads and bodies, and in this way we are able to explore different aspects of the scene in front of us. Just as the sense of the perceptual presence of the cat (in the earlier example) depends on our implicitly understanding how to gain access to more of the cat — by eye, head, body, and attention movements — so our sense of the presence of the environment as a whole depends on our being poised to direct our perceptual attention this way and that as needed (see O'Regan and Noë, 2001). The perceptual experience of the scene cannot be thought of as a snapshot-like occurrence. Rather, it is a temporally extended process of attentive engagement, and the content of experience is brought forth or enacted by this activity (Noë, 2004).

The active and attentional character of perceptual content has two important implications for how we think about the relationship between perceptual experience and the brain. First, it suggests that the content of experience itself — the content of one's temporally extended, active, and attentional encounter with the environment — cannot be represented by neural representational systems as these are standardly conceived in most of the NCC literature.[15] Second, it calls into question the minimal substrate thesis. After all, if perceptual content depends on the skillful activity of the whole animal or person, making use of its capacities for eye, head, and whole body movements, and for directed attention, then it becomes questionable whether there is any such thing as a minimal neural substrate sufficient to produce experience (see Thompson and Varela, 2001; O'Regan and Noë, 2001).

One might wonder, at this point in the argument, whether a dynamical systems notion of neural representation could be employed in defence of the matching-content doctrine. According to the dynamical systems view, the neural processes most relevant to an understanding of the brain basis of consciousness are to be found at the level of 'dynamical brain signatures' (Lutz *et al.*, 2002), understood as large-scale dynamical patterns of activity over multiple frequency bands (Thompson and Varela, 2001; Varela *et al.*, 2002), rather than the structural level of specific circuits or classes of neurons (as advocated by, for example, Crick and Koch, 1998). The dynamical approach emphasizes that perception and cognition are intrinsically temporal phenomena — they happen *in* time, not simply *over* time (Van Gelder, 1999, p. 244). Our emphasis on the temporally extended, active, and attentional character of perception complements the dynamical approach. One might think, then, that our arguments should not be taken to show that there cannot be a neural-perceptual content match at all, but only that a much richer, intrinsically dynamical notion of neural representation is needed to effect such a match. This is an interesting suggestion. It remains to be seen whether the

---

[15] This conception seems to be roughly equivalent to that of the cognitivist view of the brain, according to which, in the words of Walter Freeman (who opposes this view) 'information is given by stimuli from the environment. It is encoded in trains of action potentials as bits that represent qualities, aspects or features of the stimuli. The features are transmitted by axons that serve as channels to the brain, where they undergo binding into representations of the stimuli by synaptic networks of neurons acting as summing junctions and logical gates. The contents are stored in memory banks as representational fixed patterns. They are recalled by being read out like the content-addressable memories in computers, so they can be matched or correlated with representations of new inputs' (Freeman, 1999, p. 26).

dynamical systems approach can be deployed in this way to fashion such a dynamical notion of neural representation. Nevertheless, it is worth mentioning that ths proposal would seem to go against the grain of the dynamical approach, for two related reasons. First, dynamicists reject the representational conception of mind, and thus shun talk of 'internal representations' (e.g., Varela *et al.*, 1991; Thelen and Smith, 1994; Port and Van Gelder, 1995; Van Gelder, 1995). Second, they reject the view that cognitive processes take place wholly within the central nervous system, holding instead that cognitive processes span the nervous system, the body, and the environment (Chiel and Beer, 1997). Arguments that this perspective should be extended to consciousness have recently been presented (Thompson and Varela, 2001). Thus there seems to be little impetus within the dynamical approach to search for content matches between internal neural representational systems and perceptual experience. This is of course compatible with the possibility that one can fashion a rich dynamical account of neural representational systems. To what extent such a dynamical account would support the matching-content doctrine remains an open question.

Our considerations in this section, especially of the temporally extended, active, and attentional character of perceptual experience, give support to our position, presented elsewhere, that the content of perceptual experience is personal-level content, not subpersonal-level content (Pessoa *et al.*, 1998; Thompson *et al.*, 1999). There are familiar philosophical reasons, having to with considerations of the holism and normativity of the mental, for insisting that there is a difference in kind between personal-level content and the content of subpersonal neural states (see Dennett, 1969; 1987; McDowell, 1994). Attributions of content at the personal level, for instance, are made against the background of assumptions about the beliefs, desires, and rationality of the agent who is the subject of those attributions; such content attributions have the effect of *rationalizing* the agent's behaviour. On the other hand, no such intentional and holistic concerns arise for attributions of content at the subpersonal level. Although the reasons we have presented for being suspicious of the matching-content doctrine are independent of these considerations, they nonetheless give additional support to the position that the two sorts of content are incommensurable. Subpersonal-level content (e.g., the content of neural systems) may causally enable personal-level content (e.g., the content of perceptual experience) without being constitutive of it (i.e., neural systems may causally enable perceptual experience without its being the case that the content of perceptual experience is simply the content of subpersonal states at a certain level of processing). (See also Hurley, 1998.)

If this position is right, then cognitive neuroscience should not burden its empirical investigations with the problematic weight of the matching-content doctrine and the notion of a content NCC, but should instead seek other conceptual frameworks for articulating the relationship between consciousness and the brain, ones in which consciousness is not conceptualized as an internal occurrence in the 'mind–brain', but rather as a complex set of capacities of embodied and situated agents (Thompson and Varela, 2001). On this view, the task for neuroscience should not be to investigate consciousness according to the classical framework of psychophysical correlation, whose modern incarnation is the NCC programme, but rather to study the neurobiological processes that causally enable (but do not constitute) our embodied mental life.[16]

V: The Content Mirror Argument

Another argument in support of content matches between neural representational systems and consciousness can be found in Chalmers' article on neural correlates of consciousness (Chalmers, 2000). This argument turns on general considerations about the relationship between the contents of consciousness at the personal level and the subpersonal contents of neural systems. Chalmers believes that 'one can plausibly argue that there is reason to expect that conscious contents will be mirrored by the contents of a neural representational system at some level of abstraction' (p. 35). Here is his argument:

(1) The contents of consciousness correspond well to contents made directly available for verbal report (in creatures with language), or more generally to contents made directly available for the global voluntary control of behaviour.
(2) Therefore, there is a correspondence between the contents of consciousness and contents revealed or exhibited in certain functional roles within the system.
(3) Therefore, there is good reason to believe that conscious contents are represented at some point within the cognitive system.[17]

The burden of the argument is carried by the relation between (1) and (2). ((3) follows trivially from (2), given a certain minimal conception of representational content as equivalent to functional role.) The problem is that these two steps, and hence the argument as a whole, are equivocal with respect to the personal/subpersonal distinction we invoked at the end of the last section (and that formed the basis of an earlier argument of ours, to which Chalmers is responding). On the one hand, (1) might mean that the contents of consciousness correspond well to contents made directly available *to the person* for verbal report, or more generally to contents made directly available for the *person's* or *animal's* global voluntary control of behaviour. We accept this point; it is tantamount to a conceptual claim about what is required for the attribution of conscious contents to the person. On the other hand, (1) might mean that the contents of consciousness correspond well to contents made available to some subpersonal 'global

---

[16] For parallel arguments, see Putnam (1999).

[17] On the basis of this argument, he then writes: 'Pessoa [Thompson, and Noë] … argue … against a "uniformity of content" thesis, holding that one should not expect a match between the "personal" contents of consciousness and the "subpersonal" contents of neural systems. I agree that the existence of such a match is not automatic, but… I think that the fact that conscious contents are mirrored in specific functional roles gives reason to believe that they will be subpersonally represented at least at some level of abstraction' (p. 37).

workspace' (Baars, 2002) and hence for the subsequent subpersonal control of behaviour. This proposition cannot be a conceptual claim, but must be empirical: it could turn out to be true, but could also turn out to be false.

Now we come to (2). What is the relationship between (1) and (2)? On the first (conceptual, personal-level) interpretation of (1), (2) does not follow, and indeed seems question-begging against the thesis that the contents of consciousness at the personal level do not match the contents of functional roles at the subpersonal level. On the second (empirical, subpersonal) interpretation of (1), (2) does follow, but does nothing to establish the doctrine of content matching between the personal and subpersonal levels.

The upshot, then, is that the content mirror argument is either question-begging or insufficient to give us reason to believe that the contents of consciousness at the personal level are one and the same as the contents of subpersonal neural representations.

VI: The NCC Programme and the Internalist Conception of Experience

In this penultimate section we turn our attention to the intersection of the NCC research programme and an ongoing philosophical controversy about the nature of perceptual content. In particular, we argue that the NCC research programme rests on a philosophically controversial internalist conception of the content of perceptual experience. We cannot show here that this conception is wrong. But in showing that the NCC programme prejudges an outstanding philosophical issue, we call into question Crick's all too confident claim that scientific research on the neural correlates of consciousness rescues the problem of consciousness from the hands of philosophers (a claim that some philosophers too seem willing to embrace: Churchland, 1994). Because the internalist conception might be wrong, the fate of the NCC research programme hangs on the resolution of a significant philosophical controversy.

The NCC programme, as we have seen, is committed to the minimal substrate thesis, i.e., to the thesis that every perceptual experience has a minimal neural substrate whose activation is nomically sufficient for its occurrence. But the NCC programme is also committed, *a fortiori,* to the proposition that the activation of the NCC for a perceptual experience is sufficient for that experience whether or not the resulting experience would be veridical (whether or not things are as the experience would represent them as being). What makes this account internalist is that it views the experience as supervenient on neural processes alone. Indeed, there is a close relation between the isomorphism constraint, the substrate thesis, and this supervenience thesis. According to the isomorphism constraint, there must be a one-one mapping (at some level of description) between features of experience and features of the minimal neural substrate. As Jaegwon Kim (1993, p. 178) has discussed, it is natural to take this belief in psychophysical isomorphism to be grounded on, or at least motivated by, the following doctrine of psychophysical supervenience: *all psychological states and processes supervene on the contemporaneous and internal physical states of the* *organism.*[18] ('Contemporaneous' means that any psychological state occurring at a time supervenes on a physical state of the organism at the same time; and 'internal physical states' excludes relational properties of the organism.) On the NCC view we have been considering in this paper, perceptual experiences supervene entirely on contemporaneous, internal, physical states, in particular on 'core' or 'minimal' states of the brain (to echo Chalmers's terminology), and are accordingly supposed to be wholly independent of what is going on elsewhere in the brain and body, or in the external world.[19]

To see what this commitment comes to, suppose that you see a dagger. There is a pattern of neural activity associated with your seeing it. This pattern of neural activity is distinct from the one associated merely with the presence of a dagger to your eyes. It is the former sort of neural activity that constitutes the NCC for the experience of the dagger. Given that the NCC is thought of as that pattern of neural activity sufficient for the occurrence of the experience, it follows that the activation of the NCC gives rise to the experience whether there is a dagger there or not.

A consequence of this approach is that it must be possible for there to be perceptual experiences that are qualitatively identical but differ only in respect of their veridicality. There can be no intrinsic difference between a perceptual experience and its hallucinatory counterpart. In practice, of course, we may find that there is a difference between a given hallucination and the corresponding veridical perception. After all, what it is like to hallucinate an elephant in the corner may be rather different from what it is like actually to see a similarly placed elephant. Perhaps the hallucinatory experience seems fragile and unstable. Or perhaps the hallucination gives rise to strong feelings of fear (knowing, as one might, that there is no elephant there). But this possibility presents no problem to the NCC-programme's commitments, for the mere fact that there is a qualitative difference between the two cases shows that they are not instances of two qualitatively identical experiences, and therefore that they are supported by distinct NCCs. What the NCC approach is committed to, as we have already stated, is the possibility that there could be veridical and non-veridical experiences as of seeing that are qualitatively indistinguishable. It *must be the case* that it is possible to have an hallucinatory experience as of the elephant that is *exactly like* the experience of seeing the elephant. For it is possible, at least in principle, to produce the internal states and processes that are nomically sufficient for the experience (i.e., to produce the NCC for that experience) whatever is going on outside the head of the would-be perceiver.

This conclusion is in harmony with the line of reasoning, found persuasive by many philosophers, known as the *argument from illusion* (see Ayer, 1969,

---

[18] As Kim discusses, one standard objection to the thesis of psychophysical isomorphism is the multiple realization argument (creatures that are neurally non-isomorphic can nonetheless be psychologically isomorphic), but we have allowed the isomorphism constraint to include abstract-functional characterizations of the nervous system.

[19] The modal strength of this supervenience thesis — whether it should be taken as nomological supervenience, or metaphysical or logical supervenience — need not concern us here, for the problems with internalism we go on to discuss apply regardless.

Chapters 1–2). According to this line of reasoning, it must be the case that it is possible to have an hallucination with a certain content that is qualitatively just like a veridical experience with the same content, for otherwise we would never be fooled by our hallucinatory experiences into thinking we see what in fact we merely hallucinate. (In fact this inference does not follow, as J.L. Austin [1962] noticed. Your inclination to believe that a Vermeer is a Rembrandt does not entail that there is no qualitative difference between Rembrandts and Vermeers.) The upshot of the argument from illusion — to which it is now clear that the NCC programme too is committed — is that in seeing and undergoing the corresponding hallucination you will be in one and the same experiential state. That in one instance you are seeing, and in the other merely hallucinating, is strictly speaking irrelevant to the question of what you are experiencing. One way to put this point is to note that the argument from illusion, and the NCC programme, take as primitive the notion of an *experience-with-a-content* and then define perceiving and hallucinating in terms of this primitive notion (see Child, 1994, Chapter 5). You see that p if you have a visual experience that p when in fact p is the case (plus the satisfaction perhaps of further conditions as well). You merely hallucinate that p when you have the same experience, but without the satisfaction of those further conditions.

The problem is that the internalist conception of experience is highly problematic. Numerous philosophers have challenged it over the past fifty years.[20] One of the main lines of argument is to deny a key premise of the internalist conception, namely, that there is a common experiential content (a 'highest common factor') to perceptual experiences and the corresponding hallucinations. There is all the difference in the world, one might argue, between something's looking a certain way to one, and its merely seeming to one as if something looks a certain way to one. In the first case, one's experience involves an object in the world. In the second, it does not. If this is right, then there is no common content to veridical and hallucinatory experiences, and so we can reject the idea that we are aware of one and the same thing when we perceive and when we hallucinate.

It goes beyond the scope of this paper to review this line of argument, and we are confident that we would be unable, in doing so, to lessen the controversy and settle matters one way or the other. But nor are we content to leave the matter here with the observation that the NCC research programme rests on philosophical commitments that are fraught with controversy. We would like, briefly, to outline *one* line of criticism of the internalist conception that builds on our earlier discussion in Section IV. We will then conclude this section by considering the interplay of these issues and work on binocular rivalry and neural correlates of consciousness.

In Section IV we concluded that perceptual experience is active and attentional. The significance of this point is that the content of perceptual experience is not given all at once, in a momentary occurrence, but is made available to the perceiver by the perceiver's active exploration. This position, we believe, is

---

[20] See Sellars (1956); Austin (1962); Strawson (1979); Snowdon (1980-81); McDowell (1982); Hurley (1998); Putnam (1999).

phenomenologically plausible. If you reflect on your visual experience, now, it does not seem to you as if all the visible detail is represented in your consciousness all at once, now. Rather, you take the environment (for example, this page) to be present in all its detail and to be available, now, to your inquiry (thanks to movements of the eye and shifts of attention).

The conception of experience as a temporally extended pattern of activity has broad implications for the assessment of internalism. Given this conception, it becomes unclear precisely what it can mean to say that the experience of seeing something, and the experience of merely hallucinating that you see it, are or can be qualitatively one and the same.[21] Let us pick up a theme from Austin's criticism of the argument from illusion. From the fact that perceivers cannot tell, in the first person, whether they are perceiving or merely hallucinating, it does not follow that there is no difference in the qualitative character of these experiences. (From the fact that one might, in a certain lighting and from a certain vantage point, mistake a tree for a person, it does not follow that the experience as of the tree was exactly like the experience one would have had, had a person been there.) Experiences are not like perceptual objects that can be placed side by side and assessed for their similarity and difference. They are things we do, and they can be more or less similar in so far as what we do is more or less the same. One of the crucial points against internalism is that what we do when we engage in perceptual exploration of the environment depends not just on neural activity, but on the interplay between brain, body, and environment. If this is right, then one might reasonably doubt that neural activity alone could be sufficient for the occurrence of full-fledged experience. But this is just to entertain the possibility that the minimal substrate thesis might be false.

This point ties in with our discussion of the incommensurability of RF-content and perceptual content. Here, as there, we see that the NCC research programme — and indeed, the entire internalist, Cartesian tradition of which, we can now see, it is clearly the heir — rests on the assimilation of perceptual experience to a snapshot-like, phenomenal episode in the mind–brain.

---

[21] One might also wonder exactly what it means to say that two experiences are 'qualitatively identical'. The answer given by the internalist would presumably be that two experiences are qualitatively identical just in case they share all and only the same *phenomenal properties*. Thus, to say that there is no experiential difference between a veridical visual experience and a corresponding nonveridical one is to say that the two experiences have the same phenomenal properties (see Chalmers, 2000, pp. 22–4). In this connection, it is worth pointing out that Hilary Putnam (1999, pp. 128–32), has recently presented an important argument against the coherence of the notion of phenomenal states (states that instantiate phenomenal properties) that obey the 'highest common factor' principle, i.e., that obey the principle that 'if the two occasions seem identical to the subject as far as their appearance is concerned, then an identical "highest common factor" (an identical phenomenal state) must be present' (p. 130). The gist of the argument is that the identity condition for phenomenal states is supposed to be indistinguishability in appearance to the subject, but being in the same state is a transitive relation, whereas indistinguishability is not. Hence from the fact that the experience had on two occasions is absolutely indistinguishable to the subject, it does not follow that the subject was in the identical phenomenal state on those two occasions. If Putnam's argument is sound, then the NCC research programme would seem to rest on a problematic — arguably incoherent — conception of phenomenal properties.

These points weigh not only against the philosophical presuppositions of the NCC approach, but against the directions taken in the analysis of specific phenomena such as binocular rivalry. It would seem to follow, given the familiar assumptions of the NCC programme, that the experience of binocular rivalry ought to be qualitatively identical to the experience you have whenever the appropriate pattern of alternation is set up between (say) the butterfly-image content NCC and the sunburst-image content NCC. In particular, it ought to be the case that the experience of perceptual reversal due to endogenously produced rivalry is qualitatively identical to the experience you would have if you were presented with two alternating butterfly and sunburst stimuli (so long as the alternations had the correct temporal pattern.) Just as it ought to be possible to have a nonveridical visual experience as of a dagger that is qualitatively identical to the veridical perception of the dagger, so it ought to be possible to have an experience of binocular rivalry that is qualitatively indistinguishable from the experience encountered when presented with actually fluctuating stimuli.

The problem with this consequence, however, is that it is phenomenologically implausible. The experience of binocular rivalry does not just consist in the experience of switching back and forth between (in our simplified case) the butterfly-experience and the sunburst-experience. There is a distinctive *feel* to the experience of rivalry, stemming presumably from the endogenously produced nature of the reversals. A similar point can be made about the experience of ambiguous figures such as the Necker cube. When you experience the perceptual bistability of the different aspects of the Necker cube it does not seem to you, visually, as if the cube itself is changing. There is something in what it is like to undergo the reversals that makes it clearly distinct, experientially, from seeing a changing stimulus.

So far, however, we do not seem to have much of an argument against the NCC approach. The NCC proponent is prepared to admit that if there is this phenomenological difference in any given case, then this is a difference for which there will be a neural difference, as a result of which there will be distinct patterns of neural activity at the two respective NCCs. Indeed, one recent study has shown that there is neural activity in frontoparietal cortex associated with rivalry, but not associated with exogenously produced perceptual reversal (Lumer and Rees, 1999).

But this response, while it wards off the attack on the internalism implicit in the NCC programme, ignores a lower-level but nonetheless important weakness in the position, one which stems from the position's reliance on the internalist, snapshot conception of experience. In general, it seems a mistake to analyse the perceptual experience of binocular rivalry as the undergoing of a pattern of transitioning between the two different experiential states plus some further admixture of qualitative difference, namely, the experience of endogenously produced switching. The experience of rivalry — as phenomenological reflection on that experience itself bears out — seems to be its own autonomous state, one that does not have the normal experience of the butterfly-image (i.e., the experience of the butterfly in a nonrivalrous condition) and the normal experience of the sunburst-image (i.e., the experience of the sunburst in a nonrivalrous condition) as constituent repeating elements. In other words, the experience of rivalry is not, as it were, a *molecular* experience compounded out of butterfly-experience and sunburst-experience atoms. Rather, it seems better described as *one bistable experience*. The crucial point is that the bistable experience of seeing a butterfly/sunburst is not equivalent to the normal experience of seeing a butterfly, plus the normal experience of seeing a sunburst, plus the two alternating with the appropriate temporal dynamics. As a bistable perception, it is a unique sort of experience, which accordingly must be assessed on its own terms. (This is precisely what Kelso [1995] does in describing bistable percepts as 'metastable' states with their own proper neural and perceptual dynamics.)

In the face of this conclusion, two important points need to be made. First, the NCC programme has on the whole failed to appreciate this point about the unitary character of rivalrous experiences because of its adherence to what we are calling the snapshot conception of experience. From the standpoint of that conception, it is difficult to imagine a way of thinking about binocular rivalry and ambiguous perception in terms other than those of the alternation between two distinct snapshot-like experiential states. Second, and of considerable empirical import, once we adopt the conception of perceptual rivalry as an autonomous unitary experience of its own, one not decomposable into two separate, alternating experiential snapshots, it becomes clear that it is misguided to think that binocular rivalry can enable us to dissociate neural correlates of those two would-be constituent experiences from neural correlates of perceptual processes more generally. Only if we assume that in binocular rivalry we undergo an alternation between two atomic experiential states, would it make sense to suppose that binocular rivalry enables us to pick out NCCs for those experiences.

Not only is it *not* the case, then, that visual neuroscience rescues the problem of consciousness from the hands of the philosophers, it now seems evident that visual neuroscience in the NCC programme is tangled in the net of its own unquestioned philosophical presuppositions.

VII: Conclusion

Let us summarize the main points of this paper. Our main aim has been to cast doubt on the matching-content doctrine by arguing for the following two points. First, contrary to certain claims, there are no known examples of neural-perceptual content matches. Second, there are reasons to doubt that any subpersonal-level, neural representational system could match a personal-level, perceptual experience in content, and hence reasons to doubt the truth of the matching-content doctrine — not simply as a philosophical thesis, but as a methodological one guiding neuroscientific research.

In addition, we have challenged the minimal neural substrate thesis by emphasizing the active and attentional character of perceptual experience. If the content of perceptual experience depends crucially on the environment, as well as on skillful motor capacities and capacities for directed attention on the part of the perceiver as a situated agent in the environment, then it cannot be assumed without argument (as the NCC programme does) that there is any such thing as a minimal neural substrate sufficient to produce conscious experience. Rather, the substrates of consciousness — in particular of visual perceptual consciousness — seem to cut across the brain–body–world divisions.

Our challenge to the minimal neural substrate thesis also has implications for how to think about the isomorphism constraint (that there must be a one–one mapping, at some level of description, between features of experience and features of the minimal neural substrate). As we noted in the previous section, the isomorphism constraint is motivated by the doctrine of psychophysical supervenience (that all psychological states and processes supervene on the contemporaneous, internal physical states of the organism). But the reasons we gave for calling into question the minimal substrate thesis also call into question this doctrine. One might wonder, then, what motivation is left for the isomorphism constraint. On the one hand, if we are right that neural states (even characterized abstractly and functionally) are not sufficient for the occurrence of experience, then it follows that it will not be possible to explain every aspect of experience in terms of neural structures and processes. On the other hand, it is surely reasonable to ask, of any given aspect of an experience, whether it can be explained in terms of neural processes alone. If it can, then it seems that there must be a one–one mapping to be made between the relevant features of that experience and those neural processes. To suppose that there were no isomorphism in this sense would be just another way of supposing that there were no such explanation at all. To that extent, then, we can endorse isomorphism as a methodological constraint. But notice that an important shift has occurred. The isomorphism at issue is now, one might say, simply a shadow cast by our explanatory posture. It reflects our explanatory strategy in a given case, rather than a metaphysical commitment to psychophysical supervenience or internalism.

Finally, we have drawn attention in this paper to the fact that the NCC research programme is deeply wedded to a problematic and controversial internalist conception of the content of perceptual experience. The moral to be drawn is that neuroscience, far from having freed itself of philosophy, needs the help of philosophy now more than ever.

References

Abeles, M. (1984), *Local Circuits* (New York: Springer Verlag).

Abeles, M. and Prut, Y. (1996), 'Spatio-temporal firing patterns in the frontal cortex of behaving monkeys', *Journal of Physiology* (Paris), **90**, pp. 249–50.

Austin, J.L. (1962), *Sense and Sensibilia* (Oxford: Clarendon Press).

Ayer, A.J. (1969), *The Foundations of Empirical Knowledge* (London: Macmillan).

Baars, B.J. (2002), 'The conscious access hypothesis: Origins and recent evidence', *Trends in Cognitive Sciences*, 6 (1), pp. 47–52.

Blake, R. (2001), 'A primer on binocular rivalry, including current controversies', *Brain and Mind*, **2**, pp. 5–38.

Chalmers, D.J. (2000), 'What is a neural correlate of consciousness?', in Metzinger (2000a).

Chelazze, L., Miller, E.K., Duncan, J. and Desimone, R. (1983), 'A neural basis for visual search in inferior temporal cortex', *Nature*, **363**, pp. 345–7.

Chiel, H. and Beer, R, (1997), 'The brain has a body: Adaptive behavior emerges from interactions of nervous system, body, and environment', *Trends in Neurosciences*, **20**, pp. 553–7.

Child, W. (1994), *Causality, Interpretation and the Mind* (Oxford: Oxford University Press).

Churchland, P.S. (1994), 'Can neurobiology teach us anything about consciousness?', *Proceedings and Addresses of the American Philosophical Association*, **67**, pp. 23–40.

Crick, F. (1996), 'Visual perception: Rivalry and consciousness', *Nature*, **379**, pp. 485–6.

Crick, F. and Koch, C. (1995), 'Are we aware of neural activity in primal visual cortex?', *Nature*, **375**, pp. 121–3.

Crick, F. and Koch, C. (1998), 'Consciousness and neuroscience', *Cerebral Cortex*, **8**, pp. 97–107.

Dennett, D.C. (1969), *Content and Consciousness* (London: Routledge and Kegan Paul).

Dennett, D.C. (1987), *The Intentional Stance* (Cambridge, MA: The MIT Press/A Bradford Book).

Dennett, D.C. (1991), *Consciousness Explained* (Boston, MA: Little Brown).

Dennett, D.C. and Kinsbourne, M. (1991), 'Time and the observer: The where and when of consciousness in the brain', *Behavioral and Brain Sciences*, **15**, pp. 183–247.

Engel, A.K., König, P. Kreiter, A.K., Schillen, T.B. and Singer, W. (1992), Temporal coding in the visual cortex: New vistas on integration in the nervous system', *Trends in Neurosciences*, **15**, pp. 218–26.

Fishman, M.C. and Michael, C.R. (1973), 'Integration of auditory information in cat's visual cortex', *Vision Research*, **13**, p. 1415.

Freeman, W.J. (1999), *How Brains Make Up Their Minds* (London: Weidenfeld & Nicholson).

Goldstein, K. (1995), *The Organism: A Holistic Approach to Biology Derived from Pathological Data in Man* (New York: Zone Books).

Gurwitsch, A. (1964), *The Field of Consciousness* (Pittsburgh, PA: Dusquesne University Press).

Haenny, P.E., Maunsell, J.H.R. and Schiller, P.H. (1988), 'State dependent activity in monkey visual cortex', *Experimental Brain Research*, **69**, pp. 245–59.

Horn, G. and Hill, R.M. (1969), 'Modifications of the receptive field of cells in the visual cortex occurring spontaneously and associated with bodily tilt', *Nature*, **221**, pp. 185–7.

Hurley, S.L. (1998), *Consciousness in Action* (Cambridge, MA: Harvard University Press).

Jackendoff, R. (1987), *Consciousness and the Computational Mind* (Cambridge, MA: MIT/Bradford).

Kanizsa, G. and Gerbino, W. (1982), 'Amodal completion: seeing or thinking?', in *Organization and Representation in Perception*, ed. J. Beck (New Jersey: Lawrence Erlbaum).

Kanwisher, N. (2001), 'Neural events and perceptual awareness', *Cognition*, **79**, pp. 89–113.

Kelso, J.A.S. (1995), *Dynamic Patterns: The Self-Organization of Brain and Behavior* (Cambridge, MA: The MIT Press/A Bradford Book).

Kim, J. (1993), *Supervenience and Mind: Selected Philosophical Essays* (Cambridge: Cambridge University Press).

Köhler, W. (1947), *Gestalt Psychology* (New York: Liveright).

Leopold, D.A. and Logothetis, N.K. (1996), 'Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry', *Nature*, **379**, pp. 549–53.

Leopold, D.A. and Logothetis, N.K. (1999), 'Multistable phenomena: changing views in perception', *Trends in Cognitive Sciences*, **3**, pp. 254–63.

Logothetis, N.K. (1999), 'Vision: A window on consciousness', *Scientific American*, **281**, pp. 68–75.

Logothetis, N.K. and Schall, J.D. (1989), 'Neuronal correlates of subjective visual perception', *Science*, **245**, pp. 761–3.

Lumer, E.D., Friston, K.J. and Rees, G. (1998), 'Neural correlates of perceptual rivalry in the human brain', *Science*, **280**, pp. 1930–3.

Lumer, E.D. and Rees, G. (1999), 'Covariation of activity in visual and prefrontal cortex associated with subjective visual perception', *Proceedings of the National Academy of Sciences USA*, **96**, pp. 1169–73.

Lutz, A., Lachaux, J.-P., Martinerie, J. and Varela, F.J. (2002), 'Guiding the study of brain dynamics by using first-person data: Synchrony patterns correlate with ongoing conscious states during a simple visual task', *Proceedings of the National Academy of Sciences USA*. Early edition: www.pnas.org/cgi/doi/10.1073/pnas.032658199

McDowell, J. (1982), 'Criteria, defeasibility and knowledge', *Proceedings of the British Academy*, **68**, pp. 455–79

McDowell, J. (1994), 'The content of perceptual experience', *Philosophical Quarterly*, **44**, pp. 190–205.

Merleau-Ponty, M. (1962), *Phenomenology of Perception*, trans. Colin Smith (London: Routledge).

Metzinger, T. (ed. 2000a), *Neural Correlates of Consciousness: Empirical and Conceptual Questions* (Cambridge, MA: The MIT Press/A Bradford Book).

Metzinger, T. (2000b), 'Introduction: Consciousness research at the end of the twentieth century', in Metzinger (2000a).

Milner, A.D. and Goodale, M.A. (1995), *The Visual Brain in Action* (Oxford: Oxford University Press).

Moran, J. and Desimone, R. (1985), 'Selective attention gates visual processing in extrastriate cortex', *Science*, 229, pp. 782–4.

Morell, F. (1974), 'Visual system's view of acoustic space', *Nature*, **238**, pp. 44–6.

Myerson, J., Miezin, F. and Allman, J. (1981), 'Binocular rivalry in macaque monkeys and humans: A comparative study', *Behavior Analysis Letters*, **1**, pp. 149–59.

Noë, A. (2002), 'Is the visual world a grand illusion?', *Journal of Consciousness Studies*, **9** (5–6), pp. 1–12.

Noë, A. (2004), *Action in Perception* (Cambridge, MA: The MIT Press).

Noë, A. and O'Regan, J.K. (2000), 'Perception, attention, and the grand illusion', *Psyche*, **6** (15), URL: http://psyche.cs.monash.edu.au/v6/psyche-61-5-noe.html

Noë, A., Pessoa, L., and Thompson, E. (2000), 'Beyond the grand illusion: What change blindness really teaches us about vision', *Visual Cognition*, **7**, pp. 93–106.

O'Regan, J.K. and Noë, A. (2001), 'A sensorimotor account of vision and visual consciousness', *Behavioral and Brain Sciences*, **25** (4), pp. 883–975.

Pessoa, L., Thompson, E. and Noë, A. (1998), 'Finding out about filling-in: A guide to perceptual completion for visual science and the philosophy of perception', *Behavioral and Brain Sciences*, **21**, pp. 723–802.

Port, R.F. and van Gelder, T. (ed. 1995), *Mind as Motion: Explorations in the Dynamics of Cognition* (Cambridge, MA: The MIT Press/A Bradford Book).

Putnam, H. (1999), *The Threefold Chord: Mind, Body, and World* (New York: Columbia Univ. Press).

Rees, G., Krieman, G. and Koch, C. (2002), 'Neural correlates of consciousness in humans', *Nature Reviews Neuroscience*, **3**, pp. 261–70.

Revonsuo, A. (2000), 'Prospects for a scientific research programme on consciousness', in Metzinger (2000a).

Scheerer, E. (1994), 'Psychoneural isomorphism: Historical background and current relevance', *Philosophical Psychology*, **7**, pp. 183–210.

Sellars, W. (1956), 'Empiricism and the philosophy of mind', in *The Foundations of Science and the Concepts of Psychology and Psychoanalysis*, ed. H. Feigl and M. Scriven (Minneapolis: University of Minnesota Press).

Sheinberg D.L. and Logothetis, N.K. (1997), 'The role of temporal cortical areas in perceptual organization', *Proceedings of the National Academy of Sciences USA*, **94**, pp. 3408–13.

Singer, W. (1995), 'Putative functions of temporal correlations in neocortical processing', in *Large-Scale Neuronal Theories of the Brain*, ed. Christof Koch and Joel L. Davis (Cambridge, MA: The MIT Press/A Bradford Book).

Singer, W., Engel, A.K., Kreiter, A.K., Munk, M.H.J., Neuenschwander S. and. Roelfsma, P.R. (1997), 'Neuronal assemblies: Necessity, signature, and detectability', *Trends in Cognitive Sciences*, 1 (7), pp. 252–61.

Snowdon, P. (1980-81), 'Perception, vision and causation', *Proceedings of the Aristotelian Society*, 81, pp. 175–92.

Strawson, P.F. (1979), 'Perception and its objects', in *Perception and Identity*, ed. G.F. MacDonald (London: Macmillan).

Teller, D.Y. and Pugh, E.N. Jr. (1983), 'Linking propositions in color vision', in *Colour Vision*, ed. J.D. Mollon and L.T. Sharpe (London: Academic Press).

Thelen, E. and Smith, L.B. (1994), *A Dynamical Systems Approach to the Development of Cognition and Action* (Cambridge, MA: The MIT Press).

Thompson, E. and Varela, F.J. (2001), 'Radical embodiment: Neural dynamics and conscious experience', *Trends in Cognitive Sciences*, **5**, pp. 418–25.

Thompson, E., Noë, A., and Pessoa, L. (1999), 'Perceptual completion: A case study in phenomenology and cognitive science', in *Naturalizing Phenomenology: Issues in Contemporary Phenomenology and Cognitive Science*, ed. Jean Petitot *et al.* (Stanford, CA: Stanford University Press).

Todorovic, D. (1987), 'The Craik-O'Brien-Cornsweet effect: New varieties and their theoretical implications', *Perception & Psychophysics*, **42**, pp. 545–60.

Treue, S. and Maunsell, J.H.R. (1996), 'Attentional modulation of visual motion processing in cortical areas MT and MST', *Nature*, **382**, pp. 539–41.

Van Gelder, T. (1995), 'What might cognition be if not computation?', *Journal of Philosophy*, **XCII**, pp. 345–81.

Van Gelder, T. (1999), 'Dynamic approaches to cognition', in *The MIT Encyclopedia of Cognitive Science*, ed. R. Wilson and F. Keil (Cambridge, MA: The MIT Press), pp. 244–6.

Van Gulick, R. (1998), 'Analytic isomorphism and Marilyn Monroe', *Behavioral and Brain Sciences*, **21**, pp. 776–7.

Varela, F.J. (1984), 'Living ways of sense-waking: A middle path for neuroscience', in *Disorder and Order: Proceedings of the Stanford International Symposium*, ed. P. Livingston (Stanford Literature Series, vol. 1, Anma Libri).

Varela, F.J., Thompson, E. and Rosch, E. (1991), *The Embodied Mind: Cognitive Science and Human Experience* (Cambridge, MA: The MIT Press).

Varela, F.J., Lachaux, J.-P., Rodriguez, E. and Martinerie, J. (2001), 'The brain web: Phase synchronization and large-scale integration', *Nature Reviews Neuroscience*, **2**, pp. 229–39.

Weiskrantz, L. (1997), *Consciousness Lost and Found: A Neuropsychological Exploration* (Oxford: Oxford University Press).

Paper submitted April 2002