Workspace and Sensorimotor Theories

Complementary Approaches to Experience

Jan Degenaar and Fred Keijzer

Faculty of Philosophy
University of Groningen
Oude Boteringestraat 52
9712 GL Groningen, the Netherlands.
Email: :j.degenaar@rug.nl; f.a.keijzer@rug.nl

Abstract: A serious difficulty for theories of consciousness is to go beyond mere correlation between physical processes and experience. Currently, neural workspace and sensorimotor contingency theories are two of the most promising approaches to make any headway here. This paper explores the relation between these two sets of theories. Workspace theories build on large-scale activity within the brain. Sensorimotor theories include external processes in their explanations, stressing the sensorimotor contingencies that arise from our interaction with the environment. Despite the basic differences, we argue that workspace- and sensorimotor theories are complementary rather than competitive. By combining these theories, a number of problems that hamper these individual theories may be overcome and their strengths combined: Workspace theories have more to offer for explaining how there can be consciousness in the first place, while sensorimotor theories are strong in making sense of the specific phenomenal character of experiences.

Keywords: consciousness; sensorimotor theory; workspace theory; experience; absolute gap; comparative gap; internalism; externalism.

1. Absolute and comparative gaps

There is little doubt that conventional scientific approaches are able to find reliable correlations between neural activity and conscious experiences. The challenge is to provide more than mere correlations. We discuss the merits of two of the most promising proposals for increasing our understanding of experience currently available: neural workspace theories – or workspace theories for short – and sensorimotor contingency theories – or sensorimotor theories for short. Both come in various forms, and are sets of theories rather than specific theories. We will focus on the commonalities within these two sets of theories to assess the potential for combining the proposals.

In this paper we will assume that the approaches we discuss can increase our understanding of the physical basis of experience in a way that goes beyond mere correlation. It has been argued that conventional scientific approaches are unable to truly explain why physical processes are accompanied by phenomenal experience, a problem known as the hard problem of consciousness (Chalmers, 1995). Rather than making any direct claims on to the hard problem, our focus will be on the increase in explanatory power that a combination of theories may yield compared to the individual sets of theories.

Neural workspace theories provide a set of closely related theories, which seem promising to make significant headway toward a satisfactory empirical theory of the physical bases of experience (Baars, 1988; 2002; Tononi and Edelman, 1998; Dehaene & Naccache, 2001; Varela et al., 2001). Neural workspace theories hypothesize that conscious experience depends on a coherent pattern of neural activity that facilitates the availability of information for various processes. The mechanisms that give rise to this pattern of activity could potentially account for important features of experience. The hypothesis may even come to explain why some neural activity is correlated with consciousness while other neural activity is not.

Sensorimotor contingency theories hold that the phenomenal quality of experiences can be understood in terms of the characteristic relations between sensory input and motor action – the 'sensorimotor dependencies' or '-contingencies' (O'Regan & Noë, 2001a, b; Hurley & Noë, 2003; O'Regan, Myin & Noë, 2005; Mossio & Taraborelli, 2008). Within these theories, differences between for example visual and auditory experiences are thought of as differences in the sensorimotor dependencies, or differences in the mode of active exploration of the environment. One of the characteristics of visual experience is the way in which movement of the head enables us to look behind objects. While eye movements thereby result in large shifts of the retinal image, we experience the visual world as stable. From findings like these, sensorimotor theorists conclude that our experience is not an inner construct based on input alone, but is rather constituted by displaying the implicit knowledge of the input-output relation.

Much of present-day theorizing about conscious experience, including workspace theory, tends to focus on neural activity. Sensorimotor theories, in contrast, focus on whole patterns of interaction involving brain, body and environment.[1] The idea is that

1 While sensorimotor theory has affinity with Velmans' idea that the world as-perceived is out-there (e.g. Velmans, 1990), we like to emphasize the difference in focus. Velmans' 'reflexive' model of experience is inclined to take a brain-focused approach to the processes underlying experience, taking the contents of some experiences as a 'projection' in space by the brain (Velmans, 1990; 2007). Sensorimotor theories stress that an understanding of experience requires a focus on whole patterns of

knowledge of neural activity alone may not be enough for an explanation of conscious experience. We will refer to this claim of sensorimotor theories as *externalism* about the processes underlying consciousness and contrast it with the brain-focussed *internalism* of workspace theories. For a more elaborate discussion of internalism and externalism with respect to the processes that figure in the explanation, see Hurley (forthcoming). For present purposes, it is important to distinguish the explanatory externalism of sensorimotor theories from the more familiar representationalist commitment about the content of experience, namely that we experience aspects of the external world (e.g. Dretske, 1995). The explanatory externalism of sensorimotor theories holds that the *processes* we need to take into account to explain our experiences extend into the world, and this does not necessarily imply that we experience objective features of external objects.

Sensorimotor- and workspace theories are both major players in present-day consciousness research. Still, little work has been done on their relation. The defenders of sensorimotor theories have sometimes stressed the differences with brain-based explanations rather than searching for ways to link sensorimotor theory with explanatory paradigms like workspace theories. This differentiation has been useful to emphasize the particular and independent contribution of sensorimotor theories to the explanation of consciousness. At the same time, workspace theorists may have been disheartened by the externalist tendencies of sensorimotor theories. They may have seen little reason to relate their dominant brain-based theory to this set of relatively new theories. The general differentiation between internalist and externalist explanations (Hurley, forthcoming) – according to which both theories can be categorized – may have given the impression of intrinsic opposition. We claim that this impression is false.

At the basis of our claim lies the distinction between two fundamental problems of experience (Chalmers, 1996, p. 5; Hurley & Noë, 2003). The first problem is to understand the very *existence* of conscious experience, generally known as the *absolute gap*: Why do we have conscious experience at all? The second is to understand the *character* of conscious experience: Why do experiences have the specific qualities that they have? This problem is known as the *comparative gap* or -*gaps* (Hurley & Noë, 2003). Examples are the problem to explain experiential differences between sensory modalities (e.g. seeing versus hearing) and within modalities (e.g. the experience of red versus the experience of blue). Explaining consciousness consists at least in solving both the absolute and the comparative gap problems.

In this paper we will argue that sensorimotor- and workspace theories can each be cast as best dealing with one of the explanatory gaps. If this is right, a combination of the theories becomes highly desirable. We discuss a potential difficulty for reconciliation of the theories: Workspace and sensorimotor explanations have respectively strong internalist and externalist tendencies, which reflect deep differences in theoretical and metaphysical views. As a way to deal with these differences, we sketch three different scenarios for combining workspace- and sensorimotor theories, each reflecting different fundamental outlooks on experience. In all three cases, combining workspace- and sensorimotor theories plausibly leads to an increase of explanatory strength compared to each of the separate theories.

---

interaction with the (third-person identifiable) environment. Whether the physical constitution of experience is purely brain-based is a further question, which we touch upon in section 4.2 below.

2. Sensorimotor contingency theories of experience

Sensorimotor theories offer an empirical approach of the character of conscious sensory experience. The theories aim to answer the question that Chalmers stated thus: 'Given that conscious experience exists, why do individual experiences have their particular nature?' (Chalmers, 1996, p. 5). To answer this question, sensorimotor theorists focus on whole loops of interaction involving brain, body and world. Since our intuitions may need some stretching before we can appreciate an explanatory role for processes outside the brain, some preliminary remarks will be useful.

2.1. Not by neural activity alone

Hurley and Noë state that 'neural properties are qualitatively inscrutable' (Hurley & Noë, 2003). Indeed, it seems doubtful that the character of the experience of red or the taste of coffee can ever be explained in neuroscientific terms (Levine, 1983, 1993). Nevertheless, most people are convinced that experience occurs within the brain and that it should be explained in terms of neural activity. Hurley observes conflicting intuitions:

> If someone really has no conception of how neural or internal functional properties—or indeed any others—could explain phenomenal qualities, then how can he be so confident that if phenomenal qualities can be explained, it must be internal factors that do the job? (Hurley, forthcoming).

Given the difficulty to understand the character of experiences in terms of inner states and processes, we may need to challenge the internalist assumption. After all: 'Neural processes are normally in continuous dynamical interaction with external factors; there's nothing magical about the boundary between them' (Hurley, forthcoming). As there is no 'magical membrane' to separate the brain from the rest of the world, there is also no reason for seeking the explanation of consciousness in the brain alone, the externalist argues. Thus, the intuition that consciousness must be something within our bodies or even brains can be neutralized by an opposing intuition that consciousness is not intrinsically related to anything within the body.

When intuitions on their own are not sufficiently trustworthy, what kind of evidence do we have for an internalist interpretation of consciousness? Some theorists argue that the existence of rich experiences which are seemingly 'off-line' does provide such evidence (e.g. Koch, 2004; Revonsuo, 2006; Prinz, 2008). As experience is possible in relative isolation from the environment, one can argue that brain processes suffice for consciousness while the environment can only modulate these inner processes. In reply, we will turn to the phenomenon of dreaming, which is a show case example of this line of argument.

While dreaming may seem to provide an obvious example of strictly inner experience, the case is actually not clear cut. First, the existence of rich off-line experience does not imply that the best explanation of on-line sensory experiences must be based on internal processes alone. As Hurley says: 'If the enabling role of internal simulations in off-line cases is derivative from their role in extended dynamics, it provides no reason to hold that only internal processes can do quality-enabling work in the primary, on-line cases' (Hurley, forthcoming). The explanation of the qualitative character of dreaming experiences may ultimately derive from the explanation of sensory experiences rather than the other way around. Thus a purely neural account can be incomplete even for illusory cases and a weak form of externalism may be needed for dreaming. Second, on-line and off-line experiences

may be qualitatively different (Putnam, 1999, p. 130; Noë, 2004, p. 213-214; Noë & Thompson, 2004). Waking experiences do not typically have 'a dream-like quality' (Austin, 1962, p. 48-49). As a result, the processes that we need in order to explain the experiences may be different too, even if they partly overlap. Third, even though dreaming takes place without outward action, this does not imply that input from the environment or feedback loops extending in the body play no role. Our brains are never completely off-line as we can be wakened by noise, shaking or other stimuli. Maybe we need to take active external processes into consideration even to account for the phenomenal contents of dreams.

In addition, a positive reason for taking the externalist possibilities seriously comes from the general trend towards more dynamical, embodied and embedded explanations of cognition (e.g. Clark, 1999; Calvo & Gomila, 2008; Robbins & Aydede, 2008). Instead of decomposing dynamically interacting systems into different parts, more global and environmentally extended patterns of interaction may be involved in the workings of cognitive processes (Clark, 1997; 1999) and behaviour (Keijzer, 2001). Clark and Chalmers (1998) even introduced the notion of an *extended mind* that is literally extending into the environment. If externalism can be accepted for the processes underlying cognition, we see no reasons for a general ban *beforehand* on developing similar ideas for experience. Hurley and Noë suggest: 'To find explanations of the qualitative character of experience, our gaze should be extended outward, to the dynamic relations between brain, body, and world' (Hurley & Noë, 2003). Let us turn to the sensorimotor hypothesis and see what this brings.

2.2. Sensorimotor theory: experience in interaction

Sensorimotor theories of perceptual experience state that experiencing is best characterized as exercising our mastery, or implicit knowledge, of patterns of sensorimotor interaction with the environment (O'Regan & Noë, 2001a, b; Hurley & Noë, 2003; Noë, 2004; O'Regan, Myin & Noë, 2005). To unpack this idea, we will successively discuss the role of action, the role of the environment, and the appeal to implicit knowledge. For purposes of illustration we will mention evidence from sensory substitution.

A basic idea of sensorimotor theories is that sensory experiences are in general strongly action-dependent. If retinal stimulation shifts in the absence of eye-movements, typically the world would appear to move (or the perceiver would appear to move). However, if the same pattern of retinal stimulation occurs as the result of an eye-movement, this does not impair the apparent stability of the visual world (or perception of the perceiver's own location). Thus, action has more than an instrumental role in experience; it does not 'merely' change the input: it can directly change visual experience (Hurley, 1998). Experience is not an inner construct based on input alone.

To capture the qualitative character of sensory experiences an appeal is made to the specific way in which sensory input depends on motor action. A clear example is the experience of the softness of a sponge. According to sensorimotor theories we can understand the character of this experience if we consider the sensory consequences of motor action. For example, when we squeeze a sponge it gives little resistance and this is what its softness consists in. The experience of the hardness and softness of surfaces is not action-neutral; it rather consists in our grasp of the sensory effects of our actions (O'Regan, Myin & Noë, 2005). These dynamical patterns of sensorimotor contingencies are specific to the particular senses; feeling a sponge and seeing one

have their own, different contingencies. Sensorimotor theories aim to explain the experiential differences between modalities in this way.

Appealing to sensorimotor dependencies brings the environment into the explanation. The experience of sponge-squeezing is explained by the characteristic ways in which the sponge responds to pressure. Similarly, to understand the visual experience of objects in space, the relevant sensorimotor contingencies obtain as a result of the spatial orientation of objects and the reflective behaviour of light. For example, the distance one has to move to look behind an object depends on the relative distances of objects and perceiver. Sensorimotor theories claim that we experience the spatial relations between objects by exercising the implicit grasp of the sensory consequences of movements. For this reason, we need to take the environment into account if we are to understand perceptual experience.

Of course, sensorimotor theory does not require that we must always move in order to experience. We often see at a glance that one object is closer than another. By involving our implicit knowledge, sensorimotor theories can deal with experience in the absence of movement. Once we have the practical familiarity with the sensorimotor contingencies, we make use of this familiarity in our experience of the world. Whether we move or not, we implicitly grasp what sensory consequences are to be expected if we would make a certain move. For example, when we see the spatial orientation of objects, this experience constitutively depends on our mastery over the governing laws of sensorimotor contingency (O'Regan & Noë, 2001a). Thus, according to sensorimotor theories, experiencing is a skilful activity of sensorimotor interaction with the world (O'Regan & Noë, 2001a; Myin & O'Regan, 2002).

A good example that illustrates how sensorimotor theory goes beyond a brain-based focus comes from studies of sensory substitution devices, such as Bach-y-Rita's tactile-visual substitution system (Bach-y-Rita, 1984; 2002). This system transforms the image recorded by a camera into a tactile display, e.g. an array of vibrating pins which can be applied to the subjects back. In studies with such a devise, otherwise blind persons report vision-like experiences rather than tactile ones: They experience objects as being at a certain distance and they report experiencing spatial relationships between objects, such as that 'one is partially blocking the view' of another object. In a recent study using a different sensory substitution devise, an auditory-visual substitution system, Auvray et al. (2007) tested blindfolded sighted subjects, finding that in some cases subjects reported visual experiences despite the auditory input. Importantly, when subjects have no control over the camera – when someone else controls the camera, or when it is stationary – the change from tactile to semi-visual experience does not occur (Bach-y-Rita, 1984). This fits very well with sensorimotor theory as in this case there are no new sensorimotor contingencies to be mastered by the subject, but only passive sensory stimulation. The subjects will not acquire the practical mastery of the sensorimotor contingencies relevant to 'seeing' with the device (Hurley & Noë, 2003).

Bach-y-Rita concluded from his findings that 'we do not see with our eyes, but with our brain' (Bach-y-Rita, 2002, p. 497). However, referring to the brain does not explain the differences enabled by the device, as we can also be said to hear, feel or taste with our brain. In contrast, the change in the sensorimotor contingencies does explain why a tactile modality acquires vision-like experiential features (Hurley & Noë, 2003).

Sensory substitution provides an example of how sensorimotor contingencies can help to explain differences in the qualitative feel of sensory modalities, and possibly even how new modalities can arise (Auvray & Myin, 2009). The sensorimotor

interactions between an agent and its environment provide a systematic constraint on experience. Brain functioning is shaped by our active encounters with the environment, and sensorimotor theorists stress that it should be considered in the context of these temporally extended patterns of interaction.

2.3. Explanatory promise and limitations

Given this short description of sensorimotor theories, what can be said about their strengths and weaknesses? In particular, how does sensorimotor theory relate to the absolute and comparative gaps of consciousness?

Sensorimotor theories seem particularly strong on comparative gap issues, such as the experiential differences between different sensory modalities (Hurley & Noë, 2003). Sensorimotor theories may even provide a handle on dealing with new sensory modalities (Auvray & Myin, 2009), and some of the aspects of differences within modalities (Hurley & Noë, 2003), such as between colours (Philipona & O'Regan, 2006). In addition, a sensorimotor perspective has been applied to differences between conscious thought and sensory experiences (O'Regan, Myin & Noë, 2005). Sensorimotor theory arguably provides explanations that are more than 'mere correlation'. From a sensorimotor perspective one can understand why, e.g., visual experiences differ from tactile experiences in the way they do. In contrast, neural activity in a certain area may be reliably correlated with a particular experience, but this correlation would not explain why a particular phenomenal experience is associated with this activity. Rather than accepting intrinsically qualitative properties of neural activity with all its difficulties, sensorimotor theories characterize differences in experiences as differences in the dynamical patterns of agent-world interactions. In this way, sensorimotor theory also opens up the possibility of an evolutionary understanding of the origins and differences between sensory modalities.

Sensorimotor theorists have occasionally made claims concerning the absolute gap, stating that they explain the *presence* of experience (see especially O'Regan & Noë, 2001a, pp. 1011-1012). However, as also Noë (2004, pp. 228-31) recognizes, sensorimotor theory seems less convincing as an account of the switch from non-conscious processes to conscious ones. Sensorimotor contingencies are used by widely different systems, ranging from organisms like insects to robots, not all of them plausibly interpreted as experiencing beings. In addition, sometimes full-blown conscious experience is not present in humans despite the exercise of mastery of sensorimotor interaction with the environment. For example, when driving, absorbed in thought or conversation with a friend, you may hardly experience aspects of the environment that are used to guide your behaviour. Thus, making use of sensorimotor contingencies in itself does not seem to suffice for the presence of experience.

In their reply to the 'unconscious' driving problem, O'Regan and Noë agree that an extra ingredient is required. They write:

> A driver (…) would be said to be aware of a red traffic light if, in addition to the mastery of sensorimotor contingencies associated with the red light, his attunement to these sensorimotor contingencies is integrated into his planning, rational thought or linguistic behavior. Depending on the extent to which the seeing of the red light is incorporated into his planning or thought, the driver would be said to be aware of the red light to varying degrees (O'Regan & Noë, 2001b, p. 94).

It may be that sensorimotor dependencies are only relevant to consciousness if they play a certain role in planning, thought or speech. However, this rather seems to

describe the presence of consciousness than to explain it. As a theory of the presence of experience, sensorimotor theory seems too descriptive. As also Rowlands (2003) and Noë (2004) suggest, sensorimotor theory may not in itself be able to explain the presence of experience, but it rather has to presuppose it.

Thus, it appears that sensorimotor theory is well equipped to deal with comparative gap problems that arise once consciousness is present and enables one to make sense of various qualitative differences. However, it seems more difficult to address the absolute gap with the means provided by sensorimotor theory.

3. Workspace theories of consciousness

Conscious experience implies the availability of information: You will be able to tell when you are conscious of a stimulus, at least if you have the capacity to speak. Neural workspace theories form a set of theories that aim to identify the underlying neural mechanisms that can explain the conscious availability of information.

The central idea of neural workspace theories is that consciousness-correlated neural activity forms *a coherent pattern of neural activity that makes information globally available throughout a neuronal workspace* (Baars, 1988; 2002; Tononi & Edelman, 1998; Dehaene & Naccache, 2001; Varela et al., 2001; Metzinger, 2003).[2] This 'workspace' is 'a central information exchange that allows many different specialized processors to interact' (Baars, 1988, p. 43). By hypothesis, information that reaches this workspace will influence the processing in large parts of the brain, a bit like the information on a blackboard being available for the whole class.

Workspace theories claim that a stimulus will influence conscious experience if and only if it modulates the activity in the neural workspace. Based on criteria of availability of information for the person, typically the availability for verbal report, it offers a theory of the neural basis of experience. This theory is usually formulated in terms of the availability of information to subsystems. It builds on the subpersonal availability of information that is implied in workspace activity and, closely related, the subsequent *role* that the activity plays. This sits well with Dennett's philosophical views on consciousness. As in Dennett's (1991) theory, workspace theories approach consciousness not as an intrinsic feature of neural activity: Rather it is because of the use that is made of information that it classifies as conscious. Indeed, a neural workspace can flesh out Dennett's idea of consciousness as 'fame in the brain' (Dennett, 2001).[3] However, note that workspace theory is not necessarily committed to such an interpretation. Although the theory is based on behavioural criteria of personal-level availability, this by no means excludes the possibility to acknowledge intrinsic experiential features of neural activity. The neural workspace can be – and sometimes is – seen as 'the place where consciousness happens'.

Note that availability of information for perceptual report is not a simple criterion to judge the presence of experience. As Metzinger (2003, p. 75) points out, consciousness may come in degrees, depending on the extent to which information becomes available. He also differentiates between three dispositional properties that can exemplify availability; information can be available for guided attention, for

[2] Neural workspace theories differ in their stress on neural or informational aspects, some being almost exclusively formulated in informational terms, like Baars' original formulation, others in neural terms, as in Varela et al. Accounts that are focused on neural processes are often given an informational interpretation, as in Tononi and Edelman's models. The similarities are strong as the neurally oriented theories remain committed to some form of information processing, while those in informational terms presume neural information processing mechanisms for their implementation.

[3] In as far as Dennett's positive theory of consciousness occupies the same explanatory niche as workspace theories, it may have similar strengths and weaknesses.

cognitive processing, and for behavioural control. Thus, to study the neural basis of consciousness we may have to ask exactly in what sense information becomes available (Metzinger, 2003, p. 124). Another difficulty is that availability is not always sufficient for experience: Information may in some cases be available only unconsciously – as in the case when you respond adequately but unconsciously to a stimulus. However, when we put problematic cases aside there will remain enough reasonably uncontroversial cases that can be used in this empirical approach of experience.

3.1. Consciousness as global cortical activity

There is ample evidence that cognitive processes often occur without associated conscious experience (Dehaene & Naccache 2001; Merikle & Daneman 1999). The challenge is to determine whether there is a systematic difference between consciousness-correlated and not consciousness-correlated processing (Dehaene & Naccache, 2001). Some neural activity will be specifically associated with experience and the question is how this activity differs from the activity that isn't.[4] Neural workspace theories offer a possible answer.

A good example of a workspace hypothesis is provided by Tononi and Edelman's (1998) 'dynamic core hypothesis'. Interestingly, they start from the *character* of experience, to hypothesize on the nature of the neural processes that underlie consciousness:

> (…) our strategy is to characterize the kinds of neural processes that might account for key properties of conscious experience. We emphasize two properties: conscious experience is integrated (each conscious scene is unified) and at the same time it is highly differentiated (within a short time, one can experience any of a huge number of different conscious states) (Tononi & Edelman, 1998, p. 1846).

The *dynamic core hypothesis* proposes that which neurons are part of the 'dynamic core' can rapidly change and that the 'dynamic core' is the neural activity that correlates with consciousness. Tononi and Edelman hypothesize that the unity of experience can be explained by the 'functional integration' of the relevant neuronal activity: 'at a given timescale, these elements interact more strongly among themselves than with the rest of the system' (Tononi & Edelman, 1998, p. 1848). The differentiated character of experience is proposed to be reflected in the 'complexity' of the activity of the dynamic core, which is a function of the amount of mutual information that subsets of the dynamic core share with the rest of the core (Tononi & Edelman, 1998).

The intensive 'cross-talk' between the neurons within the dynamic core, or more in general within a neural workspace, should ensure that each part of this workspace is influenced by the other parts. If a part of the workspace activity carries information about a certain aspect of the environment, this results in the global influence of this environmental feature throughout the workspace. According to a representationalist analysis (e.g. Metzinger 2003), it results in the availability of this information for other subsystems. The workspace activity could potentially stretch out to areas

[4] Some theorists, in contrast, have suggested a link between life and experience (e.g. Noë, 2004; Thompson, 2007). For example Noë speculates that 'living beings are already, by dint of being alive, *potentially conscious*' (Noë, 2004, p. 230). Workspace theories address the contrast between conscious and unconscious processes *within* living beings.

devoted to speech so that perceptual reports can come under the influence of the environmental feature that modulates workspace activity.

Presumably, neurons throughout large parts of the brain can be part of the workspace. There are various hypotheses on the specifics of the workspace. For example, Dehaene and Naccache (2001) propose that specific 'workspace neurons' with long-distance connectivity form a neural workspace. If such neurons are sufficiently activated, they will result in brain-scale coherent activity that makes information available throughout the workspace. Tononi and Edelman (1998), in contrast, offer a more dynamical view in which it is possible that at one moment a neuron is strongly activated without being part of dynamic core, while at another moment it is part of the dynamic core. Several authors have argued that in addition to widespread cortical neurons also neurons in the thalamus may be involved (which is consistent with the re-entrant connections in the thalamocortical system) (Tononi & Edelman, 1998; Dehaene & Naccache, 2001).

Much has been written elsewhere on the evidence that is in agreement with workspace hypotheses (Dehaene & Naccache, 2001; Baars, 2002; Varela et al., 2001). An important finding is that *neural synchrony* is correlated with conscious experience (Engel et al., 1999; Engel & Singer, 2001; Varela et al., 2001). When something is consciously experienced, an associated increase of the synchronous firing of neurons may be found throughout large parts of the brain. This indicates a high level of interaction between neurons, which is exactly what workspace theory predicts.

3.2. Explanatory promise and limitations

How well does workspace theory fare with respect to the absolute and comparative gaps of consciousness? We submit that since workspace theories aim to clarify the preconditions for conscious experience, they are chiefly focussed on issues relating to the absolute gap.

Workspace theories aim to account for the presence of consciousness by explaining how the neural activity that underlies consciousness differs from activity that is not directly involved in consciousness. The basic idea is that certain ways of responding to a stimulus implicate awareness of it: When you can report that you heard a noise, you have experienced it. Neural workspace theories aim to explain features that are descriptive of experience. They do so in terms of underlying neural activity.

In their account of the mechanisms that enable persons to consciously perceive the world, workspace theories tend to use a subpersonal notion of information and they speak of the availability of information for parts of the brain. In particular, it is proposed that the workspace forms part of the subpersonal mechanisms by means of which information about the environment can become available for the subject. A way to construe an informational interpretation of subpersonal processes is as a third-person, correlation-based ascription of information. Note that such an ascription of information plausibly depends on more than neural activity alone, indeed it may only make sense in the larger sensorimotor context.

Workspace activity could make perceptual information globally available for thought and action. If a stimulus influences workspace activity, this can directly influence behaviour and information related to the stimulus will be available for report (Dehaene & Nacacche, 2001 pp. 21-22). The neural processes involved in the workspace will have a special influence on further lines of thought, action, and speech, as a direct result of the physical/functional properties of the workspace. Unconscious workspace activity is impossible, because (above a certain level) workspace activity just leads to global availability of information for neural

subsystems, and this implies that the information is available to modulate behaviour such as verbal report. (One way to construe the subpersonal 'availability of information' is as shorthand for e.g. parts of the brain coming under the influence of a certain aspect of the environment to which the person is sensitive.) Thus, workspace theories clearly aim to go beyond mere correlation: They potentially provide an explanation in which the presence of particular brain processes *implies* key psychological features signifying conscious availability.

At the same time, a detailed account is still lacking of how workspace activity can have the specific effect it has. For example, it remains to be clarified how 'speech centres' are influenced by visual processes correlating to the presence of a butterfly in such a way that it enables the person to tell others that there is a butterfly. While such an account is missing, we submit that the most important explanatory promise of workspace theory concerns the contrast between consciousness-correlated and not consciousness-correlated processes. This contrast is addressed in terms of differences in subpersonal availability, or 'fame in the brain'.

Workspace theorists do have also aimed to address the character of experience – the comparative gaps. For example, Tononi and Edelman (1998) suggested that properties of the dynamic core could account for the differentiated character of experience – the higher the complexity of the workspace activity, the more differentiated the experience. In addition, more daring attempts have been made to address the specific character of sensory experiences (Edelman & Tononi, 2000; Tononi, 2004). To approach comparative gap problems such as the differences between visual and auditory experiences, these authors appeal to the 'discriminations' made within the dynamic core. One difficulty with this attempt concerns the required understanding of the neural activity in terms of 'discriminations'. Suppose this understanding is dependent upon the theorist's knowledge of the larger pattern of interaction with the environment within which the discrimination plays its role. In that case, the understanding of experience is based in this larger pattern rather than in the workspace activity. Another difficulty is that there seems to be no intelligible link between the specific phenomenal character of experience and a set of discriminations. It is far from clear why a certain set of discriminations would result in a visual experience rather than an auditory experience, and how the experience of blue can be characterized by the way blue stimuli can be discriminated from other stimuli. Focussing on neural activity alone makes it hard to see why this activity is associated with particular experiences.

This problem becomes aggravated when workspace theories are compared to sensorimotor theories on this count. The latter theories help to understand why and even how vision and touch constitute different forms of experience in a way that is not available to workspace theories.

To conclude, workspace theories are well-equipped to address the neural mechanisms that underlie conscious experience. They can potentially explain the difference between consciousness-correlated and not consciousness-correlated processes – a difference that is important to deal with the absolute gap of consciousness. However, the specific character of experience seems to be less approachable in terms of workspace activity.

4. Compatibility of the theories

Workspace- and sensorimotor theories thus have complementary strengths and weaknesses with respect to the absolute and comparative gaps. Given this circumstance it would seem to be a good strategy to combine the two theories, turning

them into a unified framework that keeps the strong aspects of both theories and applies them to the separate gaps. However, workspace theories and sensorimotor theories are usually seen as competitors that aim to provide *different* explanations for human experience.

The division of labour suggested here is not a standard interpretation. Some defenders of both theories have claimed to address both the absolute and the comparative gaps. For example, O'Regan and Noë (2001a) enter the natural territory of workspace theory when they claim that sensorimotor theory explains the existence of experience. – They suggest that sensorimotor theory solves the problem of the absolute gap, in as far as experience is constituted by its qualities (O'Regan & Noë, 2001a, pp. 1011-1012). – The opposite also holds true: Workspace theory is sometimes used in an attempt to make sense of the specific character of experiences. For example Edelman and Tononi (2000) and Tononi (2004) suggest that the quality of experiences can be thought of as the discriminations that are made within the dynamic core.

However, against these claims, it should be noted that there is no intrinsic theoretical need to apply either of these theories to both gaps. As we argued above that in both cases the weak aspects of the one theory coincide with the strong points of the other, dividing up the territory in a combined effort is beneficial. In our view, a more fundamental difficulty for a profitable combination of workspace- and sensorimotor theories comes from deep differences in theoretical outlook involved and the very interpretation of what an explanation of consciousness amounts to. This is particularly so for the issue of localization: the question which processes, if any, are constitutive of consciousness.

In the following section, we will first discuss the issues at stake in combining internalist and externalist explanations. In section 4.2, we will turn to the issue of localization and we will sketch three different scenarios for a combination of workspace- or sensorimotor theories.

4.1 Internalist and externalist explanations

Are the ways in which workspace- and sensorimotor theories aim to explain different aspects of experience compatible? The first issue at stake concerns the commitments to respectively internalist and externalist forms of explanation. Should we approach conscious experience as something that takes place inside the head or not?

From the perspective of a workspace theorist, the discrepancy between internalist and externalist methodologies may at first not be so obvious: No one ever denied that workspace activity is embedded in a body interacting with an environment and that this impinges on consciousness, even if only via sensory input. However, sensorimotor theories envision a much more important role for the interaction with an environment than merely 'impinging'. Indeed, sensorimotor dependencies are cast as the *key* feature behind consciousness. The sensorimotor dependencies that arise from the interaction with the environment are ongoing shaping factors for brain processes. Without these shaping factors, there would not be consciousness as it occurs in normal human beings. Sensorimotor theorists claim that the explanation of consciousness needs to build on the dynamical patterns of sensorimotor interaction with the environment, and that experiencing is best seen as an activity (O'Regan & Noë, 2001a). From this perspective, a purely brain-focussed theory does not even address the problems that should be addressed. Can these internalist and externalist forms of explanation be reconciled?

A first positive reflection is that ultimately workspace activity is of course part of a larger pattern of interaction. And as sensorimotor explanations appeal to more extensive patterns of sensorimotor interaction, workspace theory can be cast as a subset of *neural* processes that co-constitute this interaction. Influencing the sensorimotor aspects of experience may be closely coordinated with workspace activity. The latter could even be cast as the very mechanism that makes sensorimotor dependencies relevant to planning, thought and language, as required by O'Regan and Noë (2001a). Thus, if both theories are correct, then those patterns of sensorimotor contingencies that underlie our experience – that is those over which we are actively exercising our mastery – are those patterns that involve workspace activity.

Against this positive reading, workspace theorists sometimes claim to characterize the processes that are *directly* involved in consciousness, while external influences work only to the extent that they impinge on the workspace. This may be read as contradicting the idea that the whole pattern of interaction with the environment is intrinsically relevant for the character of consciousness. However, we think this tension can be eased. Instead of casting the workspace as a central inner conscious domain, separated from peripheral processes that are not directly involved in consciousness (figure 1.a), workspace theory allows a different interpretation in which the connections of the workspace with certain sensorimotor processes are co-constitutive of the activity directly relevant to consciousness (figure 1.b).[5] Note that our point right now is only that this is a *possible* reading of workspace theory, which would lead to a possible combination of the explanatory means available to both sets of theories.

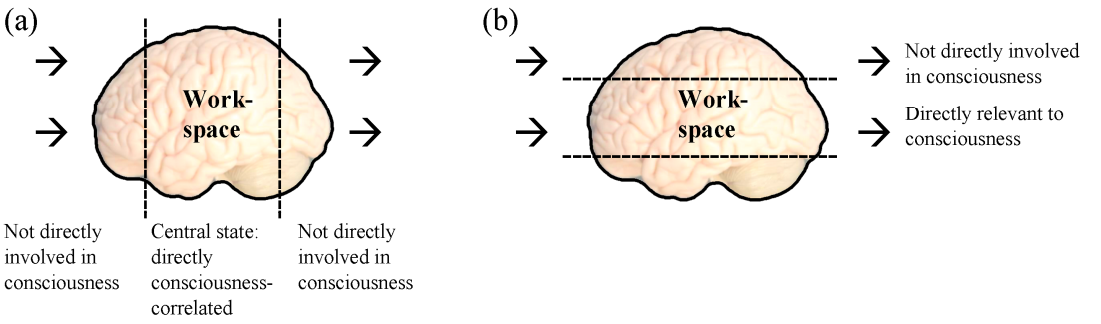

Figure 1. Schematic representation of two interpretations of the neural workspace. (a) Consciousness is interpreted as a central state. Peripheral and environmental processes are seen as not directly involved in consciousness. (b) The processes directly involved in consciousness can include processes outside the workspace. Patterns of input-output coordination that do not involve the workspace are not directly involved in experience.

The plausibility of such a combination of the theories is reinforced by the fact that a workspace approach can easily accommodate the basic *action-dependence* of experience that is stressed by sensorimotor theory. Workspace theories are informational theories of neuronal dynamics, and as such they are not committed to a particular conception of the origins of the integrated patterns of neural activity. Nothing requires workspace activity to be under influence of input alone. Indeed,

[5] In effect, the distinction is between a vertically modular 'sandwich model' reading and a dynamical reading in terms of horizontal modularity of the processes directly involved in consciousness, similar to Hurley's (1998) horizontal/vertical modularity distinction.

some work has recently been done on the integration of sensorimotor dynamics with a workspace perspective. For example, Shanahan (2006) proposes a model in which internal simulation of sensorimotor interaction with the environment is implemented within a workspace framework. Functional integration of a workspace with sensorimotor dynamics therefore seems not particularly problematic.

Some defenders of workspace theories have already stressed the need for a strong link between neural dynamics and sensorimotor coupling. Thompson and Varela (2001), for example, propose that the neural substrate of consciousness consists of 'large scale dynamical patterns' of neural activity, and that these should be considered in a broader context of sensorimotor coupling with the environment. They even suggest that 'the processes crucial for consciousness cut across brain-body-world divisions, rather than being brain-bound neural events' (Thompson & Varela, 2001). Even though Thompson and Varela did not specifically develop a sensorimotor theory, this general view is very congenial to the combination of workspace and sensorimotor theories that we propose. To conclude, from a practical point of view, there do not seem to be any intrinsic problems that forbid a combination of workspace- and sensorimotor explanations.

However, while it seems that there are possibilities for combining the explanatory focus of both sets of theories, so far we have glossed over more fundamental disagreements concerning the different interpretations of experience. For example, Deheane and Naccache (2001) identify experience with neural (workspace) activity. O'Regan and Noë (2001a, b) reject this identification, insisting that experience rather consists in a way of exploring the environment. These are deep differences in viewpoint that are difficult to reconcile. At the level of such fundamental theoretical commitments a common ground has to be found for a successful reconciliation of the theories.

4.2 Fundamental differences sorted into three scenarios

Workspace theories tend to – but are not necessarily committed to – the identification of consciousness with neural processes. Sensorimotor theories on the other hand tend to oppose this identification. A way to highlight the fundamental differences in outlook between theorists is by casting them as ideas on the localization of the processes that constitute consciousness. We will discuss three different positions on this issue, namely internal localization, external localization and no localization. We do not aim to take position here on this fundamental issue of the constitution of consciousness. Instead, we will argue that from all three positions there are systematic benefits to be had by combining both sets of theories, even when the ways in which these unions take form can be very different under these three interpretations.

The first option for a combination comes under the assumption of *internal localization*, the standard background assumption of many workspace theorists. Experience is here conceptualized as a neural process, or more specifically workspace activity. Under this assumption, a combination with sensorimotor theory would involve the explicit articulation of the systematic links between workspace activity and sensorimotor contingencies. The latter can be interpreted in terms of their direct and indirect impact on workspace activity, the place where experience comes about in this interpretation. If sensorimotor theory correctly identifies differences in the character of experiences, the processes that the internalist claims to constitute consciousness had better make appropriate contact with the patterns of sensorimotor dependencies. The systematic differences in workspace activity brought about by the contingencies of different sensorimotor modalities would allow the explanation of

comparative gaps in a way that goes beyond workspace theory. We could call this the *sensorimotor workspace hypothesis*, because ideas of sensorimotor theory are assimilated in a workspace framework, which is improved upon by this combination.

This option should be congenial for those who think that consciousness must ultimately be a brain-process. Others will see important drawbacks: Even if one accepts that the experiential relevance of sensorimotor contingencies operates through their influence on neural processes, this does not make these contingencies external to experience. Sensorimotor theorists Hurley and Noë say:

> Qualitative character may supervene on neural properties even if the qualitative expression of neural activity is determined, as we have argued, by dynamical sensorimotor considerations. (…) But if both claims are true, we hold that our account is explanatory in a way that the neural supervenience claim is not (Hurley & Noë, 2003, p. 161).

One way to elaborate this point in a slightly stronger way is by drawing an analogy between *flying* a plane and *being in* a flight simulator. Even when the pilot is not aware which condition she is in, it is only actually flying a plane that makes her fly a plane. Being in a flight simulator does not. The point here is that similarly conscious experience involves doing things and cannot be dissociated from such doings without changing the phenomenon. Even though one may preserve certain aspects of experience in a dissociated brain and body that maintains the local representatives of normal sensorimotor contingencies – as in the flight simulator – the result is not the same natural phenomenon that one initially set out to explain. Being an experiencing individual *includes* interaction with the world.

For those who see fundamental problems with inner localization there are two alternative scenarios. One of these is the option of *external localization*. External localization focuses on the realizers of conscious experience, claiming that these involve both internal and external processes. In this interpretation, consciousness is located partly outside the head in the sense that, next to brain processes, the reciprocal sensorimotor links make processes in the environment co-constitutive of conscious experience. External localization has been defended for cognitive processes by Clark and Chalmers (1998), Clark (2008), and Keijzer and Schouten (2007). Although many find external localization highly counterintuitive for consciousness, it is explored in a positive spirit by Rowlands (2003) and Hurley (1998; forthcoming).[6] An external localization scenario sets the contingencies of sensorimotor theory center stage. Within this scenario, a combination with workspace theory would be highly beneficial. The whole set of ongoing dynamical sensorimotor interaction loops, including workspace activity, could together constitute the experiential state. While sensorimotor contingencies would enable the explanations of differences in phenomenal quality, workspace theory would help to explain which interaction loops are constitutive of ongoing experience.

Both internal and external localization are subject to criticism. It is regularly argued that the criteria for applying the concept 'experience' are absent at subpersonal levels. In this view, it is a fundamental mistake to speak of physical processes as

[6] Rowlands (2003) defends the view that consciousness is partly externally located. He explicitly subscribes to a literal localization of consciousness in contrast with the possibility of having no localization of consciousness. For Hurley (forthcoming), who does not discuss the issue of no localization, drawing a boundary between the processes that are 'merely' causally involved in consciousness and the processes that constitute consciousness rests on explanatory issues.

being experiences (Putnam, 1999; Bennett & Hacker, 2003). Of course, there are localized processes that are necessary for experience, but it is denied that any of these processes themselves constitute consciousness – at the subpersonal level there is no place 'where consciousness happens' (Dennett 1991). Even if physical processes are interpreted as the vehicles of the content of experience,[7] this does not imply that consciousness can be *identified* with these processes. After all, the ascription of content to vehicles may derive from the way these vehicles are functionally integrated with the activity of the organism as a whole, so that experience is conceptually tied to the activity of the organism rather than to the necessary subpersonal preconditions. This conceptual background brings us to our final, *no localization* scenario.

Sensorimotor theory provides examples of such a non-localization interpretation, e.g. where it is said that experiencing is something we skilfully do rather than any of the underlying physical processes (O'Regan & Noë, 2001a, b; Myin & O'Regan, 2002). In this view, sensorimotor contingencies are relevant to the contrasts within experience since they characterize our perceptual engagement with the environment. While the conscious/not conscious contrast does not figure within the subject's experience, we suggest that workspace theory adds to the picture by providing a subpersonal theory of the neuronal dynamics that form a precondition for experience. Workspace theory can provide a framework to understand the mechanisms by which processes become integrated to enable thought, speech and further action, thus helping to explain some of the preconditions on which a sensorimotor theory depends.

These three scenarios provide different starting points and directions for possible combinations of workspace and sensorimotor theories. In all three cases, it is beneficial to combine both workspace and sensorimotor theories. Thus without making prior commitments to any of the three scenarios, we can hold that workspace and sensorimotor theories should be combined as it will lead to an increase in explanatory potential compared to both sets of theories separately.

5. Conclusion

We have argued that sensorimotor- and workspace theories of conscious experience can be fruitfully combined. First, they are no rivals since they have different domains of application. Sensorimotor theories can best be cast as addressing the specific quality of experiences. Workspace theories on the other hand are best seen as addressing the differences between those processes that are- and those that are not directly correlated to experience. Second, even though proponents of both sets of theories work from fundamentally different presuppositions concerning experience, it is possible to formulate different scenarios under which a combination can take place, reflecting the different background assumptions: Consciousness can be interpreted as internally localized, as partly externally localized, or as not being localized at all. In all three scenarios, advantages are to be expected from the combination of the theories.

It should be obvious that these three scenarios have different implications for the separate theories, some of which will not be acceptable for current defenders of workspace- or sensorimotor theories. In the inner localization scenario, a workspace forms the dominant basis of consciousness, while sensorimotor dependencies become relevant as factors modulating this neural activity. This relatively modest role is certainly not what sensorimotor theorists have in mind. In the case of external

[7] Note that such an interpretation does not imply internalism, as some have argued that vehicles of content need not be limited to processes within the head (Hurley, 1998; Rowlands, 2003; Noë, 2004; Thompson, 2007).

localization, some workspace theorists may be dissatisfied with the externalism which reduces the role of the brain to being merely part of the relevant interaction loops. Nevertheless, this scenario has the potential to be developed further in a way that integrates neural, bodily and environmental processes in a way that is common practice in embodied cognition. From a sensorimotor perspective it could be questioned to what extent this is about experience, rather than merely the subpersonal preconditions of experience. The scenario that rejects localization may appeal to those who are strongly committed to the sensorimotor theory as an account of experience itself. In this scenario, workspace theory will remain necessary to explain features that are descriptive of experience. In particular it could account for differences between consciousness-correlated and not consciousness-correlated processes within conscious organisms.

The three scenarios thus have different implications for the way and extent in which the sets of theories are to be integrated. In the first two cases, integration will be an important issue, requiring both sets of theories to become adapted to one another. In the no localization scenario, this need may be felt less strongly and both sets of theories can remain comparatively independent from one another. Combining workspace and sensorimotor theories is thus not a unitary affair but an enterprise that may enfold in very different ways, depending on the scenario chosen. Eventually it will become an issue which of the possible scenarios provides the most grip on the absolute and comparative gaps of consciousness.

Finally, an important issue remains how much of conscious experience will be covered if the combination succeeds, irrespective of the way the combination eventually takes shape. The sensorimotor theories we discussed focus mainly on perceptual experience and it remains to be seen to what extent for example the experiential aspects of emotions can be integrated in the approach. Workspace theories are typically developed based on *reports* of experience, leading to the question whether they are necessarily limited to the domain of reportable experiences. Other criteria besides reportability may be added here, such as availability for action and attention. Still, workspace theory remains strongly oriented on high-level human experience. It remains to be seen whether either workspace- or sensorimotor theory or their combination can be used to develop ideas on other forms of experience. While it remains an open question to what extent the range of application of the theories can be extended, we suggest that the procedure of *combining* these different kinds of theories will prove beneficial.

References

Austin, John L. (1962), *Sense and Sensibilia* (Oxford: Oxford University Press).

Auvray, Malika; Hanneton, Sylvain; O'Regan, J. Kevin (2007), 'Learning to perceive with a visuo-auditory substitution system: localisation and object recognition with 'The vOICe'', *Perception*, 36, pp. 416-30.

Auvray, Malika; Myin, Erik (2009), 'Perception with compensatory devices: from sensory substitution to sensorimotor extension', *Cognitive Science*, **33** (6), pp. 1036-58.

Baars, Bernard J. (1988), *A Cognitive Theory of Consciousness* (Cambridge: Cambridge University Press).

Baars, Bernard J. (2002), 'The conscious access hypotheses: origin and recent evidence', *Trends in Cognitive Sciences,* 6, pp. 47-52.

Bach-y-Rita, Paul (1984), 'The relationship between motor processes and cognition in tactile vision substitution', in Prinz, W., Sanders, A.F. (eds.) *Cognition and Motor Processes,* pp. 149-60 (Berlin: Springer-Verlag).

Bach-y-Rita, Paul (2002), 'Sensory substitution and qualia', in: Noë, A., Thompson, E. (ed.) *Vision and Mind: Selected Readings in the Philosophy of Perception*, pp. 419-36 (Cambridge: MIT Press).

Bennett, Maxwell R.; Hacker, Peter M.S. (2003), *Philosophical Foundations of Neuroscience* (Oxford: Blackwell Publishing).
Calvo, Paco; Gomila, Toni (eds.) (2008), *Handbook of Cognitive Science: An Embodied Approach* (Elsevier).
Chalmers, David J. (1995), 'Facing up to the problem of consciousness', *Journal of Consciousness Studies,* 2, pp. 200-19.
Chalmers, David J. (1996), *The Conscious Mind* (Oxford: Oxford University Press).
Clark, Andy (1997), *Being There: Putting Brain, Body, and World Together Again* (Cambridge: MIT Press).
Clark, Andy (1999), 'An embodied cognitive science?', *Trends in Cognitive Sciences,* 3, pp. 345-51.
Clark, Andy; Chalmers, David J. (1998), 'The extended mind', *Analysis,* 58, pp. 7-19.
Clark, Andy (2008), *Supersizing the Mind* (Oxford: Oxford University Press).
Dehaene, Stanislas; Naccache, Lionel (2001), 'Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework', *Cognition,* 79, pp. 1-37.
Dennett, Daniel C. (1991), *Consciousness Explained* (Boston: Little, Brown and Company).
Dennett, Daniel C. (2001), 'Are we explaining consciousness yet?' *Cognition,* 79, pp. 221-237.
Dretske, Fred (1995), *Naturalizing the Mind* (Cambridge: MIT Press).
Edelman, Gerald M.; Tononi, Giulio (2000), *A Universe of Consciousness. How Matter Becomes Imagination* (New York: Basic Books).
Engel, Andreas K.; Fries, Pascal; Köning, Peter; Brecht, Michael; Singer, Wolf (1999), 'Temporal binding, binocular rivalry and consciousness', *Consciousness and Cognition,* 8, pp. 128-51.
Engel, Andreas K.; Singer, Wolf (2001), 'Temporal binding and the neural correlates of sensory awareness', *Trends in Cognitive Sciences,* 5, pp. 16-25.
Hurley, Susan L. (1998), *Consciousness in Action* (Cambridge: Harvard University Press).
Hurley, Susan L. (forthcoming), 'Varieties of externalism', in Menary, Richard (ed.) *The Extended Mind* (Ashgate Publishing).
Hurley, Susan L.; Noë, Alva (2003), 'Neural plasticity and consciousness', *Biology and Philosophy,* 18, pp. 131-68.
Keijzer, Fred (2001), *Representation and Behavior* (Cambridge: MIT Press).
Keijzer, Fred; Schouten, Maurice (2007), 'Embedded cognition and mental causation: setting empirical bounds on metaphysics', *Synthese,* 158, pp. 109-25.
Koch, Christof (2004), *The Quests for Consciousness: a Neurobiological Approach* (Englewood: Roberts and Company Publishers).
Levine, Joseph (1983), 'Materialism and qualia: the explanatory gap', *Pacific Philosophical Quarterly,* 64, pp. 354-61.
Levine, Joseph (1993), 'On leaving out what it's like', in: Davies, M. Humphreys, G.W. (eds.) *Consciousness* (Oxford: Basil Blackwell).
Merikle, Philip M.; Daneman, Meredyth (1999), 'Conscious vs. unconscious perception', in: Gazzaniga, M.S. (ed.). *The New Cognitive Neurosciences, Second Edition* pp. 1295-303 (Cambridge: MIT Press).
Metzinger, Thomas (2003), *Being No One: the Self-Model Theory of Subjectivity* (Cambridge: MIT Press).
Mossio, Matteo; Taraborelli, Dario (2008), 'Action-dependent perceptual invariants: from ecological to sensorimotor approaches', *Consciousness and Cognition,* 17, pp. 1324-40.
Myin, Erik; O'Regan, J. Kevin (2002), 'Perceptual consciousness, access to modality and skill theories: a way to naturalise phenomenology?', *Journal of Consciousness Studies,* 9, pp. 27-45.
Noë, Alva (2004), *Action in Perception* (Cambridge: MIT Press).
Noë, Alva; Thompson, Evan (2004), 'Are there neural correlates of consciousness?', *Journal of Consciousness Studies,* 11, pp. 3-28.
O'Regan, J. Kevin; Myin, Erik; Noë, Alva (2005), 'Sensory consciousness explained (better) in terms of 'corporality' and 'alerting capacity'', *Phenomenology and the Cognitive Sciences,* 44, pp. 369-87.
O'Regan, J. Kevin; Noë, Alva (2001a), 'A sensorimotor account of vision and visual consciousness', *Behavioral and Brain Sciences,* 24, pp. 939-1031.
O'Regan, J. Kevin; Noë, Alva (2001b), 'What it is like to see: A sensorimotor theory of perceptual experience', *Synthese,* 129, pp. 79-103.
Philipona, David L.; O'Regan, J. Kevin (2006), 'Color naming, unique hues, and hue cancellation predicted from singularities in reflection properties', *Visual Neuroscience,* 23, pp. 331-9.
Prinz, Jesse (2008), 'Is consciousness embodied?', in: Robbins, Philip; Aydede, Murat (eds.) *Cambridge Handbook of Situated Cognition* (Cambridge: Cambridge University Press).

Putnam, Hilary (1999), *The Threefold Cord: Mind, Body, and World* (New York: Columbia University Press).
Revonsuo, Antti (2006), *Inner Presence. Consciousness as a Biological Phenomenon* (Cambridge: MIT Press).
Robbins, Philip; Aydede, Murat (eds.) (2008), *Cambridge Handbook of Situated Cognition* (Cambridge: Cambridge University Press).
Rowlands, Mark (2003), *Externalism: Putting Mind and World Back Together Again* (Chesham: Acumen Publishing).
Shanahan, Murray (2006), 'A cognitive architecture that combines internal simulation with a global workspace', *Consciousness and Cognition,* 15, pp. 433-49.
Tononi, Giulio; Edelman, Gerald M. (1998), 'Consciousness and complexity', *Science,* 282, pp. 1846-51.
Tononi, Giulio (2004), 'An information integration theory of consciousness', *BMC Neuroscience*, **5**: 42.
Thompson, Evan (2007), *Mind in Life: Biology, Phenomenology, and the Sciences of Mind* (Cambridge; London: Harvard University Press).
Thompson, Evan; Varela, Francisco J. (2001), 'Radical embodiment: neural dynamics and consciousness', *Trends in Cognitive Sciences,* 5, pp. 418-25.
Velmans, Max (1990), 'Consciousness, brain and the physical world', *Philosophical Psychology,* 3, pp. 77-99.
Velmans, Max (2007), 'Where experiences are: Dualist, Physicalist, enactive and reflexive accounts of phenomenal consciousness', *Phenomenology and the Cognitive Sciences,* 6, pp. 547-63.
Varela, Francisco J.; Lachaux, Jean-Philippe; Rodriguez, Eugenio; Martinerie, Jacques (2001), 'The brainweb: phase synchronization and large-scale integration', *Nature Reviews Neuroscience,* 2, pp. 229-39.