

David Silverman

Sensorimotor Theory and the Problems of Consciousness

Abstract: *The sensorimotor theory is an influential account of perception and phenomenal qualities that builds, in an empirically supported way, on the basic claim that conscious experience is best construed as an attribute of the whole embodied agent's skill-driven interactions with the environment. This paper, in addition to situating the theory as a response to certain well-known problems of consciousness, develops a sensorimotor account of why we are perceptually conscious rather than not.*

1.

The prevailing scientific understanding of the mind takes it that conscious phenomena such as perception and thought, though incompletely understood, are things that happen in the brain. The problem is that no intelligible connection has been made between mere brain activity and the conscious experience, seemingly quite different in character, it is believed to constitute. This difficulty is recognized by Block, who says:

All scientifically oriented accounts should agree that consciousness is in some sense based in the brain; once this fact is accepted, the problem arises of why the brain basis of this experience is the basis of this one rather than another one or none, and it becomes obvious that nothing now known gives a hint of an explanation. (Block, 2009, p. 1113)

While Block thinks the difficulty a brute fact given current understanding, others have reasoned that, since consciousness cannot be intelligibly identified with neural activity, it cannot *be* neural activity. This line of reasoning is familiar from Chalmers (1996), who rejects

Correspondence:
Email: davsil@gmail.com

physicalism altogether, although his alternative, dualism, does nothing to make consciousness more naturalistically explicable.

The sensorimotor approach makes use of the same line of reasoning, but reaches a quite different conclusion to Chalmers. Endorsing Ryle's (1949) view that the conception of the conscious mind as a 'ghost in the machine' generated by the brain is a category mistake, the approach claims that conscious experience is not something that happens in the brain, but something we do. Building on this principle, the approach's proponents have developed a contemporary view which promises to account for the phenomenal qualities of perceptual consciousness by appeal to various ways in which embodied agents skilfully interact with the environment.

One early example of an explanatory reward concerns the visual feeling of being embedded within a large and richly detailed scene. Experiments on 'change blindness' have found that subjects often fail to report seeing dramatic changes taking place in full view (e.g. O'Regan, Rensink and Clark, 1999), indicating that visual experience is likely not to involve an expansive and uniformly detailed inner model. The challenge, given this result, is to explain why visual experience nonetheless feels rich and expansive. Adopting a sensorimotor perspective, O'Regan (1992) and Noë (2002; 2004), among others, meet this challenge by claiming that visual experience is not constituted by the activation of an inner model but by skill-driven, temporally-extended exploration incorporating bodily movements such as saccades.

Pursuing the idea that we can make the phenomenal features of perceptual consciousness intelligible by identifying them with properties of bodily interaction, O'Regan and Noë (2001) make the broader suggestion that we can use sensorimotor patterns to explain in a near comprehensive way what distinguishes the varying phenomenal qualities of perceptual consciousness from one another. To this end, the authors propose that the specific qualities of experiencing, for example various colours and sounds, and more broadly the qualities of vision, audition, and other modalities, should be identified with sets of 'sensorimotor contingencies', laws of dependence between movement by perceiver or object and changes in sensory stimulation (see also Noë, 2004; O'Regan, 2011). The appeal of this proposal lies partly in its very great explanatory potential, promising as it does a systematic framework with which to make the qualities of consciousness scientifically tractable, and in turn a strong foundation for challenging the idea that they are fundamentally mysterious 'qualia'.

There are also compelling reasons to suspect that the sensorimotor contingency proposal is correct. One is an informal analogy. The softness one experiences when squeezing a sponge resides, quite unmysteriously, in its yielding easily to the touch, a sensorimotor law of just the sort that the sensorimotor framework appeals to (Myin, 2003). By analogy, it is possible that other kinds of experiential quality also reside in such laws, with the relative inscrutability of some qualities, such as the looks of colours, resulting from the fact that the corresponding laws are much more complex (O'Regan, 2011). And, in the case of colour qualities, the approach has received a measure of empirical support, O'Regan and colleagues having found that the most universally adopted colour categories (Berlin and Kay, 1969) correspond to the surfaces that yield the most easily described patterns of sensorimotor contingency (Philipona and O'Regan, 2006; Witzel, Cinotti and O'Regan, 2015), and therefore that we can explain what distinguishes these colour qualities from others if we identify them with sensorimotor laws.

These examples illustrate the progress that the sensorimotor approach is making toward accounting for the phenomenal character of perceptual consciousness. In light of the overarching limitations of identifying consciousness and its properties with brain activity, the explanatory returns the approach has started to bring in this domain give support to the still broader thesis that the conscious mind is in general best understood as an attribute of the whole embodied subject's ways of skilfully interacting with the environment. A further task, which recent work by O'Regan has begun, is to apply the sensorimotor perspective to the question of why we are perceptually conscious at all (O'Regan, 2011; 2016; Degenaar and O'Regan, 2015a,b). The bet is that with careful attention it will be possible to identify principled differences between conscious and unconscious ways of interacting, and use these to account for the existence of consciousness.

One task this paper will take up is to explicitly situate the sensorimotor approach as a response to certain well-known problems of consciousness as framed by analytic philosophy of mind. §2 identifies four different problems of consciousness, distinguishing the 'hard' from the 'easy', and the 'comparative' (why an experience has one phenomenal quality rather than another) from the 'absolute' (why it is conscious at all). §3 explains why conscious experiences are in general likely to be more usefully identified with skill-driven bodily interactions than brain states or internal representations, and §4

explains in a little more detail how the appeal to sensorimotor contingencies helps solve the comparative problems. With the approach thus situated, the paper will offer a sensorimotor approach to the question of why conscious perceivers are conscious rather than not. §5 proposes a principled difference between conscious and unconscious ways of perceptually interacting with the environment. §6 attempts to overcome potential objections that the kinds of interaction I propose constitute consciousness are in fact either unnecessary or insufficient for it.

2.

Let's remind ourselves of the problems that need to be solved in order to exhaustively naturalize consciousness. First, there are the 'comparative' and 'absolute' problems (Hurley and Noë, 2003). The comparative problem is that of explaining why a conscious experience, for example seeing the colour red, has one phenomenal quality rather than another. The absolute problem is that of explaining why there is any phenomenal quality rather than none.

There are also the 'hard' and 'easy' problems (from Chalmers, 1996, though my presentation of them will differ slightly from his). While not to be confused with the absolute/comparative distinction, the hard/easy distinction is in its own way also a distinction between the absolute and the comparative, since the easy problem requires us to explain, in comparative fashion, why one physical thing rather than another is conscious, and the hard problem requires us to explain, in absolute fashion, how physical things can be conscious at all.

Solving the hard problem entails dissolving the 'zombie' intuition, which holds that any physical feature we take to be the substrate of consciousness could in principle exist in the absence of consciousness. Solving the easy problem entails, supposing that consciousness is indeed physical, explaining which physical features constitute (e.g.) a conscious mind, or a conscious experience of the colour red. Though it does not require us to avoid the zombie intuition, the easy problem poses a considerable challenge, not least because it requires us to explain why one physical feature is conscious while another is not in a *principled* way. Mere descriptions of the physical (e.g. neural) correlates of consciousness will not suffice to solve even the easy problem unless they are accompanied by a suitably principled explanation for why these features are the correlates and not some others.

Though literature on consciousness often distinguishes the hard from the easy problem, on the one hand, and the absolute from the comparative problem, on the other, the two distinctions are not usually laid out explicitly all at once. In fact, it is useful to consider that the distinctions intersect, meaning that to explain consciousness fully we must solve four problems: (i) hard-absolute — how consciousness can in general be physical, (ii) hard-comparative — how e.g. the look of red can be physical, (iii) easy-absolute — why consciousness in general has one physical substrate rather than another, and (iv) easy-comparative — why e.g. the look of red has one physical substrate rather than another.

The sensorimotor approach, as I shall understand it, addresses the easy and hard problems in a bottom-up manner, attempting initially to solve the easy problems, such that consciousness and its qualities are intelligible and tractable from a scientific perspective. Its solution to the hard problems will involve further argumentative steps intended to show that this solution is sufficient to account exhaustively for consciousness.

3.

Before we proceed, it will be necessary to explain in some further detail why appealing to the possession and exercise of sensorimotor skills offers a more promising way to explain consciousness than brain activity. O'Regan and Noë (2001) emphasize the explanatory limitations, touched on in the last section, of neural correlates of consciousness (NCCs). They argue that whatever kinds of brain activity we may find are correlated with consciousness — for example, coherent oscillations at particular frequencies (e.g. Crick and Koch, 1990) or the quantum properties of neurons (e.g. Penrose, 1994) — there will always remain a question of why that neural feature gives rise to consciousness (with a particular phenomenal quality, or at all). This means that descriptions of NCCs are not even enough to solve the easy problems of consciousness, as I have characterized them, let alone the hard problems.

The sensorimotor approach, it should be acknowledged, is not the only available attempt to make the link between consciousness and its physical substrate more intelligible than NCCs can. There is a large body of relatively mainstream thinking which, rather than appealing directly to brain activity, appeals to representations that are *realized* by the brain and said to yield conscious experience by virtue of their

contents and in some accounts functional roles (e.g. Tye, 2000). Representationalism, as this approach is known, promises to make the existence and particular phenomenal character of conscious experiences naturalistically explicable while staying faithful to the orthodox cognitive scientific construal of the mind/brain as an information processing device, and, at least in a fashion, to the conception of conscious experience as a kind of neural activity.

As I shall understand it, the sensorimotor approach — in line with other forms of ‘enactivism’ (e.g. Thompson, 2007) — denies that we should identify conscious perception and other kinds of experience with internal representation, on pain of failing to adequately naturalize consciousness. The reason is this. Hutto and Myin (2013) have shown that the leading attempts to naturalize neurally-borne content (i.e. truth conditions) are unconvincing, because they depend on the idea that covariance between the brain and outside world is sufficient, under certain conditions, for content, whereas nothing has been done to establish that this is so. Millikan (1989), for example, attempts to naturalize content by appealing to the proper function brains have of possessing and making appropriate use of covariance relations. The problem is that even the appeal to proper function merely shows that brains have the function of exploiting patterns of covariance with the environment and does not establish that they bear content about it.

I will not argue, as Hutto and Myin (2013) do, that content should never play any role in explaining perception and action. Consider that many prominent endorsers of orthodox (i.e. content-involving) cognitive science suggest that content does not need to be naturalized but exists by stipulation, their having defined content as a kind of covariance (e.g. Rupert, 2011). This is often accompanied by the idea that content exists as an explanatory fiction, e.g. a metaphor, of one kind or another (McDowell, 1994; Sprevak, 2013; Egan, 2014). Leaving to one side Hutto and Myin’s (2013) further claim that content is not *needed* to explain basic perception and action, the availability of these strategies suggests that content is, in principle at least, admissible as a theoretical tool for modelling the physical processes that produce bodily behaviour.

The lesson I wish to draw is that content, so far as has been established, does not exist in any more metaphysically-demanding sense than as stipulation or explanatory fiction. This view is not especially left-field given the number of endorsers of orthodox cognitive science that present it as such a fiction. Consciousness, by contrast, exists independently of any stipulations or explanatory fictions offered.

Moreover, consciousness exists for the conscious subject, whereas content as a stipulation or explanatory fiction exists at most relative to a theoretical model. Since consciousness has properties that content does not, consciousness cannot be identical to content.

Bodily skills, in contrast with content, are readily observable abilities to carry out various kinds of bodily movement. There is therefore good reason to give serious consideration to the possibility that appealing in various ways to skills and skill-driven bodily interactions can help solve the four problems of consciousness.

4.

We will continue by examining the sensorimotor approach's response to the comparative problems (easy and hard). The approach claims that to perceive is to carry out bodily action that is perceptually-guided. Bodily action is perceptually-guided just in case it instantiates 'mastery' of sensorimotor contingencies (SMCs), the ways that sensory inputs are disposed to change as a result of movement by the perceiver or objects. Mastery can be glossed either as implicit knowledge of SMCs or merely as adaptedness to them. Action, as I intend, is purposeful doing, and should be read in the present paper as including cases where you are merely disposed to move your body, even if you do not actually do so.¹

Perception need not be conscious (a point we will return to shortly), but where it does occur consciously its particular phenomenal qualities are, O'Regan and Noë (2001) claim, constituted by the SMCs which are the objects of the mastery. The phenomenal quality of a sense modality (such as vision) is identical to the SMCs determined by the physical properties of the sense organs (in this case eyes), and the phenomenal qualities featured within modalities, e.g. the quality of seeing the colour red, are identical to SMCs determined jointly by the sense organs and external objects (in this case, the eyes and the surface properties distinctive of red objects).²

¹ I offer a more detailed account of sensorimotor mastery in Silverman (2017).

² Autopoietic enactivists such as Thompson (2007) claim that biological interactions instantiating autopoiesis and/or autonomy are pre-requisites for consciousness or genuine perception, a claim denied by Degenaar and O'Regan (2015b) who claim that distinctively perceptual interactions suffice for perception. This paper will not take a stance on this issue. Note, however, that even if biological interactions are pre-requisites for perception and consciousness, as Noë (2004) argues, we will still need to explain what distinguishes one phenomenal quality from another, and conscious from

Using the looks of colours as a case study, let's see how this proposal helps solve the easy-comparative problem. Phenomenal similarities and differences between colours can be quantified on the basis of verbally reported judgments of phenomenal difference and similarity, and the degrees of reported similarity and difference plotted diagrammatically as a 'colour space' (Palmer, 1999). Differences between the varying sets of SMCs that obtain when we experience various colour qualities can also in principle be quantified using mathematical laws, and plotted in what we could call SMC space, which we hypothesize will be isomorphic to quality space. To explain in a principled way why the phenomenal quality associated with red supervenes on one set of SMCs rather than another, we can appeal to the fact that the isomorphism would no longer obtain if the quality occupied a different point in SMC space.

One potential objection is that colour phenomenology could conceivably be inverted so that the isomorphism between colour space and SMC space is preserved but colour phenomenology changes. If this were a possibility, we would have no principled explanation of why the phenomenal quality of a colour such as red supervenes on one set of SMCs rather than the inverse set. Palmer, however, has argued convincingly that colour space is not symmetrical, meaning it is not possible for colour phenomenology to be inverted while degrees of reported phenomenal difference and similarity are preserved (*ibid.*).

The problem nonetheless remains, Palmer suggests, that it is possible to conceive of all colour qualities being uniformly different, for example all seeming darker than they do, without affecting the differences and similarities charted in colour space. This means the isomorphism between colour space and SMC space cannot explain why all colours do not seem, for instance, uniformly darker. This threatens to compromise the sensorimotor account's ability to comprehensively overcome the easy-comparative problem.

Moreover, even if the account just given were sufficient to solve the easy-comparative problem, it appears insufficient to solve the hard-comparative problem. The intuition motivating the latter problem, illustrated by Jackson's (1982) example of a colour scientist who has never seen colour, is that even knowing everything that could be known about the physical features (e.g. SMCs) upon which

unconscious perception — work that may be done by appeal to the non-biological kinds of interaction I describe in this paper.

phenomenal qualities supervene would not suffice to know what phenomenal qualities are like — suggesting, Jackson once claimed, that phenomenal qualities are non-physical.

The hard-comparative problem and the remaining obstacle to solving the easy-comparative problem can be addressed at once if we take it that the experience of a colour is, as Pettit (2003) proposes, simply the experience of being able to discriminate a colour appropriately from other colours. Noë (2004, p. 139), endorsing this proposal, observes that it does not require us to endorse a crude and implausible kind of behaviourism given the fineness of the discriminations required. It is also worth emphasizing that the ability to discriminate one colour from another needs to persist through changes to the way that the perceived surfaces are illuminated from the perceiver's point of view, and this illumination changes as the perceiver moves her body. This means that the abilities in question are not abilities to discriminate just any property, as a crudely behaviourist account might suggest, but are necessarily abilities to discriminate obtaining sensorimotor laws.

5.

O'Regan (2011) acknowledges that sensorimotor theory's explanation for the phenomenal character of perception, though sufficient to address the comparative problem, is insufficient to solve the absolute problem. This is because you could exercise mastery of SMCs (and hence perceive) without being conscious. O'Regan uses the old example of an absent-minded driver who, while attending to conversation with her passenger, is unconscious of the road traffic she successfully negotiates. On the basis that this is a plausible account of what happens during absent-minded driving, O'Regan treats the problem of explaining consciousness as the problem of explaining what distinguishes the absent-minded driver from one who is aware of the road. His answer is that perceptual consciousness requires, in addition to the ability to perceive and act, an appropriate kind of cognitive access, in particular the ability to think about and communicate what one sees.

Attempts to identify phenomenal consciousness with cognition are in general subject to certain well-known objections, for instance the possibility of zombies that engage in the relevant kind of cognition without being conscious (Chalmers, 1996) and the possibility that

phenomenal consciousness sometimes occurs in the absence of cognitive access (Block, 1995). I will come back to these in §6.

My focus in the present section and its associated subsections (§5.1–5.4) will be on developing O'Regan's proposal in a way that identifies consciousness with the embodied subject's skill-driven interactions. In §5.1 I will give a more detailed summary of O'Regan's response to the absolute problem, which I will treat merely as a response to the easy-absolute problem. In §5.2 I will identify two competing conceptions of cognitive access, one which is incompatible with the claim that consciousness depends on interactive abilities, and one of which is compatible with this claim, and in §5.3 I will defend a version of the latter. In §5.4 I will make some suggestions about the kinds of skill-driven interaction that cognitive access consists of.³

5.1.

Elaborating on his claim that perception takes place consciously where we have cognitive access to what we perceive, O'Regan (2011) states that this involves our being poised to make use of our mastery of SMCs for the purpose of cognitive activities such as thought, communication, and action planning. He adds that consciousness also requires a psychological 'self', given that one feature of conscious thoughts and sensations is that they always belong to us as unified subjects.

Although important disanalogies between sensorimotor and representationalist theories of consciousness will shortly emerge, it will help situate O'Regan's account if we draw some comparisons with representationalist theories. The account might be construed as analogous to Tye's (2000) PANIC theory, which claims that mental (e.g. perceptual) representations are conscious when they are, among other things, poised to enter into belief-formation. O'Regan actually draws a tentative parallel between his account and dispositionalist higher-order thought (HOT) theories (e.g. Carruthers, 1996), which claim that a perceptual state is conscious when one is poised to have the higher-order thought that one is in the perceptual state.

The key feature distinguishing HOT from a first-order theory like Tye's is that one must have cognitive access not only to the content of

³ This paper focuses on perceptual consciousness. By so doing, it aims to explain a form (arguably the paradigmatic form) of phenomenal consciousness.

the percept, but to the fact that one has this perceptual content. O'Regan's claim that consciousness requires a self could be interpreted as meaning that, in order to be conscious of a red car that you see, you must not only be poised to think about the car, but also be poised to have the thought 'I see a red car'. This would make sense of the supposed analogy between O'Regan's account and a dispositionalist HOT account.

However, there is no immediately apparent reason to think that the driver conscious of a red car in front is necessarily poised to have the thought that she herself sees a red car in front of her. On the face of it, merely being ready to think about the red car in front is sufficient to entail that she is not looking at the car in an absent-minded way and therefore that she sees it consciously. I will take it, therefore, that if the ability to ascribe one's own mental states to oneself as a subject is a pre-requisite for perceptual consciousness, this is only because it is a pre-requisite for perceiving (see Rowlands, 2016, who argues that selfhood is implied by sensorimotor mastery) or thinking (e.g. Kant, 1787/1998) rather than because it is a special feature of consciousness. The sensorimotor account as I shall understand it is more closely analogous to Tye's theory than the HOT account, requiring merely that you have an appropriate kind of cognitive access to what you see — though in the sensorimotor account, like HOT, this cognitive access is to be put to use in thought and communication, which may not be the same as belief-formation.

Since it identifies phenomenal consciousness with having cognitive access to otherwise unconscious mental activities such as perception, the sensorimotor account — in common with representationalist proposals — identifies phenomenal consciousness with a variety of what Block (1995) calls access consciousness. Whereas Block is famous for denying that access consciousness is necessary for phenomenal consciousness, even he allows that having or making use of cognitive access is constitutive of a kind of consciousness. If Block is mistaken in claiming that phenomenal consciousness can exist without access consciousness, we may freely assume that phenomenal consciousness just is a certain kind of access consciousness.

5.2.

Let's examine a crucial distinction, which O'Regan overlooks, between what is usually meant by 'cognitive access' and what proponents of the sensorimotor theory ought to mean by it. Consider, once

more, O'Regan's claim that perception takes place consciously just in case a subject has cognitive access to her sensorimotor mastery and is thereby poised to make use of it for the purpose of 'judgments, reasoning, planning and communication' (2011, p. 182). We can identify two competing ways of understanding what is meant by 'cognitive access' here, and which we choose will have a dramatic bearing on how we understand the thesis.

On one understanding (CA1), cognitive access is the accessibility of an internal representation to the brain's cognitive machinery. This may consist, for example, of the representation's being stored in working memory and consequently available for use by further processes. Given that it uses subpersonal vocabulary, the conception is committed to a particular understanding of the subpersonal machinery that underlies thought and action, namely one that involves the manipulation of neurally-realized representations. In cases where mental states such as memories and thoughts are rightly attributed to the person and not just the brain, it takes them to consist simply of being in corresponding neurally-realized subpersonal states.

On an alternative understanding (CA2), cognitive access is a personal-level (or agent-level) phenomenon constituted by the subject's readiness to have certain thoughts, plan and carry out certain actions, and so forth. It may include, for example, the readiness to have a thought about what one sees. The conception is neutral about what it is to think, plan, and communicate, and is not committed to any particular understanding of the subpersonal processes that underlie these activities.

CA1 would cement a revisionistic reading of the sensorimotor approach under which it is a version of representationalism (e.g. Seth, 2014; see also, Silverman, 2017). However, this would be contrary to what I have argued are the proper motivations of the sensorimotor approach, which preclude identifying consciousness with the deployment of content-bearing neural states. The sensorimotor theory should therefore endorse CA2. This will leave us free to understand the relevant capacities for thought and communication to consist of capacities for embodied interaction with the environment.

Before we continue, it is worth examining what philosophers usually mean by cognitive access: this will help underline the need — unrecognized by O'Regan (2011) — to explicitly distinguish the sensorimotor theory's understanding of cognitive access from the more typical understanding, and will also help ensure we do justice in our construal of cognitive access to the considerations that

legitimately motivate this more usual view. Let's begin with Block, who says:

I think that [verbal] report is a good first approximation to access consciousness. One important reason it's not adequate is that animals have access consciousness. And you know by and large, they can't report. So it's the extent to which a phenomenal state is available to cognitive machinery. And so a mouse or a dog or a cat can be access conscious of a smell or something they see. And then a representation is sent to their machinery of reasoning and deciding, even if they can't report it. (Block, 2015)

This shows that Block construes cognitive access as the availability of internal representations to cognitive machinery — by which he means, it is safe to say, wholly neural machinery. This demonstrates his commitment to CA1. It also appears to suggest that there are cases in which even access consciousness (let alone the consciousness he alleges occurs without cognitive access) cannot even in principle be identifiably manifested in behaviour. The purported advantage of this view is that it accounts for cognitive access without appealing to behaviours, in particular verbal report, which animals may lack.

However, this apparent advantage comes at a cost. Degenaar and O'Regan (2015b) suggest that talk of mental states which cannot be recognizably manifested in behaviour is incoherent on the verificationist grounds that we cannot coherently talk about entities that by nature cannot be identified empirically. Without needing to rely on verificationism, Cohen and Dennett (2011) argue compellingly that it is not scientifically respectable to posit varieties of consciousness that cannot sometimes be evidenced by behaviour.

Cohen and Dennett share, all the same, Block's aim of accounting for the consciousness of subjects who cannot express their mental states in their behaviour, this time offering the example of patients with locked-in syndrome who are conscious despite being notionally completely unable to move their bodies. With this in mind, the authors identify access consciousness, like Block, with the use of neurally-realized, functionally-defined faculties such as attention and working memory — specifying, for the sake of scientific respectability, that the functions must sometimes be evidenced by appropriate behaviours such as verbal reports in subjects who are not paralysed.

Cohen and Dennett's position is most plausibly read as suggesting that the link between the functionally-defined faculties such as working memory which constitute cognitive access and the potential bodily behaviours they are associated with is causal rather than

constitutive, with the behaviours serving only as empirical evidence that the functions are being realized. This suggests that they too endorse CA1.

The problem with CA1 is that it prevents us from adequately explaining what the faculties that constitute cognitive access, let's say attention and working memory, actually consist of. Since we are aiming to identify cognitive access with consciousness, we cannot identify the faculties with content-bearing neural states, given the prohibition (set out in §3) on identifying consciousness with such states. Even if, for the sake of argument, we ignored this prohibition and did identify the faculties with neurally-borne content, this would not circumvent the non-trivial task of identifying the behavioural criteria which Cohen and Dennett themselves argue are needed in order to ascribe access consciousness to a subject.

With these points in mind, I suggest that, endorsing CA2, we identify cognitive access with readiness to have certain thoughts or communicate certain information. To explain what this readiness involves, we must secure a suitably tight conceptual link between cognitive access and interactive capacities, while also accounting for the cognitive access enjoyed by subjects who do not have the ability to perform bodily actions.

5.3.

Keeping in mind the considerations just raised, I suggest identifying readiness to think and communicate with dispositions to exercise capacities for appropriate kinds of embodied interaction. A capacity, unlike a mere disposition, is something a subject can choose whether or not to exercise. So I am not identifying cognition, behaviouristically, with mere conditioned responses. The appeal, on the other hand, to *dispositions* serves two purposes. First, cognitive access is a dispositional notion. It does not require that we think about or communicate what we see, merely that we are poised to think about or communicate what we see. Second, the notional completely paralysed locked-in patient, strictly speaking, lacks the capacity to carry out any bodily actions. Nonetheless, if no longer being paralysed would cause her to have a capacity, she is disposed to possess this capacity, and when she so chooses will be disposed to exercise it.

On the view I am suggesting, there is thus a logical dependence between consciousness and potential bodily interaction, yet we can accommodate cases such as the locked-in patient where no deliberate

bodily movement takes place. A locked-in patient, one can reasonably hypothesize, is conscious because she is suitably disposed to perform relevant actions — for example, verbally reporting what she sees — even though she is prevented from actually performing them.

It might be objected that dead (and therefore unconscious) people are also disposed to exercise capacities like these because they would exercise them if they were, among other things, alive. To avoid this *reductio*, we must be careful to distinguish between having a disposition which cannot be activated (e.g. because of paralysis) and lacking the relevant disposition entirely (e.g. because you are dead). There is nothing to stop us making a principled distinction between these two cases by appealing, for example, to the presence or absence of appropriate internal (e.g. neural) activity.

Recall that we can neither identify consciousness with mere brain activity, on pain of failing to explain it, nor with internally-borne content, given that content, unlike consciousness, can only exist so far as has been established relative to a theoretical model and not for the subject. Since we are proposing to identify consciousness with cognitive access, we cannot identify cognitive access with contents or brain activity.

This does not prevent us from appealing to neural properties, including content if construed as an explanatory fiction, to license the claim that someone possesses a relevant disposition. The locked-in patient can intelligibly be understood to have relevant dispositions by virtue of having, for example, neural or bodily states that are such as to ground the relevant dispositions, and therefore to enable the behaviours to occur in the event that nothing prevents them from occurring. Similarly, we can appeal to facts about the brain, including neurally-borne content understood as an explanatory fiction, to explain how the relevant dispositions are enabled subpersonally.

An important consequence of the explanatory limitations of appeals to content and internal states is that appealing to them will only be useful in so far as it helps explain how the relevant kinds of bodily interaction are produced, and we cannot consider which kinds of internal state (including representation) help enable the relevant bodily interactions until we have an idea about what those interactions are. We should, moreover, be open to the possibility that non-representational and/or environmental features play relevant enabling roles in addition to or in place of internal representation — making note of the fact that sensorimotor theorists (e.g. Hurley, 1998) often talk about

interaction rather than mere behaviour to underline the ongoing role played by the outside environment in shaping the relevant activity.

5.4.

Though conscious and unconscious kinds of perception draw respectively on different parts of the brain, and perhaps on different modules and representational vehicles, we are not appealing to facts about cognitive architecture to explain what distinguishes them. Instead, we are taking them to be differing ways in which the whole embodied subject interacts (or is disposed to interact) with the environment. This means that conscious perceptual interaction should be viewed as a transformed version of unconscious perceptual interaction rather than a distinct process running in parallel. It is an activity in which the capacity for unconscious perceptually-guided action is integrated with further capacities for thought or intelligence.

Let's consider what this transformed capacity might actually be. We were acquainted earlier with the most unambiguous example of a perceptually conscious subject, namely one who is ready to verbally report what she sees. Given that we are proposing to identify the disposition to verbally report what you see with consciousness rather than merely regarding verbal report as evidence that consciousness is occurring, it would be inappropriate to consider verbal report just a matter of making appropriate noises. Consider that some chess-playing computers, for example, have the ability to announce the moves they are making, but this does not entail that they understand or are conscious of what they are doing. This is a basic example of the problem identified by Searle (1980), whose thought experiment featuring a system that manipulates symbols syntactically without understanding their meaning illustrates the difficulty of distinguishing genuine linguistic understanding from a mere simulation merely on the basis of an individual's internal states or solitary behaviour.

What we need is an account of the interactive capacities that constitute genuine linguistic understanding. Such an account is offered by Maturana and Varela (1987), who identify meaningful language use with a temporally-extended and socially-embedded activity they call 'languaging'. Communication, as those authors understand it, is the non-meaningful coordination of behaviour between organisms (we might say 'agents' since nothing in their account requires that only living beings communicate). Communication occurs 'in the linguistic domain', and is therefore proto-linguistic, where the vocalizations or

gestures, etc. used to coordinate behaviour are both learnt and arbitrarily selected, i.e. by social convention as with Peirce's 'symbols' (1857–1866/1982). Meaningful communication, i.e. language, occurs, the authors finally propose, where communication in the linguistic domain takes place recursively, meaning that the way a given utterance is used by a community evolves over time, being further shaped with each use.

Meaning itself is on this account an essentially social phenomenon. Thought, as Villalobos (2014) explains in an account based on Maturana's, is dependent on languaging, being the covert deployment of a transformed and truncated version of natural language. Although thinking does not require bodily interaction in the short term, the account implies that thought can be explained by reference to the ways the agent (including her internal states such as brain states) drives and is driven in the longer term by linguistic communication, and therefore can itself be understood as an essentially embodied and socially embedded interactive capacity. This conception of thought avoids the failure of intelligibility that arises if we attempt to identify consciousness with brain states alone.

By combining this account of meaningful language and thought with the account of perceptual consciousness presently under consideration, we can hypothesize that perceptual consciousness occurs where the subject exercises sensorimotor mastery in perceptually-guided interaction, and is ready either to coordinate this action with the action of another by linguistically communicating the SMCs mastered in the way described by Maturana and Varela (verbal report) or to perform an internally-realized extension of the same interactive behaviour (thought).

Because the linguistic interactions described by Maturana and Varela are qualitatively distinct from the perceptual interactions described by the sensorimotor account, the proposal offers a qualitative distinction between conscious and unconscious capacities, at the same time tracking a real difference between cases of conscious and absent-minded driving. In so doing, the proposal does justice to the suggestion made by higher-order thought theories that perception occurs consciously where one is ready to think about one's perceptual states (reflecting in particular Carruthers', 1996, approach, in which the thoughts in question depend on language).

The advantage of this account is that it offers a robust criterion for genuine linguistic understanding, and so thoughtfulness, without appealing to any spooky neural or biological properties. While we

should not attribute understanding to just any machine that can issue sounds that resemble words we recognize, we must have behavioural criteria for linguistic understanding in order to ascribe it to humans, animals, and, one might imagine, extra-terrestrial beings — considering how interesting it might be in this latter case to know whether they understood what they communicated. Learning to communicate using arbitrarily chosen utterances, and doing so recursively among a community of similar language users, is a plausible candidate for such a criterion. If it is the right one, then it is hard to see what value there would be to looking inside the brain for the source of meaning.

A limitation with this proposal is that it fails to explain how non-linguistic creatures could be conscious. Adopting a pluralistic understanding of perceptual consciousness, we need not rule out the possibility that non-linguistic creatures sometimes engage in a different variety of conscious perception by virtue of integrating their capacities for perceptual interaction with non-linguistic cognitive capacities. In particular, we may appeal to the phenomenon of cognitive control: the endogenous control of attention, carried out for purposes such as long-range action-planning. Cognitive control is usually understood to be responsible for the feeling that an action is volitional. In so far as such an action is perceptually-guided, we can hypothesize that cognitive control also makes the associated percept conscious.

Empirical work on the pre-motor theory of attention (see, for example, Rizzolatti and Craighero, 2010) has suggested that the neural mechanisms that underlie spatial attention are the same as those that underlie motor planning rather than distinct mechanisms as traditionally thought, with covert rather than overt attention corresponding to weak activations in those areas. This is consistent with my claim that cognitive access consists of the exercise, or disposition to exercise, capacities for bodily movement. Recent work by Clark, Schumann and Mostofsky (2015) has suggested that cognitive control, in particular, is plausibly a bodily skill, noting that interventions such as mindful movement training targeted at improving motor control also bring improvements in cognitive control.

Taking it, therefore, that cognitive control is plausibly a bodily capacity, we must explain what kinds of bodily interaction it is a capacity for. Tentatively, I suggest the following. Perceptually-guided interaction is attentive when bodily movements are predominantly aligned toward a particular object of perception. It is endogenously controlled when we cannot easily predict changes in the subject's motor behaviour using facts about her environment, but must make

use of facts about her internal states and/or the ways these have been shaped by historical sensorimotor interactions. Finally, whereas adaptedness to SMCs can be ascribed on the basis of regularities present in the subject's sensorimotor coupling with the environment occurring over varying timescales, cognitively controlled action can be distinguished, accounting for the fact that it is better planned, by regularities present only over longer timescales.

If we adopt this view, we will find that cognitive control, and so perceptual consciousness, comes in degrees. It may on this basis be appropriate to attribute the absent-minded driver, and non-linguistic creatures, degrees of consciousness in line with the extent to which their interactions (or dispositions to interact) exhibit the above characteristics.

6.

Now that we have a line of response to the easy-absolute problem, let's consider some further challenges, beginning with the hard-absolute problem. To solve this problem, we must cast doubt on Chalmers' (1996) claim that there is a conceivable and therefore metaphysically possible world in which a 'zombie' exercises the cognitive capacities that in this world give rise to consciousness, but for whom there is nothing it is like to do so.

Myin (2016), endorsing a sensorimotor view of consciousness, claims that the phenomenal and the physical are two different perspectives one can take on the same bodily doing. To explain why consciousness seems distinct from bodily doing, he appeals to Merleau-Ponty's (1945/2013) distinction between the objective and the 'lived'. If you grasp one of your own hands with the other, Merleau-Ponty observed, you can either experience the grasped hand as if it were an external object or you can experience it as part of your own body (the lived), but you can never experience it in both ways simultaneously. If consciousness is itself bodily doing, it similarly can either be experienced as something you are doing, or as a bodily activity that is taking place, but never in both ways at once.

This phenomenological distinction supports the idea behind the phenomenal concepts strategy (e.g. Loar, 1999), which claims that consciousness can appear under distinct phenomenal and physical descriptions. By appealing to this difference in how consciousness is described, we can explain why zombies are conceivable without

needing to concede that they are possible, rendering their possibility an unnecessary posit.

I will venture what I take to be a logical point, namely that what it is like to be some x is just what it is to think, perceive, or otherwise cognize in the way that x does, described from that x 's point of view. By 'point of view', I do not mean an experience or other kind of mental state, but a way of describing a mental state, for example a phenomenal concept, or an experience of something as 'lived'. To attempt to describe the activities of a rock from the rock's point of view would yield an empty description, but this is not because it lacks something called a 'point of view', but because it is not intentionally directed toward the world. I suggest that we cannot resist allowing that purported zombies have minds, since by any theory of cognition they have the same cognitive capacities we do. If a purported zombie perceives and thinks, and we describe this from its point of view, we will also necessarily arrive at a non-empty description. So zombies are metaphysically impossible.

Those who claim that zombies are possible will object that this begs the question, since they intuit that the idea of describing a creature's point of view is only coherent if that creature possesses a special property known as 'consciousness', and it is the metaphysical possibility of purely physical creatures being conscious that the zombie intuition calls into doubt. What I am trying to make compelling, however, is a competing intuition in which the order of explanation is reversed. The possibility of a physical thing's being conscious is logically derived from the possibility of describing its minded activities from its own point of view, and this descriptive stance is available in all possible worlds.

6.1.

This response to the hard-absolute problem has a bearing on how I will now respond to two related worries about the response I have given to the easy-absolute problem, firstly that it is too vague, and secondly that consciousness might occur, as Block (1995) claims, in the absence of cognitive access.

The upshot of section §5.4 was that perceptual interactions, when transformed by linguistic ones, are certainly sufficient for phenomenal consciousness. But I claimed that certain non-linguistic cognitive capacities may also be sufficient for a kind of phenomenal consciousness, the case for this being uncertain. The worry that this is

unacceptably vague loses its edge if we consider the task of explaining consciousness to be, in part, what O'Regan (2016) calls a 'definitional enterprise'. He says:

The task of the sensorimotor approach is to find out what is particular about those ways of acting or interacting that people designate as 'conscious' ways. The task is a task of being precise about definitions: what capacities to act do we want to call cases of acting consciously?... [W]e shall see that once we are precise about definitions, the task of 'explaining' consciousness melts away. Thus a better formulation of the question 'What makes an experience conscious?' would be: 'What do we mean when we say that a person is experiencing something consciously?' and 'What do we designate as conscious ways of behaving?' (*ibid.*, pp. 44–5)

This understanding does not make the need to identify consciousness with certain capacities melt away, but it does make the problem considerably less pressing, rendering the task a question of deciding which capacities it is theoretically the most productive to identify with consciousness, rather than one of uncovering the hidden essence of certain physical or cognitive states.

The suggestion that explaining consciousness is a matter of deciding how to define 'consciousness' will to many be unappealing. This is because consciousness is often conceived of as a 'what it is like-ness' that is hard or impossible to define, but with which we are directly and intimately acquainted. If this is our conception, we will only assent to identifying consciousness with behaviour if this identity turns out to be a previously hidden fact about the essential nature of consciousness which we have now discovered; we will not accept this identity as a matter of stipulation. In line with the standard 'what it is like-ness' conception, it is therefore easy to read O'Regan as being an outright denialist about consciousness.

I believe that consciousness is perfectly real, yet O'Regan is correct that accounting for consciousness is in part a definitional enterprise. Consider again the claim that the answer to 'what is it like to be an x ?' will always be 'what it is it to cognize as an x does, described from an x 's point of view'. The problem is that an x 's cognition, when described from that x 's point of view, will not necessarily deserve to be labelled as 'conscious'. Deciding which capacities constitute consciousness is the challenge posed not by the hard-absolute problem but the easy-absolute problem, and it is in our response to this problem that we should take explaining consciousness to be a definitional enterprise.

A rock, having no cognitive capacities, is not conscious. A person who perceives while thinking about and reporting what she perceives is certainly conscious. Block claims that the ability to think about what you see may merely be an ability to detect consciousness rather than a part of what it is to possess phenomenal consciousness. However, the possibility that thought and language is merely a way to detect consciousness is only compelling if you deny that there is any difference between what it is like to perceive while thinking about what you perceive and what it is like just to perceive. On the proposal I am recommending, thinking is a transformation of the whole embodied perceiver's perceptually-guided interaction rather than a distinct process. There is no reason to doubt, in view of this, that thinking about what you perceive while you perceive it makes a difference to one's perceptual consciousness rather than merely serving as a way to measure it.

It is harder to decide whether cases of perceptually-guided action such as driving in the absence of capacities for thought and language, with or without cognitive control, are cases of consciousness, because they exist somewhere in the middle of a continuum. Just as deciding whether viruses are alive will depend on what we take to be the most useful way to define 'life', deciding whether these cases instantiate consciousness will depend on what we take to be the most useful way of defining 'consciousness'. This does not mean that the existence of consciousness is under dispute.

To resolve the question of whether such in-between cases should be described as conscious, we might attempt to form some intuition. The difficulty is that, whenever we think about our own conscious experiences, they are always the kind transformed by thought. Even if it were possible to simulate in one's imagination the mind of a perceiver who cannot think about what she sees, one could not think about what this was like without transforming the imagined act. So we cannot use introspection to think about unaccessed experience and make an intuitive judgment about whether to call this conscious. Block will claim that the absent-minded driver, though lacking conscious access to what she perceives, perceives in a *phenomenally* conscious way. But the obstacle that prevents us from forming an intuition about whether absent-minded driving deserves to be described as conscious is one that applies precisely to the question of whether the driver is phenomenally conscious.

In the absence of intuition, a better approach to determining whether subjects who do not integrate their perceptual activity with thought

and language are conscious is whether it is theoretically productive to think of them as conscious. It is agreed, in any case, that a subject who perceives and acts has cognitive capacities the exercise of which can be described from the subject's point of view. With this established, it is hard to see what value is added by insisting that the exercise of these capacities, in the absence of cognitive access, deserves to be described as *conscious*. A particular worry is that describing them as such leaves open the possibility of mistaking consciousness for a distinct property or thing that exists in addition to perception and action.

There is a stronger although not unassailable case for taking perception that is integrated with cognitive control but not thought and language to be conscious. The drawback with this position is that, as just explained, it is impossible to form an intuition based on introspection about whether such episodes deserve to be described as conscious. Its advantage, however, is that it does justice to the possibility of perception and action occurring unconsciously, while also doing justice to the continuities between the cognitive capacities of linguistic and non-linguistic creatures. For now I merely want to emphasize that deciding how much integration with higher cognitive capacities is required to make perception conscious is, if my suggestion is correct, a matter of deciding what will allow us to make the most theoretically useful distinctions rather than one of closing an explanatory gap.

6.2.

Consider an advantage of the present view. We have already seen that the sensorimotor approach has the advantage of avoiding the representationalist's (including higher-order theorist's) dependence on neurally-borne content. This is necessary in order to adequately respond to the easy problem. The present approach also has an advantage over representationalism in its response to the hard problem. One notable feature of the representationalist family of theories is the degree of hair-splitting and increasingly fine distinctions offered by one version of the approach to solve problems with another. Although they all agree that consciousness is identical to a kind of mental representation, they disagree on many technical points: does consciousness depend on a single representation (Tye, 2000), and does this representation need in part to represent itself (Kriegel, 2009)? Or does consciousness instead depend on having a higher-order representation of a lower-order representation, and if so, does

consciousness consist in the relation between the higher- and lower-order representations, or just in the higher-order representation (Brown, 2015)?

These distinctions are often motivated by empirical facts about what conscious phenomenology is actually like (*ibid.*), and the question of what sorts of mental states or activities subjects such as humans actually undergo is certainly an interesting one. The problem is with using this to solve the hard problem, since one wonders why one particular sort of representation or set of representations should be endowed with ‘what it is like-ness’ and not some other. Whichever physically realized representational features we propose as candidates for consciousness, it seems possible to conceive of them existing without phenomenal consciousness. My proposal avoids this problem, because it appeals to points of view to respond to the hard problem, and to cognitive capacities to solve the easy problem. This means we only need to show that the cognitive capacities we identify with consciousness are usefully identified with consciousness, not that they suffice to close an explanatory gap.

7.

The sensorimotor approach has already brought empirically supported explanatory returns when it comes to making the phenomenal character of perceptual consciousness intelligible. We have seen that, with some further argumentative steps, its account for phenomenal qualities usefully addresses not only the easy-comparative problem, but also the hard-comparative problem of explaining how phenomenal qualities can be physical at all. The advantages of the sensorimotor approach in this domain support the view that consciousness can in general be made more intelligible by understanding it as an attribute of the whole embodied subject’s way of interacting with the world rather than a kind of neural process. Hence I have attempted here to apply a sensorimotor approach to the ‘absolute’ question of why we are in the first place phenomenally conscious of what we perceive, treating the problem as one of explaining what distinguishes the conscious from the absent-minded driver.

Building on O’Regan’s claim that perceptual consciousness consists of perceiving while having cognitive access to what one perceives, I have emphasized that, in line with what ought to be construed as the sensorimotor theory’s basic principles, cognitive access should itself be identified with a readiness to engage in certain kinds of skilful

interaction, rather than with brain activity as it more commonly is. This interaction plausibly consists, in particular, of the kind of learnt and recursive communication described by Maturana and Varela, or an internally-realized extension of the same activity. This offers a robust criterion for meaningful language and thought, and so cognitive access, without needing to make appeals to neurobiology that themselves do nothing to make consciousness or intentionality more explicable. We might also decide that less sophisticated kinds of cognitively-driven interaction, when integrated with perception, are sufficient for a more basic kind of perceptual consciousness: this will depend on what we decide is the most theoretically productive way to define ‘consciousness’.

The treatment of the task as a ‘definitional’ enterprise may seem at odds with the idea that consciousness is a matter of there being ‘something it is like’, since the question of what if anything it is like to be some creature is often thought to be a question about how things are, not a matter of arriving at a useful definition. I have suggested, in response to the hard-absolute problem, that what it is like to be some creature is just what it is to cognize in the way the creature does, described from that creature’s point of view. The limitation with this proposal is that there is no guarantee that a creature’s cognition, described from its point of view, will be of a kind we would assent to describing as phenomenally conscious. To distinguish the conscious driver from the absent-minded and so plausibly non-conscious one, I suggest therefore that we appeal to the sorts of skilful interaction I have described.

Acknowledgments

This work was funded by ERC grant 323674 ‘FEEL’. Thanks to Jan Degenaar, Kevin O’Regan, Frank Schumann, and Mario Villalobos for written comments on earlier drafts of this piece, and all my colleagues on the FEEL project for extensive discussion. Thanks also to the two anonymous referees for their useful comments, and audiences in Paris (at both Sorbonne and Descartes), Helsinki, and Tucson.

References

- Berlin, B. & Kay, P. (1969) *Basic Color Terms: Their Universality and Evolution*, Berkeley, CA: University of California Press.
- Block, N. (1995) On a confusion about a function of consciousness, *Brain and Behavioral Sciences*, **18** (2), pp. 227–247.

- Block, N. (2009) Comparing the major theories of consciousness, in Gazzaniga, M. (ed.) *The Cognitive Neurosciences IV*, Cambridge, MA: MIT Press.
- Block, N. (2015) *Ned Block on Phenomenal Consciousness, Part I*, [Online], <https://scientiasalon.wordpress.com/2015/05/18/ned-block-on-phenomenal-consciousness-part-i/>
- Brown, R. (2015) The HOROR theory of phenomenal consciousness, *Philosophical Studies*, **172** (7), pp. 1783–1794.
- Carruthers, P. (1996) *Language, Thought and Consciousness: An Essay in Philosophical Psychology*, new ed., Cambridge: Cambridge University Press.
- Chalmers, D.J. (1996) *The Conscious Mind: In Search of a Fundamental Theory*, New York: Oxford University Press.
- Clark, D., Schumann, F. & Mostofsky, S.H. (2015) Mindful movement and skilled attention, *Frontiers in Human Neuroscience*, **9**, art. 297.
- Cohen, M.A. & Dennett, D.C. (2011) Consciousness cannot be separated from function, *Trends in Cognitive Sciences*, **15** (8), pp. 358–364.
- Crick, F. & Koch, C. (1990) Toward a neurobiological theory of consciousness, *Seminars in the Neurosciences*, **2**, pp. 263–275.
- Degenaar, J. & O'Regan, J.K. (2015a) Sensorimotor theory of consciousness, *Scholarpedia*, **10** (5), 4952.
- Degenaar, J. & O'Regan, J.K. (2015b) Sensorimotor theory and enactivism, *Topoi*, [Online], <https://link.springer.com/article/10.1007%2Fs11245-015-9338-z>.
- Egan, F. (2014) How to think about mental content, *Philosophical Studies*, **170** (1), pp. 115–135.
- Hurley, S. (1998) *Consciousness in Action*, Cambridge, MA: Harvard University Press.
- Hurley, S. & Noë, A. (2003) Neural plasticity and consciousness, *Biology and Philosophy*, **18** (1), pp. 131–168.
- Hutto, D.D. & Myin, E. (2013) *Radicalizing Enactivism: Basic Minds Without Content*, Cambridge, MA: MIT Press.
- Jackson, F. (1982) Epiphenomenal qualia, *Philosophical Quarterly*, **32** (127), pp. 127–136.
- Kant, I. (1998) *Critique of Pure Reason*, Guyer, P. & Wood, A.W. (trans.), Cambridge: Cambridge University Press.
- Kriegel, U. (2009) *Subjective Consciousness: A Self-Representational Theory*, Oxford: Oxford University Press.
- Loar, B. (1999) David Chalmers's The Conscious Mind, *Philosophy and Phenomenological Research*, **59** (2), pp. 465–472.
- Maturana, H.R. & Varela, F.J. (1987) *The Tree of Knowledge: The Biological Roots of Human Understanding*, 3rd revised ed., Boston, MA: Shambhala Publications Inc.
- McDowell, J. (1994) The content of perceptual experience, *The Philosophical Quarterly*, **44** (175), pp. 190–205.
- Merleau-Ponty, M. (1945/2013) *Phenomenology of Perception*, Landes, D.A. (trans.), London: Routledge.
- Millikan, R.G. (1989) Biosemantics, *Journal of Philosophy*, **86** (July), pp. 281–297.
- Myin, E. (2003) An account of color without a subject?, *Behavioral and Brain Sciences*, **26** (1), pp. 42–43.
- Myin, E. (2016) Perception as something we do, *Journal of Consciousness Studies*, **23** (5–6), pp. 80–104.

- 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: MIT Press.
- O'Regan, J.K. (1992) Solving the 'real' mysteries of visual perception: The world as an outside memory, *Canadian Journal of Psychology*, **46** (3), pp. 461–488.
- O'Regan, J.K. (2011) *Why Red Doesn't Sound Like a Bell: Understanding the Feel of Consciousness*, New York: Oxford University Press.
- O'Regan, J.K. (2016) How the sensorimotor approach to consciousness bridges both comparative and absolute explanatory gaps, and some refinements of the theory, *Journal of Consciousness Studies*, **23** (1–2), pp. 39–65.
- O'Regan, J.K. & Noë, A. (2001) A sensorimotor account of vision and visual consciousness, *Behavioral and Brain Sciences*, **24** (5), pp. 883–917.
- O'Regan, J.K., Rensink, R.A. & Clark, J.J. (1999) Change-blindness as a result of 'mudsplashes', *Nature*, **398** (6722), 34.
- Palmer, S.E. (1999) Color, consciousness, and the isomorphism constraint, *Behavioral and Brain Sciences*, **22** (6), pp. 923–943.
- Peirce, C.S. (1857–1866/1982) *Writings of Charles S. Peirce: A Chronological Edition, Volume 1: 1857–1866*, Moore, E.C. & Fisch, M.H. (eds.), Bloomington, IN: Indiana University Press.
- Penrose, R. (1994) *Shadows of the Mind: A Search for the Missing Science of Consciousness*, Oxford: Oxford University Press.
- Pettit, P. (2003) Looks as powers, *Philosophical Issues*, **13** (1), pp. 221–252.
- Philipona, D.L. & O'Regan, J.K. (2006) Color naming, unique hues, and hue cancellation predicted from singularities in reflection properties, *Visual Neuroscience*, **23** (3–4), pp. 331–339.
- Rizzolatti, G. & Craighero, L. (2010) Premotor theory of attention, *Scholarpedia*, **5** (1), 6311.
- Rowlands, M. (2016) Are animals persons?, *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, **1** (10), [Online], <http://animalstudiesrepository.org/animsent/vol1/iss10/1>.
- Rupert, R.D. (2011) Embodiment, consciousness, and the massively representational mind, *Philosophical Topics*, **39** (1), pp. 99–120.
- Ryle, G. (1949) *The Concept of Mind*, London: Hutchinson's University Library.
- Searle, J.R. (1980) Minds, brains, and programs, *Behavioral and Brain Sciences*, **3** (3), pp. 417–424.
- Seth, A.K. (2014) A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia, *Cognitive Neuroscience*, **5** (2), pp. 97–118.
- Silverman, D. (2017) Internal representation and bodily skill in sensorimotor perception, *Phenomenology and the Cognitive Sciences*, [Online], <https://link.springer.com/article/10.1007/s11097-017-9503-5>.
- Sprevak, M. (2013) Fictionalism about neural representations, *The Monist*, **96** (4), pp. 539–560.
- Thompson, E. (2007) *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*, Cambridge, MA: Harvard University Press.
- Tye, M. (2000) *Consciousness, Color, and Content*, vol. 113, Cambridge, MA: MIT Press.
- Villalobos, M. (2014) *The Biological Roots of Cognition and the Social Origins of Mind*, PhD thesis, University of Edinburgh.

Witzel, C., Cinotti, F. & O'Regan, J.K. (2015) What determines the relationship between color naming, unique hues, and sensory singularities: Illuminations, surfaces, or photoreceptors?, *Journal of Vision*, **15** (8), 19.

Paper received October 2016; revised February 2017.